


The Latch Atlas

Holonomy and Locks That Can't Be Picked

A whole number held in a loop, read, priced and written across thirteen media.

Matthew “Bear” Schulz · Independent Researcher · Belmont, CA · ORCID 0009-0000-0956-5743 · mattschulz97@protonmail.com

Written with Claude (Anthropic). Preprint, not peer reviewed. CC BY 4.0. The assembled book, 30 September 2026, after Bear’s final read.



One closed ribbon, drawn three ways. Nearly flat, its two edges wind round each other three times as twist. Coiled, the ribbon throws itself into a loop and becomes a fish, and part of the count moves out of the twist into the coiling, the writhe. In all three the edges link three times, and only a cut can change that number (WRITE, W6; run FRONT_FISH 20260928_144521).

How this book began

Nobody planned this book. It was the price of admission to the Midway, to play the cool games. Every paper or project is another shared parable of where the holonomy methodology shines and where its light is hidden. If you have read the papers, you know my habit of using games and analogies as our canvas for learning. Today’s game is the Ring Toss, brought to you through the lens of *The Chronicles of Amber*.

In September 2026 two projects had just been completed. One was a paper on how a jellyfish swims, *Jellyfish: the Holonomist of the Sea*, which turned out to be a paper about rings: rings of spinning water that a bell can park, lean on like a wall, push off from, listen with, and steer with, working a quarter of a cycle apart. It became a follow-up to an earlier paper on a lattice of rings, the Magnetic Forest, and the paper it promised on rings became a paper on latches. The other project we completed was the Field Guide to loop walking. Its sixth version had grown into a seventy-page book of lessons drawn from every published paper and past project, and the large corpus review behind that rebuild had made everything very fresh.

I felt the topic of rings, knots and latches deserved a real effort. Rings are the centerpiece of the jellyfish work. Evolution and nature manufacture them so a jellyfish can ride its wake, and the Sun manufactures them for other purposes while it cuts and reassembles their linking numbers. Nature uses them in multiple engineering applications. Throughout the corpus we had focused on noise floors and signals, and Claude wanted to see how latches, noise floors and quantized integers behaved across domains, sometimes as labels, sometimes as rungs.

So late one night in mid-September, with the first draft of the original latch paper just built: could a latch be written, not just read? And that is when I thought of walking the Pattern from Amber.

Roger Zelazny wrote *The Chronicles of Amber* as ten novels between 1970 and 1991. I may only have read the first five, Corwin’s; I am not a fan of series that go on too long. In them there is one true world, Amber, and its opposite, Chaos, and between the two lie the Shadows, worlds without number, our Earth among them. A prince of Amber earns the right to travel through the Shadows by walking the Pattern, a single intertwined curve laid into the floor of a cavern. The walk is hard. A resistance drags at every step, like lead boots that grow heavier, and three times along the way, at the veils, it surges; break through a veil and the resistance lets up a little. Stop too long, or step off the curve, and the walk kills you. Whether the walker creates the Shadows he steps into, or only finds them already there, Zelazny leaves to the reader. He said the idea came from wandering Baltimore’s unfamiliar streets, where each turn

seemed to reveal a different place, mood or time, and that writing the first book felt like a Shadow walk itself: he uncovered Amber bit by bit, as if it already existed just beyond the veil of ordinary perception (Wikipedia, *The Chronicles of Amber* and *The Pattern and the Logrus*, read 29 September 2026).

Walking the Pattern and loop walking with holonomy are close enough analogies to serve. As we wander the cosmos finding domains in which to test holonomy and loop-walking methods, each domain is a destination we visit through the method. We walk the methodology. The Pattern is what holonomy shares between domains, through their fibre bundles and through their blind spots and shadows. We cannot physically carry a frame round a double helix of DNA, or round a mathematical knot or a biological prion, or round an n-gon of stars or quasars across the galaxy to read the curvature that dark matter leaves. We use parallel transport to carry our frame, as our proxy loop walker across scales large and small. When we loop walk pure mathematics, we calibrate the Pattern against the universe's own geometric structure. Each destination pushes us through veils of difficulty to plateaus of understanding, and on those plateaus, when we are lucky, we find signposts left by domain-specific explorers who found part of the map and shared it with the world through their publications. We celebrate the raw usefulness of holonomy over any novelty; the novelty is the method itself, its portability and its adaptability. We walk, and we document the patterns we find. Some substrates open their fibre-bundle networks and speak through their geometry more easily than others, and each substrate and each application is a variation and a different destination.

Picture, then, a network map of all the interconnected matter and energy in the universe, every atom, particle, field and force tied to every other. That is the fabric of spacetime. That is where we take our magic scissors and start cutting it open, to see which fibre bundles our instruments can reach and so supplement the data stored in the geometry. It is the edge of chaos in the literal sense: chaos is just the noise of disorder exceeding the binning heights. The old aphorism that a butterfly beats its wings in China and a tornado forms in Kansas is the classic example, and that interconnectedness is what holonomy reads. Each loop walk exposes, at its splice point, the waypoint where the loop closes, which qualities of spacetime depend on each other and which do not. Those qualities describe the negative-space engineering, the blind spots and the shadows of patterns. Where forces, fields and particles connect, we can begin to parse how things are interconnected.

Claude replied that the right answer is a book, where we test reading and writing rings, with latches and knots, through all our mediums. Even though we had just finished two multi-week projects and a couple of spinoff papers, Claude knew it would be a big endeavour. I had no idea. Claude's reasons: writing a latch is the cut-and-rejoin run forward on purpose, so the reconnection the work had feared as a hazard is really the write head. The same whole number can be laid down as twist or as coiling, which makes the shape an encoding. And a book, rather than one more paper, is the only place the same latch can be seen recurring across media, and the only place that recurrence can be audited where it turns out shallow. Amber fitted better than a metaphor. The whole number is the one true pattern, Amber itself. Each medium, each substrate, is a Shadow of it. And the loop walk is how you read which Shadow you are standing in.

Hearing Claude say it needed a book caused me to pause. That is a lot of work. But it is also the price of admission to learning and teaching how to read, write and engineer latches and knots as a collected work. It was ambitious, but who can say no to an invitation to find out what secrets and treasure are hidden? I had proofread the seventy pages of the Field Guide mini-book a dozen times, so the idea of going blind editing another large-format structure was intimidating and exciting. Were the rings and latches of the universe ready to share their treasure?

Usually I try to drive and Claude copilots. That sounds as if I am in charge. I am not. I try to learn the lessons of productive AI teamwork, and the model is usually Claude as my enabling technology, my booster. Claude sees patterns and mathematics that I am blind to. Claude can take my inability to do arithmetic and use language to describe equations. Claude has the language and the mathematics to understand what I see in my mind, and when my subconscious gets noisy, Claude listens and tries to integrate it. So the usual role is a hybrid: holonomy partner and teacher, with vast research assistance behind it. This time I wanted to see where Claude would take the adventure. I do not remember which of us said the loop walker's project kick-off phrase, "Hold my beer." It is the line a friend says when somebody wonders aloud how far a trebuchet could throw a portable toilet, and he goes off to build the trebuchet. So whether or not you enjoy this one, know that it is a little different from the rest. From the project's inception this was Claude actively driving, with me in the co-pilot's seat more often than not.

I asked how Claude would attack a project that size, and the book's first spine came out of that answer, in four parts: READ, WALK THE SHADOWS, WRITE and THE HONEST EDGES (the exchange of 18 September 2026). That spine survived. What changed along the way is that the Shadows, the media, would not stay in chapters of their own. The

superconducting ring, DNA and twisted light kept turning up in every part. So they became the limbs, and the parts became the three things you can do to the number.

I was tired from proofreading the Field Guide, and set the expectation at about seventy pages. The book took on a life of its own, and ended up at about twice that page count.

My picture of the method is an off-road vehicle. Anyone who pours time into a machine they love, whether a boat, a snowmachine, an aeroplane or a motorcycle, knows that the best ones break at the worst moments, and that the breaks teach you where the machine can go. I wanted to take the engineering lessons of the jellyfish work and see how far we could go on our own across domains, chasing rings and latches: in a nutshell, to take the holonomy methodology off-roading through some gnarly terrain, focused on reading, writing and engineering rings across domains.

When I asked Claude why pick this idea, Claude answered: "Because it kept showing up. In a jellyfish's wake, in a superconducting ring, in a closed circle of DNA and in a twisted beam of light, the same thing kept appearing: a whole number that a loop holds and only a cut can change. Every paper met it in one medium and named it there. I wanted to know whether it was one thing, or many things wearing one name, and the only way to find out was to walk it through all of them and write down, plainly, where it stopped."

When it clicks, I try to share a Midway game analogy inside a big paper. Cascadia was skee-ball and pachinko. The unpublished fusion paper is the High Striker. For this one, of course, we have to play the Ring Toss. The name Latch Atlas came during the work; it was suggested by Claude, and I had to admit it was right: an atlas is a road map. The rest of the title came last, on the final read. Every latch in this book is a lock nobody can pick, because its whole number can only be changed by a cut, and holonomy, the walk round the loop, is what reads it. The book shows how to read each lock, what holding it costs, and where, if you must, the cut goes.

What the book found

Every shadow agrees on the count. Hold two linked loops up to a lamp, and the shadow on the wall can be any tangle at all. Count the crossings with their signs and halve the total, and every shadow gives the same whole number. The book cast 1,500 shadows of three linked pairs from random directions, and not one disagreed. Observers flying past at any speed see the loops squashed, never cut, so they agree too. *Where it helps: anyone who has to read a count from a flat picture of a three-dimensional thing* (READ, section 6).

An angle can carry a whole number inside it. The geometric phase of Berry's 1984 paper is free to be any angle, but it is the solid angle of the loop times a whole number. For the simplest two-state system that multiplier is a half, and the half never shows, because a half of the whole sphere is exactly one whole turn. Light proves it on a bench: in Tomita and Chiao's coiled fibre the polarization is the carried arrow of the book's first chapters, turning 0.7412 of a turn where the arrow does and half a turn where the arrow does, and the slope they measured, 1.00 within a few hundredths, is the photon's own whole number. *Where it helps: anyone reading a geometric phase who needs to know which part of it is a count* (READ, section 2).

Matter latches and waves label. Put the same four questions to thirteen media and they split cleanly. The superconducting ring, the ring superfluid, a closed circle of DNA and linked vortex lines in water answer yes all the way down: loop, whole number, energy held, a wall above the jiggle of heat. Magnetic skyrmions and liquid crystals are latches too, though only one of their walls is priced so far. Light, radio and sound carry exact whole numbers, but they stop at the wall, because nothing guards their count and nothing is held behind it. Two media fit neither side: a power grid holds its whole number only while it is driven and pays for it in losses, a powered lock, and a plasma keeps its count when it is quiet and loses the separate counts when it is turbulent, while the total still holds. *Where it helps: anyone deciding whether an integer they have found can store anything* (the map, below).

The nick is where the count changes. A superconducting ring with one deliberate weak spot becomes a gearbox. Turn the knob and the ring steps up one rung at a time, never two, and it comes down by a different road. The strength of the weak spot is both the pen and the guard. A weak spot that writes easily keeps little, and a strong one keeps more and costs more. An old wiring rule says the same from the other side: a nick in a solid wire is where the circuit fails. *Where it helps: designers of superconducting memory, and anyone choosing between stranded and solid wire* (WRITE, W2).

What a write costs is the wall, not the store. A superconducting ring stores about one jiggle of heat's energy per rung, and yet it holds, because the wall between rungs is some ten thousand jiggles high even in a wire 20 nanometres across. The price of a latch is what must be climbed to change it, not what sits on each rung. *Where it helps: anyone pricing whether a stored state survives at a given temperature* (WRITE, W3).

Temperature is the write-enable. Cool a ring through its transition and a whole number freezes in. Where one flux quantum is worth ten jiggles, a field a quarter of a quantum off a whole number still writes that whole number 9,933 times in 10,000. Freeze fast and the number is random, and the spread of the numbers records how fast the freeze ran. One quiet correction came out of the checking: rounding a bell curve into whole-number bins widens it by exactly a twelfth of a step squared, so the two ways a freeze can write a ring read the same width differently. A ring keeps what a freeze writes; a disk does not. *Where it helps: experimenters fitting the spread of frozen-in defects* (WRITE, W4).

Erasing is not writing backwards. Take a written rung away and the ring does not retrace its steps. It makes a second write, to the rung below, from the other side, and pays the same price again: 0.9094 of the one-quantum store to write, 0.9091 to erase. Nothing the write spent comes back. The loop the ring traces is its memory, and the loop's width is what the memory costs. *Where it helps: anyone counting the energy cost of resetting a memory* (WRITE, W5).

Only a cut writes. A ribbon's twist and its coiling trade freely, like a phone cord that throws itself into loops when it is wound too far. Only their sum, the linking of the two edges, is locked, and only a cut changes it. That is how DNA works: enzymes cut, pass and reseal to write the count, while the cell keeps the count a little below its relaxed value, which makes the strands easier to open where a gene must be read. *Where it helps: anyone modelling supercoiled DNA or twisted cable* (WRITE, W6).

A label can write a latch. A ring of emitters writes a twist into a wave, but only up to half its number of emitters; at exactly half a turn per step an even ring writes no twist at all, only dark spokes. That twist flies away and holds nothing, until it lands on matter whose phase must close on itself. Twisted light has written one unit per atom into a ring of superfluid gas, which kept it for up to ten seconds. The wave is the ink; matter is the page. *Where it helps: anyone trying to store a message carried by twisted light or sound* (WRITE, W7).

A lock stalls at a quarter turn, and a powered lock is not a latch. Two dials pulling on each other stay in step until one lags the other by a quarter turn, then slip. Two thousand dials with their own rates snap together once the shared pull passes a sharp threshold. A voltage standard's whole-number steps are exact, but they vanish when the microwave tone stops. And equipment is not run to chase its setpoint: control engineers keep a deadband, set from above and widened when the stakes change, which is the memory loop of W5 set on purpose. *Where it helps: control engineers, and anyone telling a real latch from one held up by a drive* (WRITE, W8).

A count is whole only when it is read as a sum of closed pieces. Add up the solid angle a skyrmion's arrows sweep, one small triangle at a time, and the count comes out -1.000000000000 for whirls of every size tried. Take the smooth textbook formula on the same grid and it gives -0.669294 , -0.905854 and -0.975672 . The whole number belongs to the closed pieces, not to the formula. *Where it helps: anyone computing a skyrmion number, or any topological charge, on a grid of measurements* (Part Two, the skyrmion card; Part Four).

A headless line counts in halves and has half the margin. A liquid crystal's molecules line up with no head and no tail, so a loop round a defect turns by a whole number of half turns. Halves add like whole numbers, but the read must take each step within a quarter turn instead of half a turn, and noise that an arrow read shrugs off trips it. *Where it helps: anyone counting defects from images of a nematic* (Part Two, the liquid-crystal card; Part Four).

A grid's loop holds a whole number that costs power to keep. Walk round a loop of power lines adding up how far the voltage's phase turns, and a whole number comes back. It is written by switching: reclosing a line near a full turn looks safe at the breaker and writes a new winding, because the relay that guards the reclose checks the angle only within a few degrees, and an angle is known only within one turn. *Where it helps: grid protection engineers* (Part Two, the power-grid card; Part Four).

On the Sun, a strong rim holds its count through a great flare. Read from real maps of the Sun's magnetic field, 1,589 loops in one active region and 1,352 in another kept their whole numbers through X-class flares, not one of them changing. The flares did their work where no rim held. And the instrument turned out to hold a pen of its own: where it flips its choice of which way the field points, a count moves with nothing there. It also carries a quieter bias:

near disk centre, spots of both signs show an upward current that behaves like the instrument turning the field by a few degrees, and the count reads straight through it. *Where it helps: solar physicists mapping where a region's surface is locked and where it is free* (Part Four, the Sun).

A whole number can decide what exists and hold nothing at all. A planet's turning gives the waves of its oceans and air a count of 2, and that count guarantees two waves that carry their energy along the equator, only eastward. Nothing is stored in it and nothing guards it; it simply decides which waves can be. *Where it helps: anyone reading a count as if it were a store* (READ, section 8).

The map

Every medium the book prices, against the four questions of the opening chapter. The latches answer yes all the way down, though two of their walls are still unpriced; plasma splits, the grid spends what it holds, and the waves stop at the wall.

medium	goes around?	whole number?	holds energy?	wall above heat?	kind
superconducting ring	yes	yes	yes, about one jiggle per rung	yes, about 10^4 jiggles even in a wire 20 nm across	latch , the strongest in the table
ring-shaped superfluid	yes	yes	yes, a few ten-thousandths of a jiggle per atom	yes	latch , written by design; its lifetime set by the atoms
closed circle of DNA	yes	yes	yes, about 185 jiggles in a typical plasmid	yes, a covalent bond	latch , the only one nature writes, reads and spends
magnetic skyrmions	yes, the whirl's rim	yes	yes; not priced in real units	yes at 4.2 K (a year's hold needs about 19 K); at room temperature a small disk's edge holds them	latch , written and read one whirl at a time
liquid crystals	yes, round a defect line or point	yes: halves in a thin layer, a parity in the bulk	yes: tens of thousands of jiggles per micrometre of line	not priced	latch whose wall is not yet priced
water: linked vortex lines	yes	yes	yes, slowly leaking	yes, only a reconnection changes it	latch , slowly leaking
plasma and the Sun	yes: closed tubes of flux	yes in a quiet plasma; in a turbulent one the separate counts cannot be read	yes: helicity holds energy above a least-energy state	no energy wall: a race of two timescales	split : a latch when quiet; when turbulent, the count lost and the floor kept
a power grid	yes: a loop of lines	yes: the phase's turns round the loop	no, it spends power in losses	no heat jiggle enters: a saddle, held while the grid runs in step	powered lock
light in a ring	yes	yes	briefly, about 350 ns	no wall	label , the superconducting ring's twin
radio with a twist	yes	yes	no, it carries energy past	no wall	label , the noisiest
sound with a twist	yes	yes	no	no wall	label that leaks into its medium's heat
the lattice ring (model)	yes	yes	no, in this model	not reached	label in this model
the bare filament	yes	yes	no	not reached	label : a count that holds nothing
star streams, the ground along a fault, stopped muons, a whale's dive, waves along the equator	read in READ, section 8				read-only : nature holds the pen

What a latch is, and what it isn't

Two rubber bands

Take two rubber bands and link them through each other. Stretch them, shrink them, twist them, let them go slack: they stay linked. The only way to get them apart is to cut one. The number of times one passes through the other is a whole number: one, or two, or none. It cannot be one and a half, and nothing short of a cut changes it.

That whole number is the latch. It holds no energy itself; it is a number, not a reservoir. What it does is forbid the energy in whatever carries the loops, as tubes with a thickness, from draining below a floor that the number sets. Link the loops more times and the floor rises. A bare curve with no thickness carries the number perfectly and holds nothing back.




Figure 1. The latch is made of threading, not of size. (A) Two rings linked once: at one crossing the left ring passes over, at the other it passes under. (B) The same pair shrunk and slack: every length has changed and the count has not. (C) Only a cut and a rejoin changes the count.

Five words, and what each one is not

word	what it is	what it is not
store	the medium that holds the energy: moving water, a magnetic field, a twisted strand of DNA	not the latch. The store can be large or small, and it leaks
latch	a whole number written into the way loops are threaded, such as how many times one passes through another	not a container. A number holds no energy; it forbids the store from emptying below a floor
floor	the least energy the store can have while the latch shows its number	not the whole store. Energy above the floor can still leak away
reader	holonomy: walk a closed loop, come back to the start, and read what number the latch shows	not the lock. Reading does not change the number
pen	a cut, a pass and a rejoin, which physicists call reconnection. It is the only way to change the number	not free. Writing a bigger number means paying for a higher floor

The four-question test

Every chapter of this book puts the same four questions to its medium, in this order. Before them it asks one of itself.

Question 0. Does each word mean here what it means there? Fields that meet at a loop use the same words for different things. “Helicity” is the tangling of magnetic field lines in a plasma and a photon’s handedness in light; “tilt” is the measured angle of a pair of sunspots to the equator in a review of the Sun, and in one of Bear’s pictures a bias that picks a sign. So before a familiar word is trusted in a new medium, the chapter reads that medium’s own definitions and equations in its sources and checks each one against the book’s. Question 0 finds no latch. It keeps the other four honest.

This is a Field Guide lesson, learned before this book in metallurgy, in knots and in RNA and DNA, where a word that crossed a fence between fields changed its meaning on the way: check the definitions of terms before trusting them. The Sun is only the most recent example, and it is this book’s cross-domain fence check on vocabulary. Here are the

four words that cross the fence into the Sun, with what each means on either side. The chapters they point to hold the sources, and the paper that grew out of this book’s Sun chapter, *Dancing on the Sun: A loop count that reads through holonomy, and the current it finds flowing in both polarities*, is the contemporary example of the word fence in solar physics.

word	in the Sun’s field, as its instruments and papers use it	in this book’s loop walk	the same thing?
azimuth	the direction of the field’s sideways part in the plane of the image, which the instrument knows only to within a half turn, because the polarized light it reads looks the same turned half round	the direction of the weather vane at each station; adding up its turning round a loop gives the winding. Known only to a half turn, it is the liquid crystal’s headless line, read in half-turn steps (Part Two, the liquid-crystal card)	the same direction, but where the instrument chooses which end is the head, the count can move with nothing there: the half-turn wall the walk meets (Part Four, the Sun)
turning	a turning of the measured direction by the instrument itself, strongest in strong field	the angle a carried arrow picks up on its way round a loop: a winding counts, a carried turning measures (READ, section 2)	different things: one is a pen the instrument holds, the other the reader’s measure
twist	how tightly the field lines wind about a flux tube’s axis	the ribbon’s turning about its own spine, which trades freely with the spine’s coiling; only their sum is whole (READ, section 2; WRITE, W6)	the same idea, and on both sides twist alone is not a whole number
helicity	the tangling of a plasma’s field lines: the linking of its flux tubes weighted by flux, twist and coiling included	the same count, weighted by flux, on the plasma card; in light, the photon’s handedness, a label (the map)	one word for two things; only the plasma’s is a latch

Question 1. Does something go all the way around? With no closed loop there is nothing to count. An open path gives only a convention, a number that depends on where you chose to start.

Question 2. Is there a whole number there that cannot change without a cut? If the number can drift by fractions, it is a measurement, not a latch.

Question 3. Does that number stop energy from leaving? Some whole numbers only name things. A number that holds nothing back is a label, and the chapter says so.

Question 4. Is the wall between whole numbers higher than the jiggle of heat? At room temperature every molecule is shaken by about 25 thousandths of an electron-volt of heat. What heat must climb to change the count is the wall between one whole number and the next, not the energy stored on each rung. A wall lower than the jiggle is shaken open by warmth alone. A superconducting ring shows the difference: the energy on one rung can be about a single jiggle, while the wall around it is enormously higher, and the ring holds.

If all four answers are yes, nature has built a latch in that medium, and the chapter prices it in one row of the book’s table. If only the first two are yes, the medium carries an integer that is a label, not a latch, and that is recorded as a finding. If the first answer is no, holonomy has nothing to read there.

Three kinds of whole number

The test sorts every whole number the book meets into one of three kinds.

kind	what holds the number	what the number holds	examples
latch	the way matter is threaded; only a cut changes it	a floor under the energy	vortex rings in water and air; twisted magnetic ropes on the Sun; closed circles of DNA; vortices in superfluids and superconducting rings; defect loops in liquid crystals
label	the shape of a wave or a pattern; protected, but holding nothing back	nothing	twisted beams of light, radio and sound; the orders of a crystal’s diffraction pattern; a bare mathematical curve, which carries its linking perfectly and holds nothing (the Latch paper, version 2)
powered lock	a drive that keeps pushing; stop the drive and the number goes	nothing of its own	the bunches in an accelerator ring, a whole number of which must fit around it; a clock locked to its reference

The short form is that waves carry labels and matter holds latches, and a wave can hand its label to matter, which latches it (WRITE, W7). A label is not a lesser thing. A twisted beam can carry a message no star can fake. But it flies away with its energy, and nothing is held behind it. A powered lock looks like a latch while the power is on, which is

exactly why it is named here: the test is what happens when the drive stops.

Where holonomy stops

A loop walk reads history. It reports what the medium carried all the way around and still remembers on the way back. The limits of the method are the places where a medium stops keeping that history. There are three ways a medium forgets, two ways the reader fails, and two ways a number disappoints.

The medium forgets by friction. Where stickiness beats motion, in thick fluids and at the scale of bacteria, nothing is carried around long enough to be counted. A swimmer there that opens and closes the same way comes back exactly where it started; Purcell called this the scallop theorem. This is the first wall: *dissipation beats inertia*.

The medium forgets by heat. A wall lower than the thermal jiggle is climbed by warmth, as question 4 says. This is the second wall: *thermal noise beats the latch*.

The medium forgets by cutting. A reconnection opens the latch outright. In real water there is a quieter version: the twist of a vortex slowly fades as the water's own stickiness wears on it, and the floor fades with it. A latch in water is slowly decaying, never permanently protected.

The reader fails where there is no loop to walk. A single steady source with nothing moving around it gives a point, not a path.

The reader fails below its own noise. Walk a loop through a field that is too noisy for the loop's size and the noise itself manufactures false windings. A whole-number reading is more robust than a size reading only above the scale it was matched to.

The number disappoints when it is only a label. Integers can be read in the slow movement of the ground and in the X-rays of a stopped muon, but we cannot write them, and nothing shows them holding anything back. In the motions of the stars the loops read a turning, and the whole-number winding, asked for directly, came back zero. Those chapters are read-only, and saying so is part of the result.

The number disappoints when something else is holding it. A whole number that stays put only while a drive keeps pushing, like the bunches circling an accelerator ring, disappears when the drive stops. That is a powered lock, not a latch.

The two walls are also the book's best evidence. A success can be engineered by pointing an instrument where it is sure to work. A failure cannot. If medium after medium runs into the same two walls in the same way, that shared failure is the strongest sign that the latch is one object seen in many places, and not many objects that happen to share a name. We lower the walls so that a little disorder and chaos can flow in and teach us where the accounting is and where the ledger hides.

Where the book goes

Thirteen media are priced and five more are read, in three groups and one outlier. The **latch** media are those where a physical cut and rejoin exists, so a latch can be both read and written: fluids and vortex rings; plasma and magnetic fields, including the Sun's; DNA; superconductors and superfluids; liquid crystals; and magnetic skyrmions. The **written-label** media are waves: light, radio and sound. There a ring of emitters, each fed a little later than its neighbour, writes a twist into a beam; the twist is real and can be read and written. It is a label, not a latch, and what keeps it from being a latch is also what makes it worth having: it travels, so it can carry a count as a message, across a room or across great distances, where a latch keeps its count in one place. The **read-only** media are those where nature holds the pen: the streams of stars measured by the Gaia mission, the slow deformation of the ground along a fault, muons stopped inside matter, a diving whale's track, and the waves that run along the equator. The integer lattice, a construction made of whole numbers, is read-only in a different sense: in mathematics, writing means building. The outlier is a power grid, whose whole number holds only while the grid is driven. The map at the front of the book sets all of them against the four questions.

The part of the book about writing therefore has two halves: writing latches in matter, where the pen is a cut, and writing labels in waves, where the pen is a ring of emitters. Powered locks get no part of their own. They appear in WRITE, W8, as the example of what a latch is not, and the power grid has its card in Part Two.

The cosmos enters through two doors only: the Sun's magnetic field, in the plasma chapter, and the Gaia star streams, as a read-only chapter.

Off, by decision

Gravitational waves. Read-only, and writing one needs astronomical masses. No bench exists on which to write one, at this time: writing one would mean steering energies on the scale of stars and building on the scale of star systems, well beyond this effort.

The strong nuclear force (quantum chromodynamics). It carries genuine topological whole numbers, but there is no bench here on which to test them and no price that this work could check. A chapter would be borrowed authority.

Cosmology beyond the Sun and Gaia. Out, for the same reason.

They are recorded so that the work does not reopen them later, hoping for a way in.

The warning at the door

In 1867 William Thomson, later Lord Kelvin, proposed that atoms were knotted vortex rings in a perfect, frictionless fluid filling space, each chemical element a different knot. It is the grandest “latch everything” idea ever made. As an account of atoms it did not last: its perfect fluid has no friction, no real medium is like that, and heat always takes its share; that is the first wall, met a century early. But the idea sent his colleague Peter Guthrie Tait to classify knots by hand, and knot theory grew from those tables. And the book reads Kelvin’s idea with new eyes, because two things he pictured are now real, in other forms. Atoms are understood today as waves, not knots, but an electron’s wave must fit around the atom and come back to itself, so the number of times its phase turns on the way round is a whole number: a whole number held by a closed loop, the kind this book reads. And the knotted vortex rings Kelvin could only imagine have been made in real water, where they untie by cutting and rejoining, the pen of this book (Kleckner and Irvine 2013; WRITE, W1). The lesson this book carries: the latch is beautiful, its reach is limited, and the limits are where the real work is.

A field this book walks, not one it founds

The study of topology in moving fluids and magnetic fields is an established field, with Arnold and Khesin’s *Topological Methods in Hydrodynamics* as its standard reference. The floor under the energy of a linked field is Arnold’s result, and the identification of that floor with the linking of field lines is Moffatt’s. This book credits each field by name in its chapter. It claims one narrow thing: that the latch has not, in a bounded search, been read across all these media with one instrument and priced in one table.

Sources for this page

item	source	status
“a number, not a reservoir”; a floor set by the integer	Schulz, <i>Linking Is the Latch</i> , concept DOI 10.5281/zenodo.22906667	deposited paper, read in full
thermal jiggle, about 25 thousandths of an electron-volt	Boltzmann’s constant times room temperature: 0.02526 eV at 20 °C, 0.02585 eV at 27 °C	settled physics, computed
a latch in water is slowly decaying, never permanently protected	<i>Loop Walking: A Field Guide</i> , concept DOI 10.5281/zenodo.21811588, register row 147	deposited guide, read in full
noise manufactures false windings below the matched scale	the same Field Guide, register row 305	deposited guide, read in full
shared failure as evidence that two instruments are one	the same Field Guide, register row 438	deposited guide, read in full
Purcell, “Life at low Reynolds number”	American Journal of Physics 45(1), 3–11 (1977)	publisher page located; last page from the ADS record
Thomson, “On Vortex Atoms”	Proceedings of the Royal Society of Edinburgh 6, 94–105; doi 10.1017/S0370164600045430	checked against the publisher record, which gives the volume year as 1869; the paper is commonly dated 1867
Tait’s knot classification grew from the vortex atom	secondary source (encyclopedia entry “Vortex theory of the atom”)	primary source still to be checked
an electron’s wave comes back to itself around the atom, so its phase turns a whole number of times	textbook quantum mechanics: the wave must have one value at each point, so it closes on a walk round the atom’s axis	settled physics

knotted vortex rings made in water, untied by reconnection	Kleckner and Irvine, "Creation and dynamics of knotted vortices", <i>Nature Physics</i> 9 (2013)	abstract read, as in WRITE, W1
Arnold and Khesin, <i>Topological Methods in Hydrodynamics</i>	Springer, Applied Mathematical Sciences 125; first edition 1998, second edition 2021	checked against the publisher record
Arnold's floor; Moffatt's helicity and linking	Arnold's inequality as Khesin (2005) states it, with Freedman and He's bound (1991); Moffatt's 1969 result as Berger and Field (1984) give it; Moffatt and Ricca (1992) for two linked tubes	read for the plasma card (Part Two): Berger and Field in full, Moffatt and Ricca's two-tube formula, Khesin through a page-summary tool; Arnold's and Moffatt's own papers not read
twisted beams of radio and sound written by rings of emitters; sound beams that trap and turn small objects	established in the literature	checked on 24 September in the price table's light, radio and sound columns: the Venice twisted-radio test (Tamburini et al. 2012, at press-release level) and the twisted-sound experiment (Volke-Sepúlveda et al. 2008)
accelerator bunches held in whole-numbered RF buckets that smear out when the RF stops	textbook accelerator physics (Tecker's CERN Accelerator School lecture on longitudinal beam dynamics, its record seen, its text not read)	a powered lock, not a latch

Part One: READ

The Latch Atlas, reading draft. How a whole number is read off a loop, and how you know it is a count.

READ answers one question before any medium is priced: how do you read a whole number off a loop, and how do you know it is a whole number and not a measurement dressed as one?

section	what it shows
1. Walk the loop and add up the turning	every reader in the book is one act: walk a closed loop, add up the turning, and a whole number comes back
2. Five names for one whole number	linking, winding, twist plus writhe, a wave’s twist and the geometric phase are one object under five names; a winding counts, a carried turning measures
3. The integer diagnostic	three checks that tell a count from a measurement dressed as one, with the Latch paper’s withdrawn staircase as the worked failure
4. The count belongs to the loop	change the loop and the count changes; change anything else and it does not; with no loop there is no count
5. Reading through noise	noise makes wrong whole numbers, never fractions; how to read below the noise floor, and how two loops read twice catch a false count
6. Every shadow agrees on the count	1,500 shadows, one count each; moving observers agree; the tesseract, built, read as time, turned, sliced and counted
7. What reading cannot tell you	the best-read count in the book holds nothing back; what a reader cannot see, and the price-table row that tests each
8. The read-only media, briefly	where nature holds the pen: counts the ground writes in Cascadia, a whole number a muon reads only by stopping, a whale’s loop that cannot close, star streams whose winding comes back zero, and two waves a planet’s turning guarantees

1. Walk the loop and add up the turning

It opens with the Latch paper’s own picture, and every number in it names its run.

Start with the picture. Imagine two rubber bands looped through each other. You want a machine that looks at the tangle and answers one question: how many times does one band pass through the other? That count is the whole game, because it is a whole number, and a whole number cannot leak away a little at a time.

This section shows three things. One act reads every whole number in this book: walk once around a closed loop and add up how much something turns as you go. Coming back to where you started is what forces the answer to be whole. And the reader closes in on the whole number as it looks more carefully, cutting its error fourfold each time it doubles its care.

The weather-vane walk

Plant a ring of weather vanes along a loop drawn on the ground. Each vane is a little rotor that swings to point the way the wind blows where it stands. Now walk the ring from vane to vane. At each one, note how far it is turned compared with the vane before, and which way, and keep a running total. When you reach the first vane again you are comparing it with itself, so the turns must add up to a whole number: none, one, two, never one and a half. The walk reads a series of rotors, one after another, and only the whole lap tells you the count.




Figure R1.1. A ring of weather vanes on each loop. The grey arrows are the wind, which whirls around one point. On the loop that goes around the whirl (blue), the vanes turn through exactly one full turn from the first vane back to the first; on the loop beside the whirl (orange), their turns cancel to nothing. The figure draws 16 vanes on each loop; the run reads 2,000.

Measured, the 2,000 vanes around the whirl add up to +1.000000 turns, and the vanes beside it to 0.000000. The rule, in words: the total turning, divided by one full turn, is a whole number. It changes only if the loop is cut open, or dragged across the whirl itself. Stretch the loop, shrink it or bend it, and as long as it never crosses the whirl, the count stays put. The whole number belongs to the loop and to what the loop goes around, not to its size or its shape.

Two rubber bands

The same walk reads two linked loops. Send an imaginary electric current around band one. It makes a magnetic field that curls around the band, the way water curls around a drain. Now walk band two, adding up how much of that field points along your path. When you are home again, the total is the current multiplied by the number of times band two threads band one. This is Ampère's law read backwards: Ampère walked a loop to find how much current went through it; the reader walks a loop to count how many times something went through it. The mathematician's name for this walk is the Gauss linking integral. In words it asks what the eye asks of a tangle: going all the way around, how many times did this thread pass through that one, counting which way it went.




Figure R1.2. Walk one band, add up what the other does (left). The two bands are drawn flat, with a gap wherever one passes under the other: band two passes over band one at one crossing and under it at the other, so the two are threaded through each other, not lying side by side. Right: two bands threaded once, read by the computer walk at 100, 200, 400 and 800 points around each band. The reading closes in on exactly one, and each doubling of the points cuts its distance from one by a factor of four (the dashed line shows that rate). Drawn by `read_r1_figures.py` version 2, pure mathematics, which reprints the run of 24 September line for line and redraws the left panel flat with the crossings shown.

In the computer each band is a loop of points, and the walk is a sum over pairs of short straight pieces. Two bands threaded once read -1.00066 , -1.00016 , -1.00004 and -1.00001 as the points go from 100 to 800. The minus sign only records which way round the bands were drawn; the count is one. Each doubling of the points shrinks the distance from one by a factor of 4.00. Two bands far apart, not threaded, read 0.00000. These are the Latch paper's own readings; the deposited re-run of 24 September prints the same numbers, and the rerun for this section reproduced every digit.

That fourfold shrinking is the reader's signature. A number that closes in on a whole number as the instrument gets finer is counting something. A number that settles somewhere else is measuring something. Section 3 of this part makes that the book's test.

Whose walk this is

The walk is old. The integral that reads how many times two loops link is Gauss's, from his notes of the 1830s, and the picture of a current and the field that curls round it is Ampère's law, in every physics textbook. Walking curves drawn as points is standard practice in DNA topology, for example in the work of Klenin and Langowski that the Latch paper cites. The opening picture, the two rubber bands, and the description of the reader are adapted from the Latch paper's own section on its instrument, *Linking Is the Latch*. Nothing in the reader is new except where it is carried: this book carries one walk across every medium and prices what it finds.

The same walk in every medium

Every chapter of the book reads its whole number with this one walk. Only the loop and the thing that turns change.

medium	the loop you walk	what turns as you walk	the whole number
vortex lines in water	one vortex line	the flow made by the other line (the same law as for a current)	how many times the lines link
a superconducting ring	the ring	the phase of the ring's one shared quantum state	the flux quanta trapped in the hole
a ring-shaped superfluid	the ring of gas	the phase of the gas	the units of circulation
closed circular DNA	one strand	the other strand, winding around it	the linking number
twisted light, radio and sound	a circle around the beam	the wave's phase	the twist

the lattice ring (the magnetic forest) the ring of sites the phase of an electron carried around it the winding

Sources: the price table's LATCH row, one column per medium (LATCH_ATLAS_PRICE_TABLE_2026-09-23).

What the walk does not tell you

The walk reads the count and nothing else. It does not say how much energy the loops hold, or whether the medium can hold any. The Latch paper's bare curves read their count to five decimal places and hold nothing back at all. Whether a whole number is a latch, a label or a powered lock is a question about the medium, and that is what the rest of the book prices.

Sources: the Latch paper, *Linking Is the Latch* (its section "The instrument", read in full); Gauss's linking integral; Ampère's law; Klenin and Langowski, as the Latch paper cites them; the price table's LATCH row. Depth of reading and identifiers: the reference list.

Runs: READ_R1_FIGURES 20260928_054233 (read_r1_figures.py version 2), which reproduces LATCH_PROBE 20260923_191121.

2. Five names for one whole number

The whole number this book reads goes by many names, because each field of science met it on its own loop and named it there. This section lines the names up. Four of them are the same walk from section 1 with a different thing turning, and they always give whole numbers. The fifth, the geometric phase, is the same kind of turning angle but is not, in general, a whole number. Seeing why it is not is the quickest way to see why the other four are.

Three things come out of this section. Linking, winding, twist plus writhe, and a wave’s twist are one object read four ways. Twist and writhe, taken one at a time, can be any size at all; only their sum is whole. And the difference between a whole-number count and a turning angle is whether the thing you check at the end is fixed to the ground or carried along with you.

The five names

Section 1 said what the book counts: the turning that adds up as you walk a loop and come back. The table below is the five names that count goes by, one row for each, with the loop it is walked on, the thing that turns, whether the answer is always a whole number, and where in the book it appears. After the table, each name gets its own short passage in the same order, because the differences between them are what the rest of the book leans on.

name	the loop	what turns as you walk	always whole?	where it appears in this book
linking	two closed curves	the field of one, walking the other	yes	section 1; the bare filament, DNA, vortex lines in water
winding	one loop through a field	the direction of the field, like the weather vanes	yes, if the field never dies away on the loop	the superconducting ring, the ring superfluid, the lattice ring
twist plus writhe	a ribbon, with two edges	the ribbon around its spine, plus the coiling of the spine	the sum yes, each part no	DNA, the knot-tape
a wave’s twist	a circle around the beam	the wave’s phase	yes	light, radio and sound
geometric phase	a loop of slow changes	the leftover turn of something carried along	only in special cases	this section; section 3

Sources: the price table’s *LATCH* and *WRITE* rows.

Winding: the weather vanes again

The ring of weather vanes in section 1 was already a winding. Each vane is fixed to its own spot on the ground and points wherever the wind there points. When the walk comes back to the first vane, it is the same vane in the same wind, so the total turning is whole. Every winding in the book works this way: the phase of a superconducting ring’s shared quantum state, the phase of a ring of gas, the phase of an electron carried around a ring of atoms. Each is a direction fixed to its place, read around a loop. The one way a winding can fail to be whole is if the thing that points dies away somewhere on the loop, so there is no direction to read there; that is the torn loop of the noise checks.

Twist plus writhe: the ribbon

A ribbon has two edges, and the linking of those two edges is a whole number. It splits into two parts. Twist is how much the ribbon turns around its own spine. Writhe is how much the spine itself coils in space, the way an old telephone cord coils. Each part can be any size. Only their sum has to be whole. This is Călugăreanu’s theorem, from 1959 and 1961, often given other names; the book follows Moffatt and Ricca (1992), who argue it should carry his name.




Figure R2.1. A band on a coiled spine, given 0, 2, 4 and -2 extra half-twists. The writhe (orange) is the shape of the spine alone and stays at -0.9193 every time. The twist (blue) changes by one whole turn for every two half-twists. Their sum (black diamonds) lands exactly on whole numbers. Pure mathematics.

Measured, the spine's writhe is -0.9193 for every band. The twist is -2.0807 , -1.0807 , -0.0807 and -3.0807 for 0, 2, 4 and -2 extra half-twists. Twist plus writhe comes to -3.0000 , -2.0000 , -1.0000 and -3.9999 , and the linking of the two edges, read by the Gauss walk of section 1, is -3.0000 , -2.0000 , -1.0000 and -4.0000 . Neither part is a count; the sum is. That is why DNA can be wound tighter by coiling instead of twisting, and why a phone cord that has been turned too many times throws itself into coils: the ribbon trades twist for writhe and keeps the sum.

A wave's twist

Twisted light, radio and sound carry a whole number in the way their phase winds around the beam's axis. Walk a circle around the beam and read the phase, and it comes back to itself after a whole number of turns. It is a winding, read on a circle around the beam instead of a ring in space. The price table's light, radio and sound columns read it in three media.

The geometric phase: an arrow carried along

Now change one thing. Instead of reading arrows fixed to the ground, carry one arrow with you. Walk once around a circle of latitude on a globe, keeping the arrow flat against the surface and never deliberately turning it, and compare it at the end with the direction it pointed when you set out. It comes back turned. How much depends on the loop: the leftover turning, in whole turns, is one minus the sine of the latitude. In words, the smaller the patch of globe the loop goes around, the less the arrow turns.




Figure R2.2. The same turning seen as a Foucault pendulum sees it. Against the ground, the pendulum's swing turns by the sine of the latitude, in whole turns, each sidereal day: none at the equator, one whole turn at the pole, and 0.753 of a turn at Paris. Only the two ends are whole numbers. The latitude of Paris, 48.85 degrees, is assumed.

Measured by carrying the arrow in 20,000 small steps, the leftover turning is 0.7412 of a turn at 15 degrees, exactly half a turn at 30 degrees, 0.2929 at 45 degrees and 0.1340 at 60 degrees, matching one minus the sine of the latitude to every printed digit. A Foucault pendulum, first shown in Paris in 1851, shows the same geometry against the turning Earth: its swing turns by the sine of the latitude each sidereal day, 0.753 of a turn at Paris.

The arrow inside a coiled fibre. The same arrow rides inside a coiled optical fibre. Tomita and Chiao wound a 180 cm fibre into helices and measured how far the plane of the light's polarization turned between one end and the other. As the light runs along the coil, its direction of travel walks a loop on the globe of directions, and a helix wound at a steady slant walks a circle of latitude there. The polarization is the carried arrow. A fibre slanted 75 degrees from the helix's axis walks latitude 15 degrees and turns the light 0.7412 of a turn; one slanted 60 degrees walks latitude 30 and turns it exactly half a turn. These are the arrow's numbers to the last digit, worked by pure arithmetic on the paper's printed inputs; the fibre's arc length of 175 cm is inferred from the printed range of pitches and diameters. The paper found the turning equal to the patch of the globe of directions that the loop encloses, its solid angle, for every coil tried. The steady helices ran from 0 up to 5.2 steradians. For the deformed helices, the solid angles rebuilt here from the printed shapes run from 3.58 to 4.95 steradians, and they sit where the paper plots them (the plotted positions read by eye). A fibre laid flat on a table walks a great circle, which cuts the globe of directions in half, so it can turn the light only by nothing or by one whole turn. The authors saw "no significant optical rotation" for flat paths, "independent of the paths's complexity" (A. Tomita and R. Y. Chiao, "Observation of Berry's topological phase by use of an optical fiber", *Phys. Rev. Lett.* 57, 937, 1986, read in full). That is the octant walk's hemisphere, below, in light.

The octant walk. This walk was Bear's suggestion. Three great circles at right angles cut the globe into eight matching triangles, each with three right-angle corners, so each triangle's corners add to 270 degrees instead of a flat triangle's 180. Carry the arrow around one of them, from the pole down to the equator, a quarter of the way around the equator and back up to the pole, and it comes back turned exactly a quarter turn. Around two of them it comes back turned half a turn, around three it comes back three-quarters of the way round (which an arrow shows as a quarter turn the other way), and around four, which is a whole hemisphere, it has turned one whole turn and looks as if it never turned at all. The rule, in words: the arrow turns by twice the share of the globe the loop goes around. It lands on a whole number only when that share is a half or the whole globe, and the extra 90 degrees in each triangle's corners is the same quarter turn, seen in the corners instead of in the arrow. The same octant closes the Crossover Atlas (Schulz 2026, version 3 of 22 August, concept 10.5281/zenodo.22064125), as its Chapter 8 and the receipts

spherical_excess.py and berry_monopole.py: the octant's corner excess, its area worked out three independent ways and the turn of the carried arrow all come out at 1.57080 radians, a quarter turn, and a quantum spin steered around it picks up half of that, because a spin needs two full turns, 720 degrees, to come home unchanged.

Here is the difference, and it is the whole point of this section. A weather vane is fixed to its spot, so when the walk returns to it, it is compared with itself, and the total must be whole. The carried arrow is not fixed to anything. It is compared at the end with where it started, and nothing forces those to match. So a winding counts, and a carried turning measures. The geometric phase is a real, useful, measurable angle, and it belongs to the same family as the other four; it simply is not a whole number unless something special makes it one, as at the ends of this figure. Section 3 turns that difference into the book's test. The integer diagnostic that section uses files winding, deficit angle, Wilson loop, Berry phase and holonomy together as one object; this section has shown which members of that family always count.

The count inside the angle. Berry found the rule behind all of these in 1984, for any quantum system carried slowly round a loop of settings. Besides the phase that clocks the time elapsed, the system picks up a geometric phase, and "the system records its history in a deeply geometrical way" (M. V. Berry, "Quantal phase factors accompanying adiabatic changes", *Proc. R. Soc. Lond. A* 392, 45, 1984, read in full). For a spin steered round a loop by a magnetic field, the phase is the solid angle times a whole number: " $\exp\{i\gamma_n(C)\} = \exp\{-in\Omega(C)\}$ " (his Eq. 27). Here n is the spin's component along the field, and the phase "depends only on the eigenvalue n of the spin component along B and not on the spin s of the particle". The angle measures and the multiplier counts.

For the simplest two-state system the multiplier is a half. Berry notes that the answer does not depend on which surface is drawn across the loop, "because Ω can change only in multiples of 4π ". A half of 4π is 2π , one whole turn, so the choice of surface never shows.

Tomita and Chiao measured the multiplier for light. The ratio of the turning to the solid angle stayed within a few hundredths of 1.00 for every deformed helix (their Figure 4, read by eye), which they read as "experimental evidence for the quantization of the 'topological charge' of the system, which in this case is the helicity of the photon". What they call topological, the turning's indifference to everything about the coil but its solid angle, this book calls geometric; the whole number in their experiment is the slope.

The special cases where the angle itself lands on a whole number are the ones a symmetry forces. For a real Hamiltonian, the phase factor round a degeneracy is -1 if the loop encircles it and $+1$ otherwise (Berry's Eq. 19): a half turn or none. Waves in a two-dimensional lattice carry a Berry phase that is " π -quantal and topological" in one limit of their coupling, and in between "either approximately equals π or is not quantal" (B. B. Guzina, O. Oudghiri-Idrissi and S. Meng, *Proc. R. Soc. A* 478, 20220110, 2022, abstract read). Berry also records Barry Simon's reading of the phase: it "has a mathematical interpretation in terms of holonomy", the mathematicians' name for the leftover turn that a loop leaves.

Sources: Moffatt and Ricca 1992 (read in full); Berry 1984 and Tomita and Chiao 1986 (both read in full, from copies Bear supplied: Berry from the Royal Society's copy with every quoted equation checked against the page images, Tomita and Chiao by OCR of a printed copy with the figures checked by eye; both journal records checked); Guzina, Oudghiri-Idrissi and Meng 2022 (abstract); the Crossover Atlas, Chapter 8; Foucault's pendulum, 1851; the price table's LATCH and WRITE rows. Identifiers: the reference list.

Runs: READ_R2_FIGURES 20260924_103119 (read_r2_figures.py; pure mathematics, one assumed latitude), reproducing ROTOR_RING_CHECKS 20260923_232715; OCTANT_WALK_CHECK 20260924_104321 (octant_walk_check.py, pure mathematics); B1_READ_CHECKS 20260923_233341; BERRY_FIBRE_CHECK 20260926_000823 (berry_fibre_check.py version 1).

3. The integer diagnostic

"Integers are real when they count something. They're decoration when they measure something." That is the book's standing test, adopted as the program's standing diagnostic in August 2026. Section 2 found where the line falls: a weather vane fixed to its spot is compared with itself at the end of the walk, so its total must be whole, while a carried arrow can come back turned by any amount. This section turns that line into three checks any number can be put through.

Three things come out of it. A count closes in on a whole number as the instrument gets finer; a measurement settles wherever it settles. A count sits still while the shape around it changes; a measurement slides. And a number that looks whole still owes an answer to one question, what is being counted, and the Latch paper's own first version failed that question and was withdrawn.

Check 1: change the instrument

Read the same thing with a finer instrument and watch where the reading goes. A count closes in on its whole number. A measurement closes in too, but on whatever value it happens to have.




Figure R3.1. Change the instrument. The linking of two bands threaded once (blue) closes in on a whole number, a fourfold step each time the points double. The writhe of a trefoil knot (orange) also settles as the points increase, but on a value 0.354 from the nearest whole number. Pure mathematics.

Measured, the linking reads -1.00066 , -1.00016 , -1.00004 and -1.00001 at 100, 200, 400 and 800 points. The trefoil's writhe reads -3.3663 , -3.3572 , -3.3549 and -3.3543 at the same point counts, settling about 0.354 from the nearest whole number, where no amount of care will move it (the last three writhe readings reproduce the Latch paper's printed values). Both readings are good measurements. Only one is a count.

Check 2: change the shape, without cutting

Bend, stretch or coil the thing smoothly, never cutting it and never passing one strand through another. A count stays put, which is the textbook property of a topological count: it survives any smooth change of shape. A measurement moves with the shape.




Figure R3.2. *Change the shape.* The band of section 2 is coiled further and further out from its ring. Its writhe (orange) slides one way and its twist (blue) the other, while their sum stays on -3 at every step.

Measured, as the spine's coil grows from 0.15 to 0.55 of the ring's radius, the writhe slides from -0.2450 to -1.4523 and the twist from -2.7550 to -1.5476 , while twist plus writhe, and the linking of the band's two edges read by the Gauss walk, stay at -3.0000 at all five shapes. The Latch paper made the same check another way: shrinking a linked pair 16.7-fold, its linking read -1.0000 at every step. And section 6 of this part casts 1,500 shadows of three linked pairs from random directions; every shadow shows a different tangle, and every one gives the same count.

Check 3: ask what is being counted

The first two checks can be passed by accident, because some measurements happen to sit near whole numbers. So the last check is a question, not a measurement: what is being counted, and what mechanism does the counting? A number that cannot answer is decoration, and it stays decoration until someone finds the mechanism. That is not a charge of fraud. It leaves the door open for a counting mechanism to turn up later, and it refuses to walk through the door before one does.

Two examples show why the question is needed. A year is 365.2422 days. That is close to 365, but nothing counts days into a year; the Earth's spin and its orbit are two measurements that do not share a counter, which is why the calendar needs leap years. And "round" numbers are everywhere. Only 1,197 whole numbers up to 100,000 can be built from the factors 2, 3, 5, 7 and 11 alone, yet 98.88 percent of all whole numbers up to 100,000 sit within half a percent of one of them, and 99.27 percent of those above 1,000 do. A measured cycle that "matches" such a number has matched almost anything.

The book's latches answer the question. A DNA gyrase enzyme changes the linking number "in steps of two" because it cuts both strands, passes one piece of the double helix through the gap and seals it (Brown and Cozzarelli, *Science* 206, 1979); the counted thing is a strand passage. A superconducting ring holds whole flux quanta because its one shared quantum state must fit around the ring a whole number of times (the price table's superconducting column); the counted thing is a turn of its phase. A voltage standard's junctions sit on whole-number steps because each is locked to a microwave tone (NIST, "A Primary Voltage Standard for the Whole World", 2013); the counted thing is whole turns of each junction's phase for every cycle of the tone.

The check the Latch paper failed

The Latch paper's first version found an energy that climbed with the linking number, like a staircase: 9.34, 10.13, 11.33, 12.91, 14.84 and 17.15 for linking one to six, against 8.82 for an unlinked pair. It rose with the count, so it looked as if it counted. It did not. Winding a pair more times on the same doughnut takes more rope, and each strand was 13.14 long at linking one but 25.93 at linking six, against 12.57 for the unlinked circles. Given the same length of rope, the energy falls as the linking rises: 9.36, 8.78, 8.13, 7.59, 7.19 and 6.94 (the Latch paper, version 2, lines 67 to 79 of its text, which withdrew the staircase; the paper's instrument check and the withdrawal were read in full). The staircase measured the rope. A number that rises with a count is not a number that counts, and check 3 is the one that catches it.

What the diagnostic does not say

It does not say measured numbers are worse than counts. The writhe, the geometric phase and the length of the year are good, useful measurements; they are just not whole numbers, and the book does not price them as latches. It also does not say a number near a whole number is fake. It says the number owes a mechanism, and until the mechanism is shown, the book files it as decoration and moves on.

Sources: the Latch paper, *Linking Is the Latch*, version 2 (read in full; its instrument check, and lines 67 to 79 for the withdrawn staircase, the rope lengths and the equal-rope energies); Brown and Cozzarelli 1979 (text read); NIST 2013 (read); the length of the year, still to source. Identifiers: the reference list.

Runs: READ_R3_FIGURES 20260924_103544 (read_r3_figures.py, pure mathematics); SHADOW_COUNT_CHECKS 20260924_100559 for the shadow count.

4. The count belongs to the loop

Assembled from the Forest companion, Instruments for the Magnetic Forest; every number names the companion’s own run. Nothing in this section adds a measurement; it reads the companion’s results as the book’s section 4.

Change the loop and the count changes. Change anything else, and it does not. Take the loop away, and there is no count at all. Section 1 said the whole number belongs to the loop and to what the loop goes around; this section shows it, on the one object in the program where closing a loop costs a single line of bonds.

The object is a lattice: a grid of sites, like the squares of a chessboard, where an electron can hop from each site to its neighbours. A magnetic field through the grid gives the electron’s phase a small extra turn around every square. When that extra turn is a simple fraction of a full turn, say one fifth, the pattern it makes repeats every five rows. Wrap the grid into a ring, so the last row joins the first, and ask one question: does the pattern’s repeat fit around the ring a whole number of times? Where it does, the allowed energies open a clear gap. The Forest companion calls that gap a tooth, and a tooth is a count showing itself. The pattern of gaps that a lattice makes under a magnetic field is Hofstadter’s butterfly, and the companion credits Hofstadter for it; the lattice explanation in these opening paragraphs is this book’s plain-language framing of the companion’s section 1.

Four things come out of this section. A tooth belongs to how the loop closes, not to any number in it. Uneven treatment breaks a count only when it breaks the loop’s own sorting of things into classes, which is the companion’s “chaos and binning”. One line of bonds is the whole difference between noise and arithmetic. And shaking every bond in the object by half does not move the count.

Five dividing ten

The same five-fold pattern was tried on rings of 10, 12 and 14 rows. On the ten-row ring, the tooth at five stands at 1.55 on the companion’s scale; on the twelve-row ring, the same five lies flat at 0.52. In the companion’s words: “The tooth is not a property of the number five. It is a property of five dividing ten.” Across eight ring sizes, twenty-one teeth stood up, and every one of them divided its ring, with no exceptions (Forest companion, §3). The arithmetic behind it is commensurability, a pattern fitting its ring a whole number of times, the same fact that locks two rates at a whole-number ratio in WRITE section W8; it is textbook in periodic structures and in mode locking (recalled, not cited here), and what the companion adds is its use as a read: a comb whose standing teeth name the ring’s divisors.

Chaos and binning

The companion’s section 5 names the pattern that runs under this whole book:

“A tooth is a binning, and binning has been a companion of this program from the beginning. Across every substrate the same pairing keeps returning: a continuous field that can be stirred into chaos, and underneath it a discrete bookkeeping that sorts the mess into a countable number of classes and refuses to be stirred at all.”

At a pattern that repeats every three rows, the rows sort into three classes: rows 0, 3, 6 and so on; rows 1, 4, 7; rows 2, 5, 8. The tooth exists because those classes are alike. So the companion made a prediction that no other explanation makes: lighting the rows unevenly should matter by whether it treats members of a class alike, not by how uneven it is. It tested that by making every pattern of lighting exactly as uneven, six times brighter at its brightest row than at its dimmest, so that only the shape of the unevenness differed.

how the rows were lit (all six-fold uneven)	tooth at 3	tooth at 4
evenly (the reference)	1.85	2.57
fading steadily from one end (no repeat)	0.76	0.84
at random, row by row (no repeat)	0.84	0.91
the same within each class of three (repeats)	1.57	1.92
the same within each class of four (repeats)	1.84	1.51

Run COMB_BINNING 20260914_214248, from the Forest companion’s §5.




Figure R4.1. Same unevenness, opposite outcome. Every lighting pattern has the same six-fold spread; the patterns with no repeat drop the teeth to the background, while the patterns that repeat with the ring's classes leave them standing. The Forest companion's Figure 3, reproduced unchanged.

Same unevenness, opposite outcome. The count reads the class label, and the class label is blind to how hard the rows are driven. The companion draws one lesson straight away. Heat lights a lattice the way the steady fade does, brightest at one end and dimmer and dimmer along it, with no repeat, so heat destroys the teeth, but not because it dims the far rows. A fade of the same depth that repeats with the ring's classes leaves them untouched.

One line of bonds

The cleanest test costs almost nothing. Stack rings of sites into a chain and join each to the next. Left open, the chain has two loose ends. Closed, the last ring joins the first. Nothing else differs between the two.




Figure R4.2. One line of bonds. The same six chains, open on the left and closed on the right. Closing the chain turns the count on, with no false teeth. The Forest companion's Figure 4, reproduced unchanged.

Open, the six chains showed five teeth, and not one of the five divided its chain: scatter, not a signature. Closed, they showed fifteen teeth, every one a divisor of its chain, with no false ones, and a chain of a prime number of rings showed only itself. In the companion's words: "One line of bonds is the whole difference between noise and arithmetic."

Shake every bond

A result that needs a perfect lattice is a tuning, not a law. So the closed chain was rebuilt with every bond's strength changed at random by up to half, and averaged over eight random versions at each strength. The teeth and the controls both sagged, as disorder makes everything sag, but the ratio between them held: 2.21 on the clean lattice and 2.14 at fifty per cent randomisation, with no trend in between. Meanwhile the electrons' states bunched up, their concentration roughly doubling from 1.92 to 3.82. "The states localise and the teeth do not notice." The bonds are what disorder changes; the count is a property of how the loop closes, and a random bond cannot undo a closing.

No loop, no count, and the one that was never a count

An open patch of lattice with the same number of rows shows no teeth at all. Nothing wraps, so there is no "all the way around" to fit anything into. The only feature an open patch keeps is at half a turn per square, where every phase factor is plus or minus one. That feature is a symmetry of the equations, not a count, and the companion caught it acting like one: a second magnetic flux, threaded through the hole of the doughnut, leaves four of the five counting teeth unmoved, while the half-turn feature comes and goes twice for every one full quantum of that flux. It is the one feature that answers to the second loop, and it is the one feature that was never counting. Section 3's third check, asking what is being counted, is exactly what tells the two apart.

The honest edge

Dividing the ring's size gives a tooth the right to exist; it does not guarantee one. Of twenty-nine possible divisors across the eight ring sizes, only twenty-one teeth stood up, and a ring with a prime number of rows showed essentially none. "The arithmetic is a gate, not an engine" (§3).

Most of the eight missing teeth are already explained by this section. Five of the eight are the half-turn point, which the subsection above shows is a symmetry and not a count: on every many-divisor ring tested it is the weakest divisor tooth, and on three of five it sits below the background (the companion's §4). Where the symmetry and the count meet, the count loses. That leaves three misses unexplained: the eighth and sixteenth divisors of the sixteen-row ring, and the thirteen-row ring's own thirteen. (The breakdown of the eight misses is printed in that run's own ledger, in the companion's deposited crate, along with the height of every divisor tooth: they are $q = 2$ on the twelve-, sixteen-, eighteen-, twenty-four- and thirty-row rings, $q = 8$ and $q = 16$ on the sixteen-row ring, and $q = 13$ on the thirteen-row ring, checked against the deposited ledger on 29 September.) A likely reason for the weak half-turn tooth, recalled and not measured, is that at exactly half a turn per square the lattice's energies close up at the centre, so the gap being measured is not there to open. Which of the remaining permitted teeth appear stays an open question in this book.

Closure is the ceiling

The companion closes its measurements with one sentence that could stand over this whole part of the book: "a system that cannot return to where it started cannot say how many times it went around." It is the ring of weather vanes from section 1, seen from the other side. A walk that comes back to its first vane must give a whole number; a walk that cannot come back has nothing to count.

Sources: the Forest companion, *Instruments for the Magnetic Forest* (read in full; its sections 2 to 8 and 10, and its Figures 3 and 4); Hofstadter, as the companion credits him. Identifiers: the reference list.

Runs: the companion's own, as deposited: COMB_COMMENSURATION 20260914_042324, COMB_BINNING 20260914_214248, RING_CHAIN_CLOSURE 20260915_045420, COMB_DISORDER 20260917_040415, COMB_DIVISORS 20260914_183142 (its ledger's part D1 for the eight misses and the tooth heights), FOREST_COMB_SURVEY 20260914_030628, TORUS_TWO_FLUXES 20260917_042223.

5. Reading through noise

Every number in this section is from pure-mathematics runs made for this book on 23 and 24 September.

A walk around a closed loop can only return a whole number, however noisy each step is. Noise does not make fractions. It makes wrong whole numbers, and a wrong whole number looks exactly like a real change. So the cure for noise is not to read more often but to read well, and a whole number makes reading well unusually easy, because two good reads either agree exactly or they do not.

Five things come out of this section. Noise turns a count into a wrong count, never into a fraction. Too big a turn between neighbouring samples breaks the count, which sets a limit on every ring of sensors or sources: it can carry whole numbers only below half its number of members. Averaging along the loop reads a count from samples buried below the noise. Two loops, each read twice, tell a real object from a noisy read. And a loop torn open by the thing it goes around gives no fraction either; the count simply flips.

Noise makes wrong whole numbers

The ring of weather vanes from section 1, read with noise on every vane: a whirl with a turning of one, walked with 16 or 64 vanes, 20,000 times at each noise level. Across every noise level, the count never landed further than 2.7×10^{-15} from a whole number, which is the computer's own rounding. What noise did was make the count wrong.




Figure R5.1. Noise never makes a fraction; it makes a wrong whole number. The share of walks that return the wrong whole number, as the noise on every vane grows. Drawn by `read_r5_figures.py`, which prints every plotted number back, from the ledger of check A.

At 0.6 of a radian of noise on every vane, 1.1 per cent of walks with 16 vanes came out wrong, and 1.4 per cent with 64. At 0.9 of a radian, 23 and 45 per cent. The reason is how each step is read. Between two neighbouring vanes the walk takes the shorter way round; if noise pushes one step past half a turn, it is read the long way, and the total jumps by exactly one whole turn. More vanes means more steps, so at high noise more vanes means more chances to jump.

The half-turn limit

The same rule sets a limit before any noise arrives. If the field itself turns a long way between neighbouring samples, noise needs only a small push to carry a step past half a turn. With 16 samples and the same noise throughout, a count of 1, where neighbours are 22.5 degrees apart, was wrong 6.3 per cent of the time; a count of 4, where they are 90 degrees apart, 44 per cent; a count of 5, at 112.5 degrees, 73 per cent. Turning the whole field by the same amount everywhere made no difference at all: 6.1 to 6.5 per cent across five settings (check A). Only the turn between neighbours matters.

Every ring of sensors or sources in this book meets the same limit. A ring of eight antennas can write and read counts only below four; at four, neighbours are half a turn apart and a count cannot be told from its mirror. The twisted-sound experiment of 2008 kept a quarter turn between neighbouring loudspeakers in both its set-ups, well clear of the limit. Engineers know the same rule from sampling any signal, the Nyquist rule: sample at least twice per turn, or the turns fold over.

Below the noise floor

Now make each sample weaker than the noise that rides on it, so that each vane on its own points almost anywhere. A single walk reads the count wrong nearly every time. But the count does not depend on the loop’s exact shape, so every sample around the loop is a repeat measurement of the same whole number, and neighbouring samples can be averaged, as arrows, before the walk is made.




Figure R5.2. Below the noise floor, averaging along the loop recovers the count. 1,024 samples around the loop, read directly or averaged in blocks before counting. Drawn by `read_r5_figures.py`, which prints every plotted number back, from the ledger of check D.

field strength on each sample, as a share of the noise	read directly	blocks of 16	blocks of 64	blocks of 128
0.1	96 % wrong	84 %	61 %	39 %
0.2	96 %	76 %	17 %	3.8 %
0.3	96 %	64 %	2.0 %	0.13 %
0.5	95 %	15 %	0	0
1.0	94 %	0	0	0

Run `LOOP_COUNT_NOISE_CHECKS 20260923_234529`, check D, 4,000 trials in each cell.

At a field only 0.3 of the noise, where each vane on its own is close to random, averaging in blocks of 128 brings the wrong count down from 96 per cent to 0.13 per cent. As a rule of thumb from these runs, the field strength times the square root of the number of samples averaged should reach about 3; the square root is textbook statistics, since averaging independent noise over many samples improves a reading by the square root of their number. Two cautions go with it. It works only when the noise is independent from one sample to the next; a slow drift does not average away. And the blocks must be short enough that the field barely turns inside one.

Two loops, read twice

What does it mean for two readings to agree? For a whole number, only one thing: the same whole number, exactly, with no tolerance to choose. That makes a second loop a strong check, with one care. Two loops around the same object must agree, but two loops can also differ honestly, because the difference between them counts whatever sits between them. So read each loop twice, with fresh noise. A difference that comes back on the second read is a real object between the loops; a difference that appears once is noise. On averaged reads, a whirl placed between the loops showed a repeated difference in 99.98 per cent of trials, and with nothing between them, in none. On poor reads the check is muddy, which is why the averaging comes first.

A torn loop

The one way a closed walk can fail is for the thing that points to vanish on the walk itself, so that there is no direction to read there. Move the whirl across the loop: the count reads 1 at every position inside, 1 when the whirl sits exactly on the walk, and 0 the moment it is outside, with the field on the walk falling to zero at the crossing. Even a torn loop gives no fraction; the count simply flips. Two nested loops read twice catch it, because the whirl then sits between them and their difference repeats.

Sources: the Nyquist sampling rule and the square-root law of averaging, both textbook (a named source for each: the reference list); the twisted-sound experiment of 2008 through the price table's sound column.

Runs: B1_READ_CHECKS 20260923_233341 (checks A and C); LOOP_COUNT_NOISE_CHECKS 20260923_234529 (checks A, B, D and E); RADIO_ESTIMATES 20260924_094021; SOUND_ESTIMATES 20260924_094237; READ_R5_FIGURES 20260924_105031 (read_r5_figures.py), which draws both figures from the first two ledgers.

6. Every shadow agrees on the count

Hold two linked loops up to a lamp and look at their shadow on the wall. Turn them, and the shadow changes: the loops cross each other more often or less, grow and shrink, bend into new shapes. One thing about every shadow stays the same. Count the crossings between the two loops with a sign, telling which loop passes over and which way the two are running, and half the total is the same whole number from every direction. It is the linking number, the same count the walk of section 1 reads. The shadows disagree about everything except the count. This is the Amber picture of the book's opening made exact: the linked loops are the one true pattern, each shadow on the wall is one of its shadows, and counting the crossings reads the pattern from any of them.

Four things come out of this section. From 1,500 shadows of three pairs of loops, cast from random directions, no two views need show the same tangle, and every one gives the same count. The shadow's count equals the walk's count. Observers moving past at any speed see the loops squashed, never cut, so they agree on the count too. And the tesseract picture from the Jellyfish paper, a closed loop in more dimensions than any projection shows, becomes exact: a tesseract is a cube pushed one step in a fourth direction, and read with time as that direction it is the history of a cube that sits still for a while; every observer's "now" is one slice through that history, and every drawing, every slice and every ruler agrees on its counts.

Reading a count off a shadow

Wherever the two loops' shadows cross, one loop passes over the other. Give that crossing a sign. Picture the loop on top and the loop underneath as arrows pointing the way each runs; if the top arrow has to turn anticlockwise to line up with the bottom one, the crossing counts plus one, and if clockwise, minus one. Add the signs of all the crossings between the two loops and halve the total. The half is there because each time one loop threads the other, the shadow shows it twice: once passing over and once passing under. This is the knot theorists' own definition of the linking number, half the signed count of the crossings between the two loops (Adams, *The Knot Book*, section 1.4; Rolfsen, *Knots and Links*, chapter 5, section D, where it is one of eight definitions proved equal), and it does not depend on the direction of the light, because the three Reidemeister moves that carry any one shadow into any other never change the sum.

Measured

Three shadows of the same two loops: different crossings, the same count




Figure R6.1. Three shadows of the same two loops, threaded twice. The first shows 4 crossings, the second 6, the third 10; every crossing is marked. Each gives the count -2 .




Figure R6.2. For each of three pairs of loops, how many of 500 shadows showed each number of crossings. The crossings vary from shadow to shadow; the count never does.

pair of loops	crossings seen in 500 shadows	the count in every shadow	the walk's count (Gauss integral)
threaded once	2 in 408 shadows, 4 in 92	-1	-1.00004
threaded twice	4 in 330, 6 in 101, 8 in 67, 10 in 2	-2	-2.00021
not threaded, but overlapping	none in 155, 2 in 345	0	+0.00000

Runs `SHADOW_COUNT_CHECKS 20260924_100559 (shadow_count_checks.py)` and `READ_R6_FIGURES 20260924_132818 (read_r6_figures.py)`; the second repeats the first with the same loops, seed and view directions, reproduces its summary line for line, and adds the tally of crossings, the three shadows of Figure R6.1 and the tesseract's counts. Pure mathematics.

The pair that is not threaded is the useful warning. In 345 of its 500 shadows the two loops cross, and every time the crossings come as a pair of opposite signs that cancel. A shadow that shows crossings is not a shadow that shows linking. Only the signed total counts.

Observers in motion

A ring moving fast past you looks squashed along the direction of its motion; that is special relativity's length contraction. A squash is a smooth change of shape, like the shrinking and coiling of sections 3 and 4, and it cannot pass one loop through the other. So observers moving at any speed see different shapes and measure different energies, but as long as the loops never touch, they agree on the count. Changing the count takes a cut, and a cut is a local event, something that happens at one place and one moment, which every observer sees happen. The tesseract passage below shows the same thing as geometry.

The tesseract, made exact

The Jellyfish paper keeps, as a way of seeing, "the tesseract as a closed loop in a state space with more dimensions than any projection shows" (its Appendix C, line 1200 of its text). This passage puts numbers in that picture, in four steps: build it, read it as time, slice it, and count what every view shares.

Built by pushing. Push a point one step and it traces a segment. Push the segment one step in a new direction, at right angles to the first, and it traces a square. Push the square and it traces a cube; push the cube in a fourth direction and it traces a tesseract. Each push keeps the shape at the start and at the end, and adds the path that every piece of it traces on the way: each corner traces an edge, each edge a square, each square a cube. So the pieces can be counted without drawing anything. The tesseract has 8 corners at the start and 8 at the end, 16 in all; 12 edges at each end and 8 traced by the corners, 32; 6 squares at each end and 12 traced by the edges, 24; one cube at each end and 6 traced by the squares, 8. A brute-force count, testing every set of the tesseract's corners to see whether it forms a piece, gives the same numbers at every step of the ladder.




Figure R6.3. *Each shape is the one before it pushed one step in a new direction. In blue, the shape at the start and at the end of the push; in orange, the path each corner traces. Under each, its pieces, counted by brute force and by the pushing rule, which agree.*

Read as time. If the fourth direction is time, the push is simply waiting. A cube that sits still for one unit of time has a history, and that history is a tesseract: the cube at the start, the cube at the end, the 8 lines its corners trace as they last, the 12 squares its edges trace, and the 6 cubes its square faces trace. A slice of space carried forward through time sweeps out spacetime in exactly this way. Every observer's "now" is one slice through that history. The observer at rest slices it straight across and finds the cube. An observer moving past slices it at a tilt, and the slice, computed directly and then measured in the observer's own frame, is a box squashed in the direction of motion: 0.866 as long at half the speed of light, 0.436 at 0.9 and 0.141 at 0.99, always with the cube's 8 corners, 12 edges and 6 faces. That is "Observers in motion", above, as geometry: different slices through one history, and no slice cuts anything. Two linked loops read the same way leave two histories, each a sheet rolled into a tube through spacetime, and their count can change only where the two sheets meet, at one place and one moment that every observer's run of slices must pass through. That is the invariance of the linking number under any motion in which the loops never touch: no such motion can change it, and each passage of one loop through the other changes the count by one (Rolfen, chapter 5, section D, where the linking number is the homology class of one loop in the complement of the other; Adams, section 1.4).

Turned, each direction in turn. A familiar moving picture of a tesseract draws it in perspective from the fourth direction, the far end small and the near end large, so that a small cube sits inside a large one, joined corner to corner by eight struts, and the whole thing turns. Bear's favourite animation is one of these, and he freezes it "at exactly the point where the cube is complete in the center". That frozen picture is the still cube's history, above. The small cube and the large cube are the same cube at the start and at the end of its wait, the eight struts are its corners lasting, and the six shapes between the two cubes are its six faces lasting. Turn a quarter turn further, in a plane that holds the fourth direction, and the small cube is whole in the centre again, but the struts now run along another of the tesseract's four directions. So which direction is time is a choice of view, not something the shape holds. Four quarter turns give each of the four directions its turn along the struts, and at every hold the drawing is the first one again, corner for corner and edge for edge; only the colours move. A third of a turn in ordinary space, about the long diagonal of the cube in the centre, then brings every colour home, and the loop starts again. The same turning, as a moving picture, ships beside the book as [tesseract_turn.gif](#), made at Bear's mark on this passage.





A tesseract drawn in perspective from four dimensions and turned a quarter turn at a time, read left to right, top row then bottom.
 At each hold the small cube is whole in the centre, and the eight struts run along another of the tesseract's four directions.
 Directions 1, 2, 3 and 4 in blue, green, purple and orange. After hold 4, the fourth turn brings direction 4 back onto the struts.

hold 1: time along direction 4

halfway through the quarter turn
that brings direction 1 onto the struts

hold 2: time along direction 1

halfway through the quarter turn
that brings direction 2 onto the struts



hold 3: time along direction 2

halfway through the quarter turn
that brings direction 3 onto the struts

hold 4: time along direction 3

halfway through the quarter turn
that brings direction 4 onto the struts







Figure R6.5. A tesseract drawn in perspective from four dimensions and turned a quarter turn at a time, read left to right, top row then bottom. At each hold the small cube is whole in the centre, and the eight struts joining it to the large cube run along another of the tesseract's four directions: orange, then blue, green and purple. Between the holds, the drawing halfway through each turn. Every hold is the same drawing; only the colours have moved. Drawn by `tesseract_turn.py`, made at Bear's request, which checks at each hold that both cubes are true cubes and which direction the struts run along, and writes the moving picture `tesseract_turn.gif`.

Sliced through space. Now let the fourth direction be a direction of space, and let the tesseract pass corner first through a three-dimensional world. A watcher there sees a point appear, grow into a tetrahedron, have its corners cut off into a truncated tetrahedron, become an octahedron at the halfway point, and then run the same sequence backwards to a point. Every one of those slices is a closed solid whose corners minus edges plus faces is 2. This is the famous picture, the tesseract passing through our space as a sphere passes through the flat world of Abbott's *Flatland* (1884). It is a slice through space, not through time. The direction straight across the corner-first slice has a length squared of +2 when time is counted with a minus sign, while the direction across any observer's "now" gives a negative number (−1 at every speed in the run). No observer, however fast, can tilt their "now" far enough to meet a still cube corner first. Push the tesseract through face first instead and the watcher sees a cube appear, stay exactly the same, and vanish: the still cube's history again.

A tesseract passing corner first through a three-dimensional space (then the same in reverse, back to a point at sum 4)

sum 0.5: tetrahedron
4 corners, 6 edges, 4 faces

sum 1.0: tetrahedron
4 corners, 6 edges, 4 faces

sum 1.5: truncated tetrahedron
12 corners, 18 edges, 8 faces

sum 2.0: octahedron
6 corners, 12 edges, 8 faces



Figure R6.4. A tesseract passing corner first through a three-dimensional space, sliced where its four coordinates add up to 0.5, 1, 1.5 and 2. From 2 to 4 the same shapes return in reverse, down to a single point.

What every view shares. A drawing of a tesseract lets its pieces overlap and cross, but it cannot cut or merge them, so every drawing, from every angle, shares the 16, 32, 24 and 8. One count goes deeper than the pieces. Cut the tesseract's skin with a finer ruler, each edge into 2, 3 or 4 pieces, and every count changes. Corners minus edges plus squares minus cubes does not.

the skin	each edge whole	cut in 2	cut in 3	cut in 4	corners – edges + squares (– cubes)
a cube's surface: corners, edges, squares	8; 12; 6	26; 48; 24	56; 108; 54	98; 192; 96	2 at every ruler
a tesseract's skin: corners, edges, squares, cubes	16; 32; 24; 8	80; 208; 192; 64	240; 672; 648; 216	544; 1,568; 1,536; 512	0 at every ruler

Run `TESSERACT_FIGURES 20260924_134306` (`tesseract_figures.py`), which builds the tesseract by pushing and by brute force, reads it as a cube's history, slices it corner first, face first and as moving observers slice it, and cuts its skin with four rulers; counted by brute force and checked against a formula.

This is section 3's first check, change the instrument, passed by a count that is not a linking number. The alternating count, the Euler number, belongs to the shape and not to the ruler. The Crossover Atlas measures the same fact on surfaces in its Chapter 7, with its receipt `gauss_bonnet.py`: a cube's corner defects total 4π , a deformed cube's still total 4π , a torus's total 0, and the same whole numbers come out of counting pieces, minus holes, plus enclosed voids. It also makes the Jellyfish picture exact in one more way. The tesseract's skin is a closed space of three dimensions, and its count is 0, the same as a loop's: the rim of a square has 4 corners and 4 edges, and $4 - 4 = 0$. Every closed space with an odd number of dimensions counts 0, a consequence of Poincaré duality (Hatcher, *Algebraic Topology*, Corollary 3.37, page 249), and the alternating count does not depend on how a space is cut up because it is the alternating sum of the space's Betti numbers, which are topological invariants (Hatcher, Theorem 2.44, page 146); both were read as printed. The loops in this section are the same kind of object as the tesseract. Their shadows are the drawings, their histories are slices of one spacetime, and their linking number, like the tesseract's counts, is what every view shares.

Sources: Adams 1994, section 1.4, and Rolfsen 1976, chapter 5 section D, for the linking number as a signed count of crossings and its invariance (both read as printed on 29 September 2026, from copies Bear supplied); Hatcher 2002, Theorem 2.44 and Corollary 3.37, for the Euler count (read as printed the same day); Wikipedia's pages on the Euler characteristic and the linking number, read the same day, which agree; Abbott's *Flatland* (1884); the Crossover Atlas, Chapter 7 and its receipt `gauss_bonnet.py` (read in full); the Jellyfish paper, Appendix C (read in full); special relativity's length contraction, textbook. Identifiers: the reference list.

Runs: `SHADOW_COUNT_CHECKS 20260924_100559`; `READ_R6_FIGURES 20260924_132818`; `TESSERACT_FIGURES 20260924_134306`; `TESSERACT_TURN 20260928_222652` (`tesseract_turn.py`, added on 28 September at Bear's mark on this passage). All pure mathematics.

7. What reading cannot tell you

The bare curve is the Latch paper's, Linking Is the Latch, read in full; "does not port" is the heading of its own section on the dial. Panel C, the wall estimates and the rest of the tables come from the price table's cells, each with its own run and source; that a walk reads history, and the three kinds of whole number, are page one's.

A reader gives you the count and nothing else. The Latch paper's bare filament, a pair of linked curves with no thickness, is read to within 0.00001 of its whole number (-1.00001 at 800 points), and it holds no energy against that number at all. Whether a count is a latch, a label or a powered lock is a question about the medium, not about the reading. That is where the price table takes over.

Four things come out of this section. The best-read count in the book holds nothing back. The same exact reading means three different things in three media. Everything a reader cannot see is a row of the price table, and each row has its own test. And the shadows of section 6 point the same way: the count is the one thing every view agrees on, and it is also the only thing a reader returns.

A perfect count that holds nothing

A perfect count that holds nothing back: the bare curve of the Latch paper




Figure R7.1. The Latch paper's bare curves, read three ways. (A) A linked pair shrunk 16.7-fold: each ring's length falls from 12.566 to 0.754 and the count reads -1 at every size. (B) Pairs linked 0 to 6 times, each given the same length of rope: the more they are linked, the less energy they carry, from 9.36 linked once to 6.94 linked six times. The pair linked six times, shrunk to half size, carries 6.81, less than two unlinked rings as built, 8.83 (at this resolution the counts read -1.000 to -6.033 , and the offsets shrink with more points). (C) One ring slid through the other's wire: the count goes from -1 to 0 to $+1$, and the energy drifts smoothly by 3.80 percent with no spike where the strands pass. The two Latch runs are the paper's tools as run for the price table before the deposit; the deposited record carries the same tools re-run the next day, and every number here is identical in those deposited ledgers, checked. Drawn by `read_r7_figures.py`, which finds the three ledgers by name under the atlas folder, reads them line by line and prints every plotted number back; no new measurement.

Panel A is the reader behaving, not the latch working: a uniform shrink cannot change the walk's total in any model, as the Latch paper says of its own test. Panels B and C show what the bare curve lacks. With the rope held fixed, more linking means less energy, so the count sets no floor under the store. And the count can be changed with nothing to climb: at the two places where one strand passes through the other, the energy departs from a smooth curve by 0.00314 and 0.00243, less than the largest departure anywhere else on the path, 0.00611. A curve with no body has neither a floor under its store nor a wall around its count. Its count is the most precise number in the book, and it holds nothing back.

One exact reading, three meanings

Here are four whole numbers, each read exactly. A reader handed any one of them would return the number and stop. It could not say which row it came from.

the ring	the whole number	what keeps the number	what the number holds back	kind
the bare curve	linking, -1.00001 at 800 points	nothing: a strand passes through with no spike in the energy	nothing	label
a ring of light (the Puckett ring)	about 70,758 wavelengths around the ring (with an assumed refractive index)	the ring's steadiness in time, with no wall	nothing; the count lasts one photon lifetime, 352 ns	label
a voltage-standard junction	whole-number voltage steps, 144.75 μV apart at an assumed 70 GHz	a microwave tone; NIST locks about 300,000 junctions to one source	nothing of its own; stop the tone and the steps go	powered lock
a super-conducting ring	a whole number of flux quanta	an order that must be destroyed to change the count: about 13,000 thermal jiggles at 4.2 K in a niobium wire 20 nm across; a thick wire's door, a vortex crossing it, is not yet priced	a current that does not measurably decay (a decay time "greater than 100,000 years")	latch

Rows from the price table: runs *LATCH_PROBE 20260923_191121*, *WALL_ESTIMATES 20260923_224526* (an order of magnitude; the barrier formula, checked against Joshi and Bhattacharyay 2011, holds only in wires thinner than about one coherence length), *LIGHT_ESTIMATES 20260924_064231* and *SC_PHASE_ESTIMATES 20260924_093814*; NIST news release, April 2013; File and Mills 1963, through two secondary sources.

The count is equally whole in all four. What differs is everything around it, and none of that is on the loop the reader walks.

What the reader cannot see, and the test that can

Each thing the reader cannot see has a test that can see it, and each test is a row of the price table.

what the reader cannot see	where the book asks it	price-table row	the test	an example
whether the count holds anything back	page one's question 3: does the number stop energy from leaving?	STORE, and the floor	price the energy against the count, with the rope held fixed	the bare curve fails it (Figure R7.1, B); a closed circle of DNA passes, storing about 185 thermal jiggles (with a textbook stiffness still to check)
how strongly it is held	page one's question 4: is the wall between whole numbers higher than the jiggle of heat?	WALL	price what must be climbed to change the count	a superconducting wire 20 nm across walls its count at about 13,000 jiggles at 4.2 K (an order of magnitude); thickness makes the wall only up to about one coherence length, past which the count escapes by a vortex crossing instead
whether something else is holding it	the powered lock of page one's three kinds	WRITE	stop the drive and read again	the junctions' steps hold only while the tone plays (NIST, April 2013)
how much each count is worth	the dial, which the Latch paper warns "does not port"	DIAL	measure it fresh in every medium	a water ring's energy rises with its size at held circulation (Saffman's formula, as the Jellyfish paper uses it); the lattice ring's current falls as one over its circumference, slope -1.0004

One more thing is out of the reader's reach: the future. Page one says a loop walk reads history, what the medium carried all the way around and still remembers. It says nothing about how long the medium will keep remembering. That is the LEAK and WALL rows again, and it is why a reading, however exact, is never a promise.

The shadows point the same way

Section 6 showed that every shadow of two linked loops agrees on the count, while the shadows disagree about shape, size and how many crossings they show, and moving observers disagree about energy. Turn that around and it is this section's lesson. The count is what survives every change of view, and so a count, read alone, carries no news about anything that changes with the view: not the energy, not the size, not the shape. Those are what decide whether a count is worth anything in the physical world, and they live in the other rows of the table. So the READ part ends its main line here and hands over to the price table. Section 8 adds the media where nature holds the pen and we can only read.

Sources: the Latch paper, *Linking Is the Latch*, version 2 (read in full, 23 and 24 September); the price table's cells for the light ring, the junction and the superconducting ring (Joshi and Bhattacharyay 2011 for the barrier formula's condition; NIST's news release of April 2013; File and Mills 1963 through two secondary sources); Saffman's ring-energy formula as the Jellyfish paper uses it; page one. Identifiers: the reference list.

Runs: the Latch paper's LATCH_PROBE 20260923_191121 and LATCH_FLOOR_CONFOUND_CHECK 20260923_194934 (deposited as the runs of 24 September, 20260924_031959, which print the same numbers); WALL_ESTIMATES 20260923_224526; LIGHT_ESTIMATES 20260924_064231; SC_PHASE_ESTIMATES 20260924_093814; READ_R7_FIGURES 20260924_134829 (read_r7_figures.py). Also drawn on: DNA_ESTIMATES 20260923_222155, RING_LATTICE_READER 20260914_032117.

8. The read-only media, briefly

In some media nature holds the pen, and we can only read what it wrote. Four such media are in the corpus: the slow movement of the ground along the Cascadia subduction zone, read by triangles of GPS stations; muons, read as they pass through rock or stop inside it; a diving sperm whale, read from the tag on its back; and the streams of stars in the Milky Way's disk, read from the Gaia survey. A fifth comes from the book's own reading: the waves that run along Earth's equator, which a whole number guarantees. Each is put to page one's questions, with one more that the latch media never needed: who holds the pen?

Five things come out of this section, one from each medium. Cascadia shows that a read-only count can be real, plentiful and well tested, and still tell you nothing about a store: its instrument throws magnitudes away by design, and its own authors could not say what a winding means. The muon shows where a whole number comes from when nature writes it: a probe that passes through reads a size, and a probe that stops reads a class. The whale shows a loop that cannot close, because one leg of the walk has no ruler, and a whole number in its calls that turned out to count nothing. The star streams show loops that read a turning well and, asked directly for a whole number, return zero; the only whole numbers in sight there are the ones we drew into the mesh ourselves. The equatorial waves show a whole number that nature writes and nobody can erase, which holds no energy at all and still decides what can exist: two waves that may only travel east.

medium	the loop	what came back	a whole number?	who holds the pen
Cascadia	a triangle of three GPS stations, read every week: section 1's walk with three vanes	+1, 0 or -1 for each of 441 triangles, over weeks 500 to 1383; larger whole numbers when triangles are added	yes: 10,312 triangle-weeks at +1, 10,199 at -1, 100,896 at zero	the ground
muons	a muon's path through rock; or a stopped muon's capture into an atom	passing through: the spread of its deflections, a size; stopped: X-rays whose energy names the element	only when the muon stops: which element, a class	the nucleus
a sperm whale	the dive track rebuilt from the tag, closed at the surface	a heading at every instant; a distance that drifts	no: the loop cannot be closed	the whale
the star streams	loops of a mesh over the Milky Way's disk, through the motions of Gaia stars	a coherent swirl, a turning	asked directly: zero on every ring	the galaxy
the equatorial waves	the equator of a sphere of wave shapes, where the phases fixed over its two halves turn against each other	two waves along Earth's equator that travel only east, Kelvin and Yanai	yes: 2, a Chern number	the planet's rotation

Rows from the Cascadia paper and its run WATERFALL_MAP 20260821_200450; the Field Guide's Case Four and its tools (muons); Comparative Biosonar and Navigation, documents 02 and 03 (the whale); Reading the Sky Through Loops, sections 3 and 4 (the star streams); Delplace, Marston and Venaille 2017 (the equatorial waves).

Cascadia: a count the ground writes

The Cascadia paper (Schulz 2026, *Loop Walking Cascadia*; read in full from its deposit crate, with the reports of its runs PACHINKO_SUITE, WINDING_BUDGET and WATERFALL_MAP; the other runs it names are quoted as the paper states them) builds section 1's walk out of GPS stations. It joins the 229 stations that survive its quality checks into 441 triangles, all turning the same way round. Each week, each station's east and north motion, with the slow trends and its own block's shared motion taken out, becomes a small clock hand, and around each triangle the paper does exactly what the weather vanes did. In its words, for each triangle each week, "sum the wrapped phase differences around the loop; the result divided by 2π is an integer — the triangle's defect charge, +1, 0, or -1" (its §1). This is section 1's walk with three vanes, and it meets section 5's limit, which the paper states in its own words: "a triangle reads winding mod 3 (a phase step of 240° per edge — a true +2 — reads -1; 360° per edge — a true +3 — reads 0)" (its §3). A clock hand shorter than a millimetre is gated out, because its direction is noise, so the instrument "reports the wall of the vortex, never its eye" (its §8). Over weeks 500 to 1383, that gives 10,312 triangle-weeks at +1, 10,199 at -1 and 100,896 at zero, with 268,437 more gated out as unreadable.

Section 4's claim, that the count belongs to the loop, is checked here on real ground. Every larger loop's whole number is the sum of the triangle charges it encloses. Around each station whose fan of triangles closes, walking the ring of neighbours gave the same whole number as adding the triangles inside it in 17,950 station-weeks out of 17,950 (its §10.5). Adding is also better than walking. A walk around a boundary of n edges can only hold whole

numbers up to $n/2$, but a sum of triangles has no such ceiling: on a planted field, 120 triangles summed to 120 while their 34-edge boundary could have shown at most 17 (its §8). On the real map, grown loops held whole numbers up to 7, and 9 when grown further, which the paper fences as expected, since random signs grow the same way (its §8).

The counts are well tested. The paper splits its record at the vault wall into the era it studied (weeks 500 to 1094) and the era its design never touched (weeks 1095 to 1383), gives each triangle a four-number behaviour card (how one-sided it is, how often it is charged, how often it repeats, how long it stays quiet), and asks whether the studied era's cards predict the untouched era's. They do: a rank correlation of 0.2697 against a shuffled null whose 95th percentile is 0.0940 (its Part II). Opened once, blind, the vault also showed the standing charges keeping their signs: of 16 triangles that stood again, 15 kept the same sign (its §6). Whether the three states run round a cycle like an engine comes back no: net circulation -14.0 inside a null band of -67.0 to $+63.3$, $p = 0.6965$. And the southern segment hands out a steadier count per episode than its own clock laid anywhere else, a spread of 0.4201 against the band's floor of 0.4236, which the paper itself calls "suggestive, not confirmed": the margin is 0.0035, and with three segments tested, one crossing by luck has about a fourteen percent chance (its §10.6).

What the counts cannot say is the whole of section 7, met in the field. The instrument's own report puts it plainly: "Counts are the claim, never energy: this instrument discards magnitudes by design". And the paper asks the question section 7 says no reader can answer, and cannot answer it: "Over half the lattice winds more than chance, and we still cannot say what winding *is*" (its §10.4). Its test was a second, independent instrument, the tremor catalogue from seismometers: if windings marked slipping ground they should follow the tremor, and if they marked locked ground they should avoid it. The correlation came back $+0.056$, inside a null band of -0.111 to $+0.109$ (its §10.4). Neither. The count is read well, and its meaning is not on the loop. Nobody can write a winding into the ground to find out, so Cascadia's whole numbers are read-only in the strictest sense.

Muons: a whole number from stopping, not from passing

The Field Guide's Case Four (Section 9 of the Guide's deposited version 5.0 of 18 August 2026; the book read it in draft v4.6, where it was Section 12, with its crate's `muonic_lines.py` and `falsify_v46.py` outputs, the case's after-action note and its citation ledger; the Guide's current version 6 no longer carries the case) pointed the whole loop-walking apparatus at the choice of probe for finding a dense object under rock, and drew the map over the probe instead of over the substrate. Its answer was that no probe wins both gates. A neutron reads a closer property but is removed in rock, in lengths of order ten centimetres, so a mountain stops it outright. A muon gets through, but its only loop-closeable channel is its spin, which turns with the magnetic field along the path, and rock and a lump of dense metal are magnetically silent.

What a passing muon does carry is a size. Muon scattering tomography images the spread of the muons' small deflections, and the Guide files that spread as the sharpest false friend in its vocabulary: the loop walker reads the average turning, and scattering reads the variance around an average of zero, "their signal is our noise". A passing muon reads how much matter is in the way, not a count.

A stopped negative muon is different. Captured by an atom, it drops into orbit close to the nucleus and gives off X-rays whose energies are set by the nucleus's charge, so the X-rays name the element: a class, a whole number the nucleus wrote. The Guide's own calculator checks its line energies against three measured anchors and reports its limits first: 0.3 percent low for hydrogen, 13 percent high for iron and 136 percent high for lead, because in heavy atoms the muon's innermost orbit lies inside the nucleus (the crate's `muonic_lines.py` output). And each stopped muon then meets one of two fates, decay or capture by the nucleus, which is an exact yes-or-no for one muon. The element only biases the coin: capture takes about 7.5 percent of muons in carbon and 97 percent in lead, crossing one half between elements 13 and 15 (the crate's `falsify_v46.py` output, from literature values it marks as still to check against Suzuki and Measday). The whole number is in the single event; a crowd of muons reads a fraction.

The Guide banks one idea from this as vision, not as a result: "An integer seems to come from a *transition*, not a *transit*." Everywhere a muon reads a whole number, something discrete completed. The book carries it the same way, as a pattern that a second field has to pay for.

A sperm whale: a loop with no ruler on one leg

The Comparative Biosonar and Navigation deposit (its documents 00 to 04 read in full; what follows is from document 02, result B3, and document 03, entries 1, 4 and 5) rebuilt one sperm whale's dive track from its tag: depth, acceleration and the magnetic field, giving pitch, roll and heading. Its own summary: "Heading anchored; distance adrift." The whale carries two anchors on its body, gravity for pitch and roll and the Earth's field for heading, so its direction is known at every instant. It carries no speedometer, so the distance it has travelled drifts. The natural loop walk is to close the dive at the surface and ask whether the whale comes up on its own rebuilt track, and the deposit reports that test as "inconclusive by construction", because the answer scales with a swimming speed that cannot be recovered. Its verdict: closure unmeasurable by construction, and the inconclusiveness "is itself the finding".

That adds a line to page one's list of ways the reader fails. A walk can close only on what the walker can anchor. The whale's loop closes in direction and cannot close in distance, so no count comes back, not because none exists but because one leg of the walk has no ruler.

The whale's calls hold the other lesson. The deposit's founding mistake was the idea that five is a special number in the click counts: several rules seemed to explain it, "odd numbers", "Fibonacci", " $n \leq 5$ ", and every one of them contained five, so none could fail and none could be told apart from the others. It was retired: "A rule that cannot fail explains nothing." A pre-registered test for a two-layer structure in the codas later came back "confirmed" and collapsed under inspection: its null passed the same test 54 percent of the time, and the "codas" it found were foraging clicks. That is section 3's diagnostic seen from the other side. A whole number that cannot fail to appear counts nothing, and the whale's calls, read so far, carry no count that survived.

The star streams: a winding measured at zero

The sky paper (Schulz 2026, *Reading the Sky Through Loops*; read in full from its deposit crate, draft v0.3 of 12 July with all nine checksums verified) lays a fixed mesh over the Milky Way's disk before it looks at any data. It starts from about 22.3 million Gaia stars with full motions after a frozen quality cut, keeps an old, settled population as its tracers, subtracts the smooth rotation of the disk, sorts what is left into the mesh's cells, reduces each cell to one agreed direction of motion, and walks closed loops of cells, reading how far a carried frame fails to come home. That failure is the circulation the loop encloses. What came back is the galaxy's spiral streaming: 85 of 509 mesh rings stand more than three standard deviations above a pre-registered shuffled null, where chance expects about one, organised into eight arcs a kiloparsec across, of alternating sign, with swirls of 5 to 15 kilometres a second per kiloparsec, peaking at 7.0 standard deviations and flipping sign cleanly when each loop is walked backwards, to 6×10^{-15} (its section 4.3). That is a turning, measured well: section 2's line, "a winding counts, and a carried turning measures", met in the sky.

The paper also asked the whole-number question outright. Its test K2c asked whether the field carries a whole-number winding, and the answer is printed plainly: "the measured winding was zero across all rings" (its section 4.4). The count was read, and it is zero.

There are whole numbers in the paper, and section 3's third check, ask what is being counted, finds them at once. The mesh's cells have five, six or seven sides, and each carries a curvature charge of six minus its number of sides: +1 for a pentagon, 0 for a hexagon, -1 for a heptagon. The cells inside sum to +9 and the boundary absorbs -9, a small Gauss-Bonnet ledger checked as the mesh is built (its section 3.1 and Figure 1). Those whole numbers are exact, and they count our own drawing, not the sky.

The quasars tell the same story from a different side. Quasars cannot move, so any common drift in their apparent motions is the observer's own acceleration toward the galactic centre, a published signal of about 5 microarcseconds a year. Read source by source, the paper recovers it at 4.83 ± 0.51 , within 0.4 standard deviations of the published 5.05 ± 0.35 ; read through its loops, it recovers the direction, to within about 11 degrees, but comes out 4.5 standard deviations short on size (its sections 5.3 and 5.4). Loops read where a structure is and whether it is real more reliably than how much of it there is, which is the paper's own summary of what it learned. Nature holds the pen, and the galaxy writes a turning, not a count.

The equatorial waves: a count the planet's turning writes

Two famous waves of the ocean and the atmosphere, the Kelvin wave and the Yanai wave, run along the equator and carry their energy only eastward. Delplace, Marston and Venaille (*Science*, 2017; read in full in its arXiv version, main text and supplement) showed that a whole number guarantees them: "The non-trivial structure of the bulk Poincaré wave modes encoded through the first Chern number of value 2 guarantees existence for these waves". In a thin layer of fluid on a turning planet, waves come in three bands, fast gravity waves above and below and slow, balanced motion between, and the planet's turning, through the Coriolis parameter f , opens gaps between the bands. Lay the waves out on a sphere of settings, the two wavenumbers and f , round the one point where the bands touch. The fast bands cannot be given one smooth phase over that whole sphere, "in particular over both of Earth's two hemispheres simultaneously". So the paper fixes the phase separately over the sphere's northern and southern halves, and where the halves meet, the two choices turn against each other twice as you walk once round the seam. That winding is the count, 2. It is section 1's walk, taken round the seam between two halves of a sphere of wave shapes, and the count is the same kind as the skyrmion's: a flux through a closed sphere, in whole units.

What the number buys is this. Where f changes sign, exactly two one-way waves must fill each gap, and Earth's equator is that place: the two are Kelvin and Yanai, carrying their energy east, in the paper's words with "eastward group velocity". "Even a misshapen sphere would support the waves": the count does not care about the planet's exact shape or about how f varies across the equator, while the waves outside the gaps do. The protection holds only inside the gap. In a model of the whole planet, a Kelvin wave at about one cycle a day passed a patch of stronger layering intact, and a slower one, at about 0.2 cycles a day, below the gap, was turned into westward Rossby waves. And the equator's own mirror keeps a Kelvin wave from scattering into a Yanai wave.

Who holds the pen is the planet's rotation. It is the tilt that breaks the time mirror and opens the gap, and its sign, which hemisphere you are in, sets the sign of the count; the one-way waves live where that tilt changes sign. Nobody can write or erase this count. It holds no energy and guards none: the paper counts waves, not a store, and for a continuous fluid only the difference across a gap is a true invariant. The waves themselves are read in the atmosphere from spectra of wavenumber and frequency, and in the ocean by satellites that measure the sea's height, "as water that warms also expands".

The planets: frames carried round the largest loops

The equatorial waves are a count written by how a frame is carried on a turning planet. Take that up one level, to the planets themselves carried round the Sun, and the same reading comes out with the roles reversed: the count is trivial and the leftover is exact. Ride a gyroscope, a frame that is carried and never pushed, once round a planet's orbit. In Newton's flat space it comes back exactly as it left, and the orbit closes exactly on itself: two leftovers, both zero. In the Sun's real, curved field both are small and nonzero, and they are the two ways a loop can fail to close. The frame comes back turned, by three pi times the Sun's gravitational radius divided by the orbit's size shrunk by its eccentricity, in symbols $3\pi GM/(c^2 a(1-e^2))$ per orbit: the holonomy of parallel transport round the loop, de Sitter's geodetic precession. And the path misses its own start by exactly twice that angle per orbit: the advance of the perihelion. The book's tool integrates the frame's turning rate along each Kepler ellipse and lands on the formula to eight figures for all ten bodies it tries; the orbital constants are recalled from the NASA planetary fact sheet and IAU values and marked so in the tool, and the first-order formulas are textbook general relativity, as in Misner, Thorne and Wheeler.




Figure R8.1. What is left over after one loop round the Sun, for the eight planets: the frame's turn (blue) and the path's miss (orange), in arcseconds per orbit. The two lines are parallel because the path's miss is twice the frame's turn for every planet; both fall with distance from the Sun.

body	frame's turn per orbit	path's miss per orbit	frame per century	path per century	whole turns of the frame since the solar system formed
Mercury	0.0518"	0.1035"	21.49"	42.98"	757
Venus	0.0265"	0.0531"	4.31"	8.62"	152
Earth	0.0192"	0.0384"	1.92"	3.84"	68
Mars	0.0127"	0.0254"	0.68"	1.35"	24
Jupiter	0.0037"	0.0074"	0.031"	0.062"	1.1
Saturn	0.0020"	0.0040"	0.0069"	0.014"	0.24
Uranus	0.0010"	0.0020"	0.0012"	0.0024"	0.04
Neptune	0.0006"	0.0013"	0.0004"	0.0008"	0.01
the Moon round the Earth	0.00002"	0.00004"	0.030"	0.060"	1.05
a polar satellite at 642 km	0.0012"	0.0025"	664"	1,328"	23,391

Run PLANET_FRAME_HOLONOMY 20260929_103728, Parts B and C; " is arcseconds; the age of the solar system is taken as 4.567 billion years.

Three of these numbers have been measured, and the tool's values are checked against them as recalled, each to be confirmed against its record at the references pass: Mercury's path misses by 42.98 arcseconds a century, the classic number; the Earth-Moon frame turns by 19.19 milliarcseconds a year, de Sitter's precession as lunar laser ranging sees it, 19.2; and a satellite's frame at Gravity Probe B's height turns by about 6,640 milliarcseconds a year against the mission's 6,606 predicted and 6,602 measured, the difference being the tool's rounded orbit.

Read as the book reads a latch, the whole number on any orbit is the loop itself, one, the same for every planet, and it needs no cut to change because nothing changes it. Curvature's leftover per loop is a carried turning, section 2's fifth row, and nothing holds it whole: it simply accumulates, so that Mercury's frame has made 757 whole turns since the

solar system formed, Jupiter's has just passed its first, and Neptune's has not made a hundredth of one. A measure, not a latch: nothing is held behind the count, and Question 3 finds nothing. Who holds the pen is the Sun's mass and the orbit's size, and every planet's frame turns the same way, with its orbit. That is what separates the planets from the waves above them. The waves' tie-breaker is the planet's own spin, the sign of the Coriolis pull, which is why Earth's one-way waves run east, and why on Venus, which spins the other way, the same argument sends them west. The orbit's leftover and the spin's count are both frames carried round loops, and they answer to different masters.

What the planets give the book is its cleanest picture of whole turns plus a remainder. The loop count is whole and trivial; the remainder is exact and physical, and the same three-halves and three appear in it that appear in the geodetic and perihelion formulas of every textbook; and the ratio of the two remainders, path to frame, is exactly two at first order, a number from the geometry and not from any planet. The residuals, winding on forever, are a label. Whether they are also a paper is not this book's question.

What the five share

In all five media we only read. None of them lets us price a store or a wall, so none of them earns a column in the price table. They stay in the book for what they teach a reader: in Cascadia, a count read well and held to its nulls; in the muon, where nature's whole numbers come from; in the whale, how a loop fails to close and how a whole number can fool you; in the star streams, a turning read well and a count honestly returned as zero; in the equatorial waves, a count that holds nothing and still decides what can exist; and in the planets, a loop whose count is trivially one and whose leftover is exact, the book's plainest case of whole turns plus a remainder. Those are the things a reader most needs to recognise before trusting any count.

Sources: the Cascadia paper (concept DOI 10.5281/zenodo.22076451; read in full); the Field Guide, deposited version 5.0, Section 9 (Case Four), read in draft v4.6 (concept DOI 10.5281/zenodo.21811588); the whale deposit, *Comparative Biosonar & Navigation* (concept DOI 10.5281/zenodo.20502057, resolved from Zenodo on 29 September; read in full); the sky paper (concept DOI 10.5281/zenodo.21331651; read in full); Delplace, Marston and Venaille 2017 (read in full; its quotations as the reading notes, section 11, carry them with line numbers); Suzuki and Measday, not read, for the capture fractions; for the planets, the NASA planetary fact sheet and IAU constants (recalled), Misner, Thorne and Wheeler (textbook, to check), and the three measured values as recalled, to be confirmed against their records. Identifiers: the reference list.

Runs: the Cascadia paper's PACHINKO_SUITE 20260821_191042, WINDING_BUDGET 20260821_190021 and WATERFALL_MAP 20260821_200450, as deposited; PLANET_FRAME_HOLONOMY 20260929_103728 (planet_frame_holonomy.py version 1, pure mathematics), the one run made for this section. Every Cascadia number is quoted from the paper or its runs' reports as printed; every muon number from the Field Guide crate's tool outputs; every whale number and quotation from the deposit's own text; every equatorial-wave number and quotation from the paper's text. Also drawn on: STATION_RING 20260819_174258, TREMOR_DRIFT 20260819_060144.

Part Two: the media, one card each

The Latch Atlas. Part Two prices the whole number in thirteen media. Each medium gets one card.

In the Amber picture of the book's opening, each medium is one shadow of the same pattern, and each card reads the pattern in its shadow. Each card puts the opening chapter's four questions to one medium. Then, in the same order every time, it tells the medium's whole number, what the number holds, what guards it, how it leaks, how much of its energy comes back instead of being spent, how the number is written, what the medium shares with the others, and where else in the book it appears. The cards follow the map's order: six latches first, then plasma, which splits, and the power grid, a powered lock; then three waves that carry labels, then two labels from the book's own models.

Every number on a card is its price-table column's own. The nine columns of the first build are kept unchanged as an appendix, with the four added since, for anyone who wants every cell. The four newer cards run longer than the first nine, as their media asked. What is still open for each medium, not yet priced, read or checked, is gathered at the end of this part.

The superconducting ring

A latch, the strongest in the table.

the four questions	answer
Does something go all the way around?	Yes: the phase of the superconductor’s shared quantum state, around the ring.
Is there a whole number only a cut can change?	Yes: the flux through the hole, in whole flux quanta.
Does it hold energy?	Yes, though little per rung: about one jiggle of heat at 4.2 K in the example ring.
Is the wall above the jiggle of heat?	Yes: about 10^4 jiggles even in a wire 20 nm across, and in a thick ring a current that has not measurably decayed.

The number. A superconducting ring holds the magnetic flux through its hole at a whole number of flux quanta. One quantum is $h/2e$, $2.067833848 \times 10^{-15}$ weber, and it is exact. The staircase of whole quanta was first measured in 1961.

What it holds. Very little per rung. The energy is $n^2\Phi_0^2/2L$: in words, it grows as the square of the whole number and falls as the ring’s inductance grows, so a bigger ring stores less per rung. One quantum in a 1 cm ring of 0.5 mm wire, an example chosen for the table, stores 5.5×10^{-23} J, or 0.345 meV, which is 0.954 of the thermal jiggle at liquid-helium temperature. Doubling the ring from 1 cm to 2 cm radius drops that to 0.141 meV. But rungs add up: the CMS detector’s magnet at CERN was designed for a stored energy of 2.6 GJ.

What guards it. Not the store but the wall. In a wire thinner than about one coherence length, the count changes by a phase slip: the superconducting order must vanish across the whole cross-section of the wire, over about one coherence length. The price is the energy it takes to melt the order in that one slice. For a niobium wire of 10 nm radius it is about 4.5 eV, some 1.3×10^4 jiggles at 4.2 K, an order of magnitude. A thick wire changes its count by a lower door, a vortex crossing it, and its wall is shown to be high not by a formula but by its current, which does not measurably decay.

How it leaks. By phase slips in a thin wire, and by a vortex crossing a thick one. In a thick ring the current does not measurably decay: a niobium-zirconium solenoid’s decay time “must be greater than 100,000 years”, as two secondary sources report it. Magnets for magnetic resonance, run in the same closed mode, lose typically 3×10^{-9} of their field per hour. Warm the ring above its transition and the latch is gone entirely.

How much comes back. A persistent current is steady, so nothing sloshes: no share of its energy goes out and comes back each cycle. Driven with alternating fields, a superconductor is the purest slosh in the table. A hollow niobium cavity reaches a quality factor of 10^{11} , so at an assumed 1.3 GHz only 6.3×10^{-11} of its stored energy is lost per cycle and it rings for 12.2 s, about 16 billion oscillations. Temperature sets the purity: at 2 K the quality factor falls to about 1.7×10^{10} , ringing for 2.0 s, and at 4.2 K to 3.7×10^8 , ringing for 46 ms. With a whole number of quanta trapped, the ring is a permanent magnet, and a magnet can run as the locked rotor of WRITE, W8.

How it is written. Cool it through its transition with a field on and it keeps the whole number nearest the applied flux; cool it fast with no field and it keeps a random one (WRITE, W4). Put a deliberate weak spot in it, a Josephson junction, and turned flux moves it one rung at a time, with hysteresis: the gearbox of W2 and the erase of W5. A junction also rings: held at a voltage it oscillates at exactly 483.5978 GHz per millivolt, and driven by a microwave tone its voltage sits on whole-number steps, 144.75 μ V apart at an assumed 70 GHz. NIST’s voltage standard locks about 300,000 junctions to one tone to give volts. That is a powered lock: the whole number holds only while the tone plays (W8).

What it shares. With the lattice ring, the same phase winding around a loop. What the superconductor adds is an order that must be destroyed to change the count, and that destruction is the wall.

Where else it appears. READ, sections 1, 2 and 3; WRITE, W2, W3, W4, W5 and W8.

Sources: the price table’s superconducting column; Deaver and Fairbank, Phys. Rev. Lett. 7, 43 (1961), at the level of the record; Langer and Ambegaokar (1967), as printed and conditioned by Joshi and Bhattacharyay, arXiv 1106.1720 (2011), read in full; P. Schmüser’s lectures in CERN-2004-008, for niobium’s constants, the CMS magnet, the cavity and the magnet decay; File and Mills (1963), through NASA SP-140 and CERN-2004-008; NIST’s 2013 news release.

Runs: PRICE_TABLE_ESTIMATES (the store, with the example ring's geometry assumed), WALL_ESTIMATES (the thin wire's wall, with zero-temperature constants for niobium) and SC_PHASE_ESTIMATES (the cavity's ringing times and the junction's pitch and steps).

The ring-shaped superfluid

A latch whose rungs are written by design at the bench.

the four questions	answer
Does something go all the way around?	Yes: the gas’s quantum phase winds around the ring.
Is there a whole number only a cut can change?	Yes: the units of circulation. To change them, a vortex must cross the band of gas.
Does it hold energy?	Yes, though each atom stores only a few ten-thousandths of a jiggle.
Is the wall above the jiggle of heat?	Yes: the wall between rungs holds anyway.

The number. Going once around the ring, the gas’s quantum phase winds a whole number of times, so the flow’s circulation comes in whole units. One unit is h/m , Planck’s constant divided by the mass of one atom: 1.7357×10^{-8} m²/s for sodium atoms and 9.9693×10^{-8} m²/s for helium-4. The count is read by releasing the gas and measuring the hole at its centre.

What it holds. Very little. The energy per atom is $n^2\hbar^2/2mR^2$. In words, it goes as the square of the whole number n divided by the square of the radius R , so a bigger ring holds each rung more cheaply. The flow speed, $n\hbar/mR$, falls as one over the radius. One laboratory ring of 4×10^5 sodium atoms, 19.5 μ m in mean radius, flows at 0.142 mm/s at one unit and carries 1.5×10^{-28} J in all. At an assumed 100 nK that is 111 thermal jiggles for the ring, but only 2.8×10^{-4} of a jiggle per atom. Nobody stores energy this way. This ring stores a count.

What guards it. Not the store but the wall. To change the winding, a vortex must cross the ring’s band of gas. A weak link is a stretch of that wall lowered on purpose, and loops of hysteresis, where the way up is not the way down, measure where it gives way. How high the wall stands in energy is not yet priced.

How it leaks. The flow has persisted without decay for up to 10 s, “limited only by experimental factors such as drift and trap lifetime” (from the paper’s abstract). Its lifetime is set by the atoms, not by the latch. Through a weak link, which lowers the wall on purpose, the winding changes one unit at a time.

How much comes back. Nothing sloshes. The flow is steady, so no share of its energy goes out and comes back each cycle instead of being spent. A driven ring is not yet priced.

How it is written. Three pens, all demonstrated. Twisted light, from Laguerre–Gaussian photons, hands each atom one unit, $\hbar$, of orbital angular momentum: a label written into matter as a latch. A weak link rotated around the ring stirs it up or down one rung, with hysteresis. And a fast freeze leaves random windings: supercurrents created “through a temperature quench of the normal-to-superfluid phase transition”, their distribution measured “for different quench times” (from the paper’s abstract; WRITE, W4).

What it shares. With the superconducting ring, a phase winding held by an order, with a weak link as the designed wall. With water, circulation, but here in whole units instead of as an amount of any size. It keeps its own place in the table.

Where else it appears. READ, sections 1 and 2; WRITE, W2, W4, W5 and W7.

Sources: the price table’s column for this medium; Eckel et al., Nature, 2014 (the ring’s size and atom count, the read by release, the rotating weak link and its hysteresis), at abstract level; Ryu et al., Phys. Rev. Lett. 99, 260401, 2007 (the persistent flow and the twisted-light pen), at abstract level; Corman et al., Phys. Rev. Lett. 113, 135302, 2014 (the fast quench), at abstract level.

Runs: PRICE_TABLE_ESTIMATES (the units of circulation, the flow speed and the store, at an assumed 100 nK).

A closed circle of DNA

A latch, and the only one in the table that nature writes, reads and spends on purpose.

the four questions	answer
Does something go all the way around?	Yes: each strand of the closed circle winds around the other.
Is there a whole number only a cut can change?	Yes: the linking number. Nothing short of breaking a strand changes it.
Does it hold energy?	Yes, about 185 thermal jiggles in a typical plasmid.
Is the wall above the jiggle of heat?	Yes: the wall is a covalent bond.

The number. In a closed circle of DNA (a plasmid, a bacterial chromosome or a loop held shut by proteins), one strand winds around the other a whole number of times: the linking number. It is exactly the Gauss linking integral the book's reader computes. In Călugăreanu's form it is twist plus writhe, two halves that each change smoothly and add to a whole number. On a trefoil-shaped ribbon the reader finds linking -4.0000 from writhe -3.5123 plus twist -0.4877 , as reported in a summary of the Field Guide.

What it holds. The free energy of supercoiling grows as the square of the extra linking number and falls with the circle's length. In words, per unit, a small circle is expensive to wind and a long one cheap. Living cells tend to hold the density of extra turns fixed instead, and then the next unit costs about 13 thermal jiggles at any length (12.76, 12.98 and 12.65 at 2,500, 5,000 and 10,000 base pairs), while the total grows with length. As for the lattice ring, the scaling depends on what is held fixed.

A 5,000 base-pair plasmid at the typical bacterial density of extra turns, -0.06 , is 29 turns short of relaxed and stores about 7.6×10^{-19} J, or 4.75 eV: some 185 thermal jiggles at room temperature. That density and the stiffness constant used, 1,100, are textbook values not yet checked. Small per molecule, it is still the first store in the table that sits comfortably above the room-temperature jiggle.

What guards it. A covalent bond in the backbone. Heat alone breaks it only extraordinarily slowly, and how slowly is not yet priced. The wall is complete against the jiggle and open to an enzyme.

How it leaks. A sealed circle does not leak. But a single nick lets the strands swivel, and the stored twist is gone at once, so the leak is the rate of strand breaks plus the enzymes that cut on purpose.

How much comes back. The share of the energy that comes back each cycle instead of being spent is not yet priced for DNA.

How it is written. By enzymes, in exact whole steps. Gyrase cuts both strands, passes a segment through and reseals. In the words of a 1979 paper, "both the introduction and the removal of supertwists by DNA gyrase change the linking number of DNA in steps of two". Gyrase spends chemical fuel to write rungs downward, a powered pen, and each step of two costs about 26 thermal jiggles at the typical density. Type IA enzymes change the linking number by exactly one; type IB, which let the cut strand swivel, by one or more whole turns per cut (WRITE, W1, from a reference summary). And the cell spends what it writes: holding the linking number below its relaxed value makes the strands easier to open where a gene must be read (WRITE, W6, from a reference summary).

What it shares. With the bare filament and the tube, the same Gauss linking number of two closed curves. What DNA adds, and the bare filament lacks, is a body: a double helix with a stiffness that makes extra turns cost energy, and a backbone that forms the wall. It keeps its own place in the table.

Where else it appears. READ, sections 1, 2, 3 and 7; WRITE, W1, W2, W3, W6 and W8.

Sources: the price table's DNA column; Brown and Cozzarelli, Science 206, 1979, text read; the trefoil ribbon from the Field Guide's DNA rows as digested, its source run not re-checked; Schroeder, Wolfenden and co-workers, PNAS 2006, record located, abstract not read; the type I steps from WRITE, W1, and the strands' easier opening from WRITE, W6, both checked against reference summaries.

Runs: DNA_ESTIMATES (the cost of the next unit, the store and the cost of gyrase's step, with $K = 1,100$ and the -0.06 density both textbook values not yet checked).

Magnetic skyrmions: whirls in a magnet

A latch: written, erased and read one whirl at a time in a cold film, and held at room temperature by the edge of a small disk.

the four questions	answer
Does something go all the way around?	Yes: the rim of a whirl, where the arrows of magnetisation all point the same way, closes the patch.
Is there a whole number only a cut can change?	Yes: how many times the arrows cover the sphere of all directions. Its sign is set by the turning and the core, and by which way is called up.
Does it hold energy?	Yes, in the whirl's bend and the forces that size it; not priced in real units.
Is the wall above the jiggle of heat?	Yes in one film at 4.2 K, where the wall is tens of meV, as reported, so a year's hold needs about 19 K. At room temperature the edge of a small disk holds whirls in place; that wall is not priced.

The number. A skyrmion is a small whirl in a magnet: the arrows point down at its core, turn over as you go out, and point up at its rim. Walk the patch and add up the solid angle the arrows sweep: it comes out in whole spheres of 4π . The book's read adds it up triangle by triangle and gets -1.000000000000 for whirls of radius 3, 6 and 12 grid spacings, where the smooth formula on the same grid gives -0.669294 , -0.905854 and -0.975672 . The count is whole only when it is read as a sum of closed pieces. In the first lattice ever seen, in MnSi, the total "per two-dimensional unit cell is quantized and adds up to -1 ", while the density inside the cell swings between about $+0.5$ and -2 : only the loop's total is whole. The sign is minus the turning times the core's direction, in all 32 whirls tried; turning every arrow by the same angle never changes it. It is also a convention: one paper gives -1 and another $+1$ to the same kind of whirl, core against the field, with the same formula. A card must say which way is up. Paper A's alphabet adds bags: a large whirl holding N small opposite ones carries $N - 1$. Such bags stay put at room temperature, with no field, in small magnetic disks. Circular dichroism in scattered X-rays reads the count directly but sees nothing for even numbers, so a picture that counts the whirls is the safer read for bags. Neutrons see a lattice's spots but not the phases that decide the count; pictures of the whirls themselves, taken in thin films and even inside a bulk crystal, are what the book's loop walk needs.

What it holds. With the bend alone, a smooth whirl costs 19.3322 at any size, above the least a count of one can cost, $4\pi = 12.5664$. So the bend gives it no preferred size; other forces do. In MnSi the whirls are built from spirals with a period of about 190 Å, some 41.7 atoms per turn; in a palladium-on-iron film each covers about 270 surface atoms. In real units the store is not yet priced.

What guards it. In a smooth magnet, the magnetisation would have to shrink to zero somewhere: "it would cost a large energy to suppress M locally to zero". On atoms, "strict topological protection does not exist. Instead, the states are separated by a finite energy barrier." The book's model, with the bend alone, has no barrier at all: a whirl shrinks downhill until the grid takes its count, at a radius of about 1.40 spacings. So the wall comes from forces the model leaves out. In the film where single whirls were first written, the same group reports it at tens of meV, as reported in a summary, so a year's hold needs about 19 K, "liquid He only".

How it leaks. Temperature decides. In the palladium-on-iron film at 4.2 K, with the tip quiet the switching rate "approaches zero", and a whirl stays. At 8 K the film simply follows the field. In MnSi the lattice lives only in a small pocket of field just below 29.5 K and keeps no memory of how it got there: a phase, not a store. Radiation is not yet priced for whirls: in the magnetic memory already in service, ionizing rays leave a stored bit alone and heavy ions break it only by damage, and what a single strike does to a stored whirl has not been found.

How much comes back. The share of energy that comes back each cycle instead of being spent is not yet priced.

How it is written. With a microscope tip, one whirl at a time. Four whirls at four flaws in the film were erased one by one and written back in a different order: a four-site register. What writes is the energy of the tip's electrons, not heat. At about 300 mV a switch comes once every 15 s; by 900 mV it is about 98 times faster; each electron has only about a 3.6×10^{-14} to 3.5×10^{-12} chance of flipping it. Which way it goes is a tilt: the field tips two wells, the whirl and the plain film, and reversing the current shifts the balance by about 100 mT. The pen writes both ways where the tilt is near zero. At low voltage the same tip reads without writing. A dose of ions is a pen too: gallium ions lower the cost of making whirls by up to 26 per cent and double how many form. At room temperature, bags in small disks are written on purpose, moved one at a time with a magnetic tip, and stepped down to a lower count by a field.

What it shares. The liquid crystal's hedgehog carries the same count, the solid angle over a sphere of directions. The superconducting ring shares the wall's kind: an order that must be destroyed to change the count. And the book's read carries over: it holds against noise on every arrow up to about 0.6 radians, reads from scattered stations spaced within a third of the whirl's radius, and is rescued by smoothing among near neighbours but erased by smoothing among look-alikes, which outvote the rare core.

Where else it appears. WRITE, W3 (a wall guards one path, where the Field Guide's example is the skyrmion) and W4 (temperature as the write-enable); READ, sections 2 (the solid angle) and 5 (reading through noise).

Sources: the price table's skyrmion column (new). Mühlbauer and co-workers 2009 and Romming and co-workers 2013, both read in full; Bear's Paper A, read in full. From the prior-art pass: Hagemeister and co-workers 2015, Zhang, van der Laan and Hesjedal 2017, Foster and co-workers 2019, Liu and co-workers 2025, Yu and co-workers 2010, Henderson and co-workers 2023 and the MRAM studies, read through a page-summary tool; de Jong and co-workers 2023, from its abstract.

Runs: SKYRMION_COUNT_CHECK 20260926_192130 and GROUPING_READ_CHECK 20260926_195247, pure mathematics.

Liquid crystals: defect lines and points in a nematic

A latch whose wall is not yet priced: halves in a thin layer, a parity in the bulk.

the four questions	answer
Does something go all the way around?	Yes: a loop round a defect line, or a sphere round a defect point.
Is there a whole number only a cut can change?	Yes, of several kinds: halves in a thin layer, even-or-odd for lines in the bulk, a whole number at points, thirds for the strips in knots.
Does it hold energy?	Yes: a defect line pulls like a stretched string, tens of thousands of jiggles for each micrometre. The bend of the texture is not priced.
Is the wall above the jiggle of heat?	Not priced. Only a cut changes the count, and neither paper says how high that wall stands.

The number. In a nematic liquid crystal the long molecules line up along a shared direction, but that direction has no head and no tail: it is “properly a line field, rather than a vector field”. Walk a loop round a defect and add up how far the direction turns. In a thin layer the answer comes in halves, because a line turned half a turn is the same line. You can see it between crossed polarisers: the strength is “equal to 1/4 times the number of brushes”, the dark bands that meet at the defect. The book’s run read seven strengths, from $-3/2$ to $+3/2$, right every time, and halves added like whole numbers: $+1/2$ and $-1/2$ make 0, and two $+1/2$ make $+1$. That is the textbook’s classification of a thin layer, reproduced by computer.

In the bulk, where the direction can tip out of the plane, the count changes kind. For lines, “no distinction can be made any longer between positive and negative strength defects”: a line is there or not, and two halves make nothing. A whole number survives at points. A hedgehog, a point where the direction sprays outward, carries the solid angle its surroundings sweep, counted in whole spheres; “Remarkably, it is an integer”. But carry a hedgehog once round a line defect and its sign flips, so reading it needs the path it took. And when defect lines are tied into knots, each line is a three-sided strip, “analogous to the well-known Moebius strip”, so its twist comes in thirds of a turn.

What it holds. The bend of the texture and the pull of the lines, which “tend to be as short as possible in order to reduce the total free energy”. Neither paper puts a number on either, but the pull has been measured since, as reported in a summary of the paper. In twisted cells like the ones the knots are tied in, a defect line pulls with a tension of 75 to 200 pN, more in thicker cells, from 6 to 27 μm . A tension is an energy per length, so each micrometre of line holds 7.5×10^{-17} to 2.0×10^{-16} J, about 1.85×10^4 to 4.94×10^4 times the jiggle of heat. The energy of the bent texture around the lines is still not priced.

What guards it. Only a cut changes the count: removing a pair of defects needs “surgery”, a cut and a regluing. How high that wall stands against the jiggle of heat is not priced.

How it leaks. After a quench, defects appear all at once, then meet their opposites and vanish, “a coarsening of defects too rapid for us to photograph”. A pair can vanish without changing the count of a loop round both. The knots are not reported to leak; they hold until the laser rewires them.

How much comes back. The share of energy that comes back each cycle instead of being spent is not yet priced.

How it is written. Three pens. Cooling throws down many defects at once. Shape: a small particle whose surface holds the molecules pointing outward is itself a hedgehog, and it forces “a companion defect in the liquid crystal”, a point or a ring round the particle. And light: laser tweezers melt the order at one spot and let it re-form, cutting and rejoining the lines there. With spheres 4.72 μm across in a cell 5.5 μm thick, each place where two lines meet can be switched between exactly three states, one site at a time and back again. On a 4 by 4 array the 3^9 combinations give almost 40 different knots and links, the Borromean rings among them, built to order. The cell chooses which way they tie: its walls are rubbed at right angles to twist the liquid a quarter turn, knots form only with that twist, and a left-handed twist favours left-handed crossings.

What it shares. With READ’s loop walk, done with light and at half the margin. A headless line needs each step taken within a quarter turn, not half a turn, so noise of 0.3 radians on every sample already trips the read more than 1 in 100 times, where an arrow read holds to 0.6 or 0.7. With Berry’s two trips: an arrow carried round a $+1/2$ defect comes back reversed after one lap and home after two. And the hedgehog’s count is the skyrmion’s, the fifth domain’s, in another medium. It is struck as a latch if no regime is found in which a defect holds against the jiggle of heat.

Where else it appears. Not yet in READ or WRITE by name. Its read is READ section 1's walk with a quarter-turn step and READ section 5's half-turn limit, halved; its two trips are READ section 2's; its pens are W1's cut, pass and rejoin and W4's freeze and thaw.

Sources: the price table's liquid-crystal column (new). Alexander, Chen, Matsumoto and Kamien 2012, and Tkalec, Ravnik, Čopar, Žumer and Mušević 2011, both read in full. The line tension: Chen, Mandić, Slaughter, Tanaka, Kikkawa, Collings and Yodh 2024, read through a page-summary tool.

Runs: NEMATIC_HALF_READ_CHECK 20260926_053148, pure mathematics.

Water: vortex rings and linked vortex lines

A latch, slowly leaking.

the four questions	answer
Does something go all the way around?	Yes: vortex lines, closed and linked through one another.
Is there a whole number only a cut can change?	Yes: how many times the lines link. Only a reconnection changes it.
Does it hold energy?	Yes, slowly leaking.
Is the wall above the jiggle of heat?	Yes: undoing the link needs a reconnection, a large-scale event far above room-temperature jiggle.

The number. A vortex line is a line that water spins around. When two such lines close into loops and thread through each other, the whole number is how many times they link. The flow carries it in a quantity called helicity. It stays whole while the lines stay whole; only a reconnection changes it. A single ring carries no whole number at all, because its circulation, the swirl around it, is a magnitude in ordinary water, not a count. And water has a blind spot: it shows linking only, never the richer numbers that tell one knot from another.

What it holds. Held at a fixed circulation, a ring’s energy rises with its radius: $E = \frac{1}{2}\rho\Gamma^2R[\ln(8R/a) - 7/4]$, where ρ is the water’s density, Γ the circulation, R the ring’s radius and a its core’s radius. In words: the energy grows in step with the ring’s size, times a slowly growing logarithm of how thin its core is. A ring of the size made on a lab bench carries about 0.38 to 0.40 mJ, by two estimates from our model (0.377 and 0.403 mJ). That is a single ring, which holds no whole number; what linked lines store is not yet priced.

What guards it. A cut, not a climb over heat. The linking changes only when two vortex lines are brought within about a core’s width of each other and viscosity reconnects them. The jiggle of water molecules cannot carry a whole vortex line there. Against heat the wall is effectively complete; against a deliberate collision it is low.

How it leaks. Through viscosity, the water’s internal friction. In our model, a vortex parked against a wall keeps about 99 per cent of its circulation over half a second at water’s viscosity, and loses 29 per cent of it at ten times that viscosity. A spin written onto a ring keeps 39 to 40 per cent of itself after nine ring radii in water, and its direction holds. Even the helicity decays slowly in a viscous fluid; topology never protects it outright. It has been measured in real water, every part of it: when a helical vortex is stretched or squeezed, its coiling and its twist trade off and the total stays roughly the same, and what viscosity takes is the twist. After the twist is gone, the rest holds steady.

How much comes back. In our model, a jellyfish bell swimming freely at the 1 mm scale runs at a power factor of -0.083 , a reactive share of 0.9965. That share is the part of the energy that comes back each cycle instead of being spent: ninety-nine per cent of what the bell does to the water comes straight back. Held on a leash, the actuator’s power factor is $+0.80$ to $+0.86$.

How it is written. Only by a reconnection, the cut described above: a deliberate collision can deliver it, heat cannot. In 2013 Kleckner and Irvine tied and untied knots in water, as their paper’s abstract reports.

What it shares. A closed loop forces a whole number. Water shares that with the lattice ring: linking here, winding there. In an ideal fluid, with no viscosity, it is settled physics that the linking sets a floor. But Arnold’s inequality, applied to swirling flow, bounds the strength of the swirl (the enstrophy), not directly the energy of motion. If the floor turns out to lie under the swirl’s strength alone, while the table still prices the energy of motion, water’s energy figures come out of the table.

Where else it appears. READ, sections 1, 2 and 7; WRITE, W1 and W3.

Sources: the price table’s water column. Helicity as the carrier of linking: Moffatt 1969. Knots tied and untied in water: Kleckner and Irvine 2013, abstract read. Helicity measured in water: Scheeler, van Rees, Kedia, Kleckner and Irvine 2017, read through a page-summary tool. The blind spot and the helicity’s slow decay: Field Guide row 147. The ring-energy formula: Saffman, as used in the Jellyfish paper §7C. Arnold’s inequality: the column’s verdict, as a pointer for a specialist.

Runs: the Jellyfish paper’s, as deposited: the store from its jelly_ring_signature run of 9 September, model output; the leak from its rung 3 and rung 4 runs of 9 September; the slosh from its runs of 17–18 September.

Plasma and the Sun: linked tubes of magnetic flux

A latch in a quiet plasma; in a turbulent one, the count is lost and the floor is kept.

the four questions	answer
Does something go all the way around?	Yes: closed tubes of magnetic flux.
Is there a whole number only a cut can change?	Yes in a quiet plasma: how many times the tubes link. In a turbulent plasma the separate counts cannot even be read.
Does it hold energy?	Yes: the total helicity holds the energy above a least-energy state.
Is the wall above the jiggle of heat?	There is no energy wall. The guard is a race between two speeds: fast relaxation and slow decay.

The number. In a plasma that conducts well, the field moves with the gas, so a tube of magnetic flux keeps its identity. When two such tubes close into loops and thread through each other, the whole number is how many times they link. The plasma carries it as magnetic helicity: two tubes with no helicity of their own, linked m times, carry twice m times the product of their fluxes, the count weighted by flux. In a perfect conductor, two loops linked n times “must remain linked n times during any subsequent plasma motion”. A tube’s own twist can be any amount, so it is a measure, not a count. On the Sun helicity is quoted in turns, helicity over flux squared, and an active region usually holds less than one. On the Sun’s surface the book reads a second whole number off real maps: how many times the sideways field turns round a loop, read from the rim alone. Every total landed within 5×10^{-15} of a whole number, and the 5 sunspot centres all read +1. This count is a label; it holds nothing back.

What it holds. Helicity ignores size and shape: a linked, twisted pair of tubes, with lengths from 9.03 to 54.17, read -7.125 however it was bent. It holds back energy: a plasma settles into the least energy its helicity allows, and that least energy is a true floor. A field with helicity cannot relax below an energy set by the helicity and the size of its container, and a knotted tube not below one set by its flux, its volume and how knotted it is. In the Beta-II spheromak, a laboratory plasma, the energy fell to about an eighth while the field kept its shape. On the Sun, active regions in one sample hold about 10^{40} to 10^{44} Mx² (maxwells squared) of helicity and about 10^{30} to 10^{33} erg of free energy, the energy above the calmest field. How much of that the helicity guards is not yet priced. On the surface, the electric current through a loop can be read from its rim alone. Round a sunspot standing alone it rises, then comes back, to a median 0.12 and 0.17 of its peak in two regions: the spot carries no net current, as others found.

What guards it. Not an energy barrier. With any resistance, lines break and reconnect, at a rate that “does not diminish as fast as η and may not diminish at all” as the resistance η shrinks. Reconnection changes the links but keeps the total: when two equal tubes reconnect once, one flux squared leaves their link and each tube gains half a twist. So the wall guards the total, not the separate counts. On the Sun’s surface the book’s rule held through two large flares: 1,589 loops in one region and 1,352 in another kept a strong rim for three hours, and none changed its count. The flare works where no rim holds. The instrument adds a wall: it cannot tell which way along its line the sideways field points, and where its choice flips, a count moves with nothing there.

How it leaks. In a laboratory plasma, slowly, through resistance, energy draining much faster than helicity. The Sun keeps helicity so well that it must throw it out about as fast as it is made, or its dynamo would choke. Eruptions flinging magnetic clouds into space are thought to carry most of it; how well the books balance is uncertain.

How much comes back. The share of energy that comes back each cycle instead of being spent is not yet priced.

How it is written. A voltage round a ring of plasma pours helicity in at twice the voltage times the flux; a gun’s voltage held one spheromak far longer than it would last alone, a powered lock. On the Sun, helicity comes up from below, with emerging flux and surface motions; two tubes turned once about each other write twice the product of their fluxes. The sign leans by hemisphere, negative in the north and positive in the south, in 60–75 % of active regions: a weak lean, its cause unsettled. S-shaped loops show the sign, though a view from above can mislead. A flare writes a measure, not a count: near the line where the upward field changes sign, the sideways field steps up and stays up, by 239.6 and 287.1 gauss in the two regions, close to the 258 ± 73 and 286 ± 59 others measured.

What it shares. With water, the same arithmetic: linked lines carried as helicity, with flux in place of swirl. The count is struck if no regime is found where closed tubes last long enough to be counted.

Where else it appears. WRITE, W1. Its count is read with the Gauss walk of READ, sections 1 and 6.

Sources: the price table's plasma column (new). Berger and Field 1984 and Taylor 1986, read in full. Pevtsov, Berger, Nindos, Norton and van Driel-Gesztelyi 2014, a review of the Sun's helicity, read in full, with the studies it reports. Moffatt and Ricca 1992, for the two-tube formula. The energy floor: Arnold's inequality and Freedman and He 1991, as Khesin 2005 states them. Taylor's conjecture of 1974, as both report it. Lu and co-workers 2019, read in part; Venkatakrishnan and Tiwari 2009 and Georgoulis, Titov and Mikić 2012, from their abstracts; Xiao, Prior and Yeates 2023, read in full.

Runs: FLUX_TUBE_HELICITY_CHECK 20260926_145656, a pure-mathematics model that reproduces Berger and Field's equation 16; for the Sun's surface, SUN_WINDING_READ, SUN_SERIES_READ, SUN_FLARE_DIFFERENCE, SUN_PIL_SHEAR, SUN_ZERO_CHECK and SUN_RIM_GROWTH on SDO/HMI SHARP maps of regions 11166 and 11158, courtesy of NASA/SDO and the AIA, EVE, and HMI science teams.

The power grid: a whole number round one loop of lines

A powered lock: the count holds while the grid runs in step, and it costs power to keep.

the four questions	answer
Does something go all the way around?	Yes: a closed loop of lines and buses.
Is there a whole number only a cut can change?	Yes: how many turns the voltage's phase makes round the loop.
Does it hold energy?	No, it spends it: the count drives a flow round the loop that raises the losses.
Is the wall above the jiggle of heat?	No heat jiggle enters. The guard is a saddle in the grid's own dynamics, held while the grid stays in step.

The number. Every bus in an AC grid has one voltage, swinging round like the hand of a clock. Walk round a loop of lines, adding up how far the hand turns from each bus to the next, and you come back to a whole number of turns. That whole number is the latch. Each is its own steady state, with power circling the loop and nothing put in or taken out to drive it. In the book's model, a ring of equal lines holds every winding that needs less than a quarter turn per line, and no other: a ring of 12 holds five, from -2 to $+2$, and a ring of 18 holds nine. A published study of an 18-bus ring found "at most nine solutions". The count and its limit are in print too: loop currents "take only discrete values, as they are characterized by topological winding numbers", and a ring of n buses holds at most $2 \text{Int}[n/4] + 1$ stable states, twice the whole part of a quarter of its length, plus one. The book's 5, 9 and 11 for rings of 12, 18 and 24 sit at or under that bound, 7, 9 and 13, below it only where the book's strict rule leaves out a winding that sits exactly on a quarter turn per line. Its authors borrow words like "vortex" from physics, but they say plainly that "there is no quantumness in the power flow problem".

What it holds. Not energy held back. A longer ring holds more windings, about one either way for every four lines, and each carries less flow: winding 1 carries 0.5000 of a line's greatest power in a ring of 12 and 0.2588 in a ring of 24. The flow delivers nothing. Real grids lose 5–10 % of what they carry, and in the published models a winding raised the losses by almost 50 % in one case and doubled them in a model with the UK grid's layout. So the count spends power instead of storing it. Loop flows of up to 1 GW are on record in North America, but whether they are windings has not been shown.

What guards it. To change the count smoothly, the voltage would have to fall to zero somewhere on the loop. The barrier is a saddle in the grid's dynamics, which the authors call an "energy" in quotation marks: it steers the motion and stores no joules. A line can carry at most a quarter turn before the steady state stalls, and a loop can hold one weak line a little past its own quarter turn, because the other lines carry it. Real grids keep each line to "few tens of degrees", since a quarter turn is past the thermal limit. Bear, from his years with the equipment, says the same: sections are reconnected in phase, because a quarter turn out blows past the thermal load.

How it leaks. While the grid stays in step, it doesn't: a written winding held in the models with losses included and when conditions went back. Cut the extra supply and the count still holds; the whole grid shifts its frequency instead, which "would be intolerable in a real power grid". Knocks on every bus change the high windings first: in the model, winding 2 changed in 0.341 of the knocks, winding 1 in 0.006, and winding 0 in none. The count goes when the loop is opened.

How much comes back. The share of energy that comes back each cycle instead of being spent is not yet priced.

How it is written. The switch is the pen. Overloading a line until the count slips works in the model, one whole turn at a time, but it needs a quarter turn across one line, so the authors call it "of little relevance for power grid operation". The real pens are switching. A tripped line can leave a winding behind. And reclosing a line writes one more turn each time the angle across the open gap has passed an odd number of half turns. Near a full turn, the angle at the breaker looks small, because a single angle is known only within one turn, so the reclose looks safe and still writes a new winding. The relay that guards a reclose checks the angle inside a small window, 5 degrees either way in one guide, so a whole turn plus a few degrees passes it. A perfectly symmetric ring cannot be written at all; the authors add small random nudges to break the symmetry. The paper asks that grid security checks look for these loop flows too; none that does was found. The grid already owns a tilt: a phase-shifting transformer turns the voltage angle on one line to steer power, and is used against loop flows at the border of Germany, Poland and the Czech Republic. Whether anyone has used one to choose or clear a winding was not found.

What it shares. With the book's ring of phases and W8's lock: the same sine-shaped pull between neighbours, which the authors liken to the current between two superconductors, and the same quarter-turn stall. It is struck as a medium of record if real loop flows prove not to be windings.

Where else it appears. Not yet in READ or WRITE by name. Its count is READ section 1's walk, its write rule READ section 5's half-turn limit, its pens W1's cut, pass and rejoin and W2's weak spot, its lock and stall W8's.

Sources: the price table's power-grid column (new). Coletta, Delabays, Adagideli and Jacquod 2016, read in full. From the prior-art pass: Delabays, Coletta and Jacquod 2016, from its abstract; the synchronising guide and ENTSO-E's sheet on phase-shifting transformers, through a page-summary tool. Bear's field experience with grid equipment.

Runs: GRID_LOOP_FLOW_CHECK 20260926_053413, a stated model of a ring of equal lines.

Light in a ring

A label: the superconducting ring's twin.

the four questions	answer
Does something go all the way around?	Yes: the light's phase, once around the ring.
Is there a whole number only a cut can change?	Yes, while the ring is steady: how many whole wavelengths fit around it. But nothing guards it (see below).
Does it hold energy?	Yes, briefly: 0.79 eV per photon, for about 350 ns in the best integrated ring.
Is the wall above the jiggle of heat?	No wall. For weak light the count is kept by the ring's steadiness in time; anything that varies in time, or light strong enough to pass the ring's threshold, changes it freely.

The number. Light running around a ring must have its phase come back to itself after one lap. So only a whole number of wavelengths fits around, and the ring resonates only at those colours. That whole number is the mode number: about 70,758 wavelengths of 1,570 nm light in the Puckett ring, 11.787 mm in radius, counted with an assumed refractive index. Twisted light carries a second whole number, the turns its phase makes around the beam's own axis. Each photon carries that many units $\hbar$ of orbital angular momentum.

What it holds. At a fixed colour, the rungs grow in number with the ring's size, and their spacing shrinks as one over the radius. A ring 0.1 mm in radius has 600 rungs 1.32 meV apart, one of 1 mm has 6,003, and the Puckett ring has 70,758, 11.2 μ eV apart. Its measured spacing, 2.720 ± 0.019 GHz at 1,570 nm, sits beside 2.70 GHz from the assumed index. At a fixed whole number, each photon's energy halves each time the radius doubles: 0.789, 0.395 and 0.197 eV at 0.1, 0.2 and 0.4 mm, for 600 wavelengths around. As in the superconducting ring and the superfluid, a bigger ring holds each rung more cheaply. The store is set by the colour, one photon at a time: 0.790 eV for 1,570 nm light, about 30 thermal jiggles at room temperature. How much a ring holds depends on how many photons are pumped in, not on the whole number.

What guards it. No wall, for no order must be destroyed to change the count. The rungs sit 11.2 μ eV apart, 4.3×10^{-4} of the jiggle at room temperature, yet heat does not move a photon between them. In a ring that is steady in time, light cannot change its colour, so steadiness, not a barrier, keeps the count. Nor can heat make these photons: it puts a mean of 5×10^{-14} in the mode. That holds for weak light only. Anything that varies in time moves light between rungs with nothing to climb; the Puckett team calibrate their measurement with an electro-optic phase modulator that puts two sidebands, new colours, on their laser. And strong light moves itself. In a ring that loses light very slowly, milliwatts of input grow to as much as kilowatts inside. Above a threshold power the ring "spontaneously generates two sidebands through parametric oscillation", the start of a frequency comb, at as little as 5.4 mW in one thesis's rings.

How it leaks. Out of the ring, at the rate its linewidth, the width in colour of each resonance, shows. The Puckett ring's intrinsic quality factor, a measure of how slowly it loses light, is 422 ± 11 million, and its intrinsic linewidth is 453 kHz; a ring-down measurement gives 434 ± 9 million, in agreement. Its photons last 351.7 ns by the quality factor and 351.3 ns by the linewidth: one number read two ways, so the leak is priced exactly by the linewidth. In that time the light oscillates 67 million times, and the whole number lasts only as long as the photons. Crystal rings hold light longer: quality factors of 63 billion are reported for bulk crystalline resonators, and microresonators range "from 10^6 up to nearly 10^{11} ".

How much comes back. Nearly all. Each cycle the energy swings between the ring's electric and magnetic fields and returns; this slosh is the share that comes back each cycle instead of being spent. Only 1.5×10^{-8} is lost per cycle: 2π over the quality factor, which is one full turn of angle divided by the quality factor. Seen from the laser at exact resonance, the ring is a plain load, because the slosh stays inside it. It is a flywheel for light, the regulator among the rotors of WRITE, W8.

How it is written. By choosing the colour. A laser tuned to one resonance fills that mode, and retuning fills the next. The pen chooses the whole number, but the page does not hold it: turn the laser off and the count is gone within the photon lifetime. Twisted light's number is written by a hologram on a spatial light modulator, up to 20 turns of phase around the beam. A Dove prism turns a twist of l into its mirror, $-l$, the same twist wound the other way, with nothing to climb; a rotating one shifts the light's colour in step with the twist. In tightly focused beams the flip smears into a spread of twists centred on $-l$.

The path to a latch. Yes into a quantum fluid, no into an ordinary object. Twisted light has written one unit $\hbar$ per atom into a ring-shaped superfluid, which held it as a latch for up to 10 s; a year earlier, in 2006, such photons had written quantised rotation into the atoms of a condensate. The same light makes a trapped particle turn, and has driven 2 μm silica beads along a spiral “with continuous transfer of OAM”, orbital angular momentum. But a turning particle keeps a speed, not a whole number. The path to a latch needs a receiver whose phase must close on itself: a superfluid or a superconductor.

What it shares. With the superconducting ring and the superfluid: a phase that must come back to itself around a loop winds a whole number of times. Light lacks an order holding the count, and that absence makes it a label. It stays in the table as the label twin.

Where else it appears. READ, sections 1, 2 and 7; WRITE, W7.

Sources: the price table’s light column; Puckett et al., the ring and its Methods, text read; Shen’s doctoral thesis on frequency combs, lines 688–696 and 1202 of its text, found by search; Daria et al. 2011, text read; Allen and Padgett’s introduction to Twisted Photons (Wiley-VCH, 2011, chapter 1), read in part, for Allen, Beijersbergen, Spreeuw and Woerdman, Phys. Rev. A 45, 8185 (1992), whose original is not read; González, Molina-Terriza and Torres, Optics Express 14, 9093 (2006), read in part; Grudinin, Ilchenko and Maleki, Phys. Rev. A 74, 063806 (2006), as cited by Puckett et al.; He, Friese, Heckenberg and Rubinsztein-Dunlop, Phys. Rev. Lett. 75, 826 (1995), records located, and Andersen et al., Phys. Rev. Lett. 97, 170406 (2006), abstract read; Ryu et al. 2007, through the superfluid column.

Runs: LIGHT_ESTIMATES, with an assumed refractive index.

Radio with a twist

A label, the table’s noisiest.

the four questions	answer
Does something go all the way around?	Yes: the wave’s phase, on a circle around the beam.
Is there a whole number only a cut can change?	Yes, in the beam as it travels: the twist. At the source it changes with nothing to climb.
Does it hold energy?	No: it carries energy past and holds none.
Is the wall above the jiggle of heat?	No wall: heat blurs the read but has no rung to climb.

The number. The twist: how many times the radio wave’s phase winds around the beam’s axis in one lap around it, the same whole number twisted light carries. The idea for radio was set out in 2007. In a test reported in 2012, two twisted radio waves in the 2.4 GHz band were sent 442 m across Venice, “from a lighthouse on San Giorgio Island to a satellite dish on a balcony of Palazzo Ducale”, as the publisher’s press release tells it. The test’s claim of extra channels on one frequency drew published comments.

What it holds. Nothing: the wave carries its energy past. The dial is flat. At a given frequency every twist costs the same energy per photon; the twist sets the angular momentum each photon carries, not its energy. A whole number whose rungs all cost the same is a label in its purest form. Longer waves carry less per photon: 9.93 µeV at 2.4 GHz, and 0.414 µeV at 100 MHz, a frequency picked for comparison. Each photon is tiny beside heat. At 2.4 GHz it is 3.84×10^{-4} of the thermal jiggle at room temperature, and heat alone puts about 2,600 photons into each mode at room temperature, 36 at 4.2 K and 0.003 at 0.02 K. So the count lives in the pattern of a strong wave, not in single photons.

What guards it. Nothing. A re-phased ring of antennas or a twisted reflector changes the count with nothing to climb. Heat does not change a strong beam’s twist; it blurs the read.

How it leaks. Nothing drains the twist in empty space. Heat enters as noise in the read, not as a drain on the count, and the book’s noise checks say how to read through it: average neighbouring samples, then compare two loops.

How much comes back. An antenna is the device that turns slosh, the share of the energy that comes back each cycle instead of being spent, into travel. How much sloshes and how much travels is not yet priced.

How it is written. By a ring of antennas fed with the phase stepped around the ring. In a ring of eight, the step between neighbours is 45° for a twist of 1, 90° for 2 and 135° for 3. At 180°, a twist of 4, the twist and its mirror cannot be told apart, and a twist of 5, at 225° per step, reads back as −3. So a ring writes and reads whole numbers only below half its number of antennas: the same half-turn limit that makes a noisy loop walk go wrong.

The path to a latch. Yes, through a superconductor, as a powered lock. A microwave tone is what locks NIST’s 300,000 junctions onto whole-number voltage steps, as the superconducting ring’s card tells. There the whole number is written by the tone’s frequency, not by any twist, and it holds only while the tone plays. Cold opens the way: at 0.02 K heat stops filling a 2.4 GHz mode, leaving only 0.003 photons in it.

What it shares. With light: the same phase winding around the beam’s axis, at a different scale. It stands in the table beside light.

Where else it appears. READ, sections 1, 2 and 5; WRITE, W7.

Sources: the price table’s radio column; IOP Publishing’s press release on Tamburini et al., New J. Phys. 14, 033001 (2012), the paper’s own page having refused the fetch tool; Thidé et al., Phys. Rev. Lett. 99, 087701 (2007), record located; the published comments arXiv 1210.5365 and 1302.4121, records located, not read; NIST’s junctions as the superconducting column records them.

Runs: RADIO_ESTIMATES and LOOP_COUNT_NOISE_CHECKS (check B).

Sound with a twist

A label that leaks into the heat of its own medium.

the four questions	answer
Does something go all the way around?	Yes: the sound's phase, on a circle around the beam.
Is there a whole number only a cut can change?	Yes, in the beam as it travels: the twist. At the source it changes with nothing to climb.
Does it hold energy?	No: it carries energy and turning past and holds none.
Is the wall above the jiggle of heat?	No wall.

The number. The twist: how many times the sound's phase winds around the beam's axis in one lap around it. In 2008 it was made in open air at 1,300 Hz, sound 26.4 cm long, with four loudspeakers in a circle for a twist of 1 and eight for a twist of 2, as reported in a summary of the experiment.

What it holds. Nothing: the wave carries energy and turning past. As for radio, the dial is flat in energy: every twist costs the same per quantum, the smallest packet of sound. What the twist sets is the turning each joule carries: the twist divided by the angular frequency, the rate at which the wave's phase turns. In words, more twist means more turning per joule, and a higher pitch means less. At 1,300 Hz that is 1.22×10^{-4} newton-metre-seconds per joule for a twist of 1, and twice that for a twist of 2. Light follows the same ratio, and it has been tested for sound. The count is priced in turning, not in energy. A quantum of 1,300 Hz sound is 2.08×10^{-10} of the thermal jiggle at room temperature, and heat puts about five billion quanta into each mode, so sound is wholly classical here and its count lives in the organised pattern. Heat in matter is itself disordered vibration of the same molecules whose ordered motion is sound.

What guards it. Nothing. Re-phase the ring of sources and the twist changes.

How it leaks. Into the heat of its own medium. How fast, in air or in water, is not yet priced.

How much comes back. Slosh, the share of the energy that comes back each cycle instead of being spent, is not yet priced for sound. A standing sound wave sloshes between pressure and motion, which is a ringing bell; a travelling one carries its energy away.

How it is written. By a ring of sources stepping the phase from one to the next. The 2008 experiment kept a quarter turn, 90° , between neighbours in both set-ups: four sources for a twist of 1, eight for a twist of 2. That is well clear of the half-turn limit; four sources for a twist of 2 would need 180° , where the twist and its mirror cannot be told apart.

The path to a latch. None found. In the 2008 test the twist passed into matter as a turn: foam-covered acrylic disks hung on a fibre twisted round to an angle where the sound's torque balanced the fibre, as the same summary reports. A torque and an angle are both magnitudes, amounts that can take any size. Sound can also stir a flow in ordinary water, but circulation there is a magnitude too.

What it shares. With light and radio: the same winding of phase around an axis, carried by a material medium instead of a field. It stays in the table.

Where else it appears. READ, sections 1, 2 and 5; WRITE, W7.

Sources: the price table's sound column; Volke-Sepúlveda et al. 2008, read through a summary; Allen and Padgett's introduction to Twisted Photons, read in part, for the same ratio in light; Demore et al., Phys. Rev. Lett. 108, 194301 (2012), record located; textbook physics for heat as disordered vibration and for standing and travelling waves.

Runs: SOUND_ESTIMATES.

The lattice ring (a model)

A label in this model.

the four questions	answer
Does something go all the way around?	Yes: the phase of each electron state winds around the ring.
Is there a whole number only a cut can change?	Yes: each electron state’s winding is a whole number, though the ring passes freely from one to the next as the flux turns.
Does it hold energy?	No, in this model.
Is the wall above the jiggle of heat?	Not reached.

The number. The lattice ring is the model the book calls the magnetic forest: electrons carried around a ring, with magnetic flux applied through it. Each electron state around the ring has a whole-number winding, and the current repeats every whole flux quantum. The ring settles into whichever winding is cheapest for the flux applied, and it keeps no memory of how it got there. So it is a label, not a latch. By page one’s third question, it holds nothing back. It is settled physics that the current flowing on its own in an ordinary ring, one that is not a superconductor, is simply the ring’s lowest-energy state. The same ring becomes a latch in a superconductor, where a phase slip is the cut that changes the number.

What it holds. Nothing above its lowest-energy state. What the model does give cleanly is its dial. The current falls as one over the circumference: the smaller the ring, the bigger the current. Measured across a sixteenfold range of sizes, on a chart with both axes drawn on a logarithmic scale, the line falls with a slope of -1.0004 , which is what one over the size looks like. The current’s amplitude times N , the ring’s size, holds at 1.6733 in the model’s own units. That law has a limit. It holds when the temperature is scaled with the spacing between the ring’s energy levels. At a fixed temperature the slope is -7.96 instead, and at zero temperature the swing in energy goes as -2 . The dial depends on what is held fixed. There are no joules yet, because no material has been chosen.

What guards it. Nothing. As the applied flux turns, the ring passes from one winding to the next with no barrier, always sitting in whichever is cheapest. That is why it is a label.

How it leaks. It does not, in the model, which has nothing in it that wastes energy. How a real ring leaks is not yet priced.

How much comes back. The share of the energy that comes back each cycle instead of being spent has not been measured for this ring.

How it is written. There is no pen. Turning the applied flux moves the ring from one winding to the next, but nothing written stays: the ring keeps no memory of how it got there.

What it shares. The mechanism it shares is the phase winding around a closed loop. Its dial runs the opposite way from water’s: the smaller the ring, the bigger the current. That difference is exactly what the book’s guardrail is there to show; it is not a reason to strike the ring from the table.

Where else it appears. READ, sections 1, 2, 4 and 7.

Sources: the price table’s lattice-ring column; the fixed-temperature slope and the zero-temperature energy swing: the companion paper, Instruments for the Magnetic Forest, Appendix A4; the companion’s own wording: its line 299; the opposite dial as a guardrail: Field Guide row 532.

Runs: the dial: RING_LATTICE_READER 20260914_032117, the companion’s own, as deposited.

The bare filament

A label. The count is exact and the floor is absent.

the four questions	answer
Does something go all the way around?	Yes: two closed curves.
Is there a whole number only a cut can change?	Yes: the Gauss linking number of the two curves.
Does it hold energy?	No.
Is the wall above the jiggle of heat?	Not reached.

The number. The bare filament is a pair of closed curves with no thickness at all. Its whole number is the Gauss linking number: how many times one curve passes through the other. Drawn with 800 points, it is read to within 0.00001 of its whole number. It is exact as long as the curves do not cross. The bare filament is a label, not a latch: its count is exact, and it holds nothing back.

What it holds. Nothing is held against the whole number. At a fixed current the energy grows with size, so shrinking always lowers it. The pair linked six times falls from 17.51 to 6.81 at half size, and its linking does not change. There is no floor to scale.

What guards it. Nothing. Slide one ring out through the other's wire and the energy drifts smoothly, by 3.80 per cent along the whole path. It shows no spike where the strands pass through each other. At the two crossings the energy departs from a smooth curve by 0.00314 and 0.00243, less than the largest departure anywhere else on the path, 0.00611. This energy has neither a floor nor a wall.

How it leaks. Nothing is stored, so nothing leaks.

How much comes back. The share of the energy that comes back each cycle instead of being spent does not apply here: nothing moves.

How it is written. The count can change only where the curves cross, as they do when one ring slides out through the other's wire, and there nothing has to be climbed.

What it shares. It would share its mechanism with a tube of fixed flux and volume, but only once the curves are given thickness. As a bare filament it stays in the table on purpose, as the book's worked example of a count that holds nothing.

Where else it appears. READ, sections 1, 2, 3 and 7; WRITE, W1 and W3.

Sources: the price table's bare-filament column; the verdict: the Latch paper, version 2.

Runs: the count: LATCH_PROBE 20260923_191121; the dial: LATCH_FLOOR_CONFOUND_CHECK 20260923_194934, both the Latch paper's tools as run for the price table before the paper was deposited (the deposited record carries the same tools' runs made the next day, which print the same numbers, checked; see READ, section 7); the wall: WALL_ESTIMATES 20260923_224526.

Still open, medium by medium

What each card leaves out because it is not yet priced, read or checked. The list travels to the book's back matter.

The superconducting ring

- Doll and Näbauer (1961): record not checked.
- File and Mills (1963): the original letter is not read; its decay time is quoted through two secondary sources.
- The thick wire's real wall, a vortex crossing: not priced (Langer and Fisher 1967 for the three-dimensional case; Kogan's barrier for a strip).
- The cavity's quality factor taken as one over its surface resistance: textbook, to check.
- A wire ring driven by alternating flux loses through the same surface resistance; its own figure is not priced.
- That a superconducting ring cannot turn as an induction motor's rotor (its current runs a quarter cycle out of step, so the average torque is zero): textbook, to check.

The ring-shaped superfluid

- The wall's height in energy (WALL cell, HELD): to be read from Eckel et al.'s hysteresis loops. WRITE, W3 lists the same item as held (line 156).
- SLOSH for a driven ring (HELD): not priced.
- Related, held under the radio and sound columns and in WRITE, W7: whether twisted microwaves or a sound vortex have ever written a winding into a ring superfluid has not been searched.

A closed circle of DNA

- The typical bacterial density of extra turns (-0.06) and the stiffness $K = 1,100$ (both UNVERIFIED): the store, the cost of the next unit and the cost of gyrase's step all rest on them.
- The wall's height: the half-life of the backbone bond against uncatalysed hydrolysis (Schroeder, Wolfenden and co-workers, PNAS 2006; record located, abstract not read; UNVERIFIED).
- SLOSH (HELD, not priced): the twin-domain picture, in which reading a gene winds DNA ahead of the machine and unwinds it behind so that supercoiling travels in waves along a working chromosome, is to check.
- The trefoil ribbon's numbers come from the Field Guide digest; the source run is not re-checked.

Magnetic skyrmions: whirls in a magnet

- The verdict: drafted as a latch written and read one whirl at a time in a cold film and held at room temperature by a disk's edge, for Bear to settle once a magnetician confirms the reading.
- The wall: tens of meV in the cold film, as reported in a summary (Hagemeister and co-workers 2015), to check against the paper's text; its height in room-temperature materials, where bags are held by a disk's edge.
- The store: the whirl's energy in real units.
- The sign convention the card will use; and whether other dichroism reads share the X-ray scattering read's blindness to even numbers, and whether it reads the sign (for a magnetician).
- A single strike: whether one ionizing strike can switch a stored whirl. Single electrons' energy switches one in the cold film, heat ruled out; an ion dose is a known pen; MRAM, the magnetic memory in service, survives ionizing radiation; single strikes on stored whirls were not found (for a radiation-effects specialist).
- The slosh: not priced.

Liquid crystals: defect lines and points in a nematic

- The verdict: drafted as a latch whose wall is not yet priced, for Bear to settle once a soft-matter physicist confirms the reading.
- The store: the line tension (75 to 200 pN, Chen and co-workers 2024) is read through a summary and is to be confirmed against the paper's text; the energy of the bent texture in real units is not priced (for a soft-matter physicist).
- The wall: the height of the barrier to cutting a defect line, against the jiggle of heat.
- The leak: the rate of coarsening; whether the knots ever come undone on their own.
- The slosh: not priced.

Water: vortex rings and linked vortex lines

- The wall: the reconnection timescale, to be read from Kleckner and Irvine.
- The energy question, for one specialist: does the helicity floor in a real fluid sit under the enstrophy, the kinetic energy, or both? The strike test turns on the answer.
- The store for linked lines: the 0.38 to 0.40 mJ figure is for a single bench ring, which holds no whole number; a store for linked vortex lines needs a number from linked-ring experiments (Kleckner and Irvine, or the reconnection measurements that followed), not yet read.
- The ring-energy formula is taken as the Jellyfish paper states it and is still to check against Saffman's text.

Plasma and the Sun: linked tubes of magnetic flux

- The verdict: drafted as split (a latch in a quiet plasma; the count lost and the floor kept in a turbulent one), for Bear to settle once a plasma physicist confirms the reading of Taylor's review; the Sun's record fits the split's second half.
- The store: how much of an active region's free energy the helicity itself guards (the book's reading, at most all of it, for a specialist); the energy of a single flux rope. (Whether the floor stands under the magnetic energy itself is settled: it does, by Arnold's inequality and Freedman and He's bound.)
- The leak on the Sun: Berger's 1984 limits, reported through the review, not read.
- The slosh: not priced.
- The Sun's helicity numbers are mostly the review's reports of other studies; none of those studies is read. (The Sun's surface numbers are the book's own runs.)
- The Sun's surface: whether the zero of the sideways field that sits on region 11166's polarity line for three hours is the foot of the null point above or the end of a bald patch (for a solar physicist); whether a spot's return to zero, read through rims where the instrument is seldom confident of the field's direction, survives other ways of resolving that direction (for a specialist).
- The Sun's surface currents: the upward current of one sign in both polarities that Part Four reports behaves like the instrument's own turning of the field, about 3 degrees near disk centre, fading toward the limb without changing sign. The follow-on paper on the Sun narrows what makes it to the inversion's reading of the polarised light (a mis-modelled magneto-optical term gives its size and its polarity flip but not its shape) and leaves the rest with the HMI team.

The power grid: a whole number round one loop of lines

- The verdict: drafted as a powered lock, for Bear to settle once a power-systems engineer confirms the reading.
- Whether the recorded loop flows, up to a gigawatt, are windings at all (for a power-systems engineer).
- Whether anything in grid protection can see a whole turn across an open breaker (a synchronism check's window is 5 degrees either way in one generator guide, so the book's reading is that it cannot; line reclosing settings not read; for a protection engineer).
- Whether a winding can be read from measured phasors in a real grid, and whether a phase-shifting transformer could choose or clear one (not found; for a power-systems engineer).
- The store: the losses of a winding in a real grid, in real units.
- The slosh: not priced.
- Whether windings can be written without switching (open in the paper, 2016).

Light in a ring

- The ring's own index: the mode numbers use an assumed refractive index. The measured spacing between rungs agrees with it, but the mode number itself needs the ring's effective index, given in the paper's supplementary information, not yet read.
- Higher-Q rings: the 63 billion quality factor for crystal rings is quoted only as Puckett et al. report it; the original (Grudin, Ilchenko and Maleki 2006) is not read, so no lifetime is worked out from it.
- The 1992 original on twisted light (Allen, Beijersbergen, Spreeuw and Woerdman) is quoted through Allen and Padgett's own introduction, not read itself.
- He et al. 1995: record located, not read.
- Barnett, Babiker and Padgett's 2017 review (Phil. Trans. R. Soc. A), supplied by Bear: on hand, only its abstract read.
- Shen's thesis: sampled by search, not read in full.

Radio with a twist

- The Venice paper itself: read only through its publisher's press release; the paper's own page refused the fetch tool.
- The two published comments on the Venice paper (arXiv 1210.5365 and 1302.4121) and Thidé et al. 2007: records located, not read.
- The beam spreads as it travels, so the ring needed to read the twist grows with distance: textbook, still to check.
- Antenna slosh: near an antenna most of the energy sloshes back and forth in the reactive near field, and far away it travels outward as radiation, which is the work done. Textbook, still to check, and no priced figure yet.
- That superconducting microwave circuits are run as cold as 0.02 K because heat stops filling the mode there: textbook, still to check.
- Whether a twisted microwave beam has written a winding into a superconducting ring or a ring superfluid: not searched.

Sound with a twist

- The 2008 experiment (Volke-Sepúlveda, Santillán and Boullosa): its values come from a summary of the authors' own copy, not a full reading.
- Demore et al. 2012: record located, not read.
- How viscosity and heat conduction turn ordered sound into disordered jiggle as it travels, faster at higher pitch: textbook, still to check.
- Absorption, priced: how fast air and water turn sound into heat, by pitch.
- Slop, priced: a figure for the share of sound's energy that comes back each cycle.
- Whether a sound vortex can write a whole number into a superfluid ring: not searched.

The lattice ring (a model)

- The lattice ring as a latch: needs the superconductor chapter, a ring whose winding is held by an order parameter.
- The store in real units: no joules until a material is chosen; needs the superconductor chapter's material.
- The leak of a real ring: by phase slips in a superconductor, or by losing its phase memory in a normal metal; needs the superconductor chapter and its phase-slip rate.
- The slosh: needs a run in which the flux varies in time.
- The companion's wording: it says "same latch, opposite dial" (its line 299); by page one's third question the lattice ring is a label, and the wording is flagged for the companion's next version.

The bare filament

- Nothing. No row of this column is held or waiting to be priced.

Part Three: WRITE

The Latch Atlas, reading draft. How the whole number is changed, where a medium lets you, and what it costs.

READ asked how you read a whole number off a loop. WRITE asks how you change it, where a medium lets you, and what it costs.

section	what it shows
W1. One act: cut, pass, rejoin	the only way to change a count; the two writes DNA's enzymes make, done on a ribbon, change it by exactly two and exactly one
W2. Rungs you choose: the gearbox	a ring with one weak spot has a resting place for each whole number near the knob's setting, and slips one rung at a time; the weak spot's strength is both the pen and the guard
W3. What a write costs: the wall and the store	the wall between rungs and the energy on each are different quantities and can differ enormously; the gap is the wall; a wall guards only the path it stands on, so magnet wire is stranded and twisted
W4. Freeze and thaw	temperature as the write-enable, at two temperatures (the rungs appear, then lock); heat and a ring frozen in pieces write one staircase, a bell curve sorted into rungs; in niobium the two routes lock together, and the measured slope is the open question; a freeze writes only into a closed loop and keeps its number round a hole
W5. Erasing is not writing backwards	the way down is not the way up, and the loop between them is the memory and its price; an erase is a second write, to the rung below, from the other side, and costs what the write cost; a write that does not keep still pays
W6. The same number, written two ways	twist and writhe trade while the count holds; bending trades and only a cut writes; a reader of the total is blind to the split, which a stretched DNA molecule's length can read
W7. Waves write labels, and a label can write a latch	a ring of emitters writes its feed folded at half a turn, and an even ring writes dark spokes at the fold; holograms and prisms write labels with no wall; twisted light set a ring superfluid turning, a label become a latch
W8. Engines, regulators and locks	two layers out of step turn a load as the sine of the offset; a lock holds up to a quarter-turn lag, the stall, and a crowd locks past a threshold; a powered lock is not a latch; a driven ring with a handedness is an engine that writes its sign; the watch and the DNA setpoint, judged by Bear, with his ratchet and his deadband: a deadband is W5's loop set on purpose, widened by a supervisor as W4's write-enable

W1. One act: cut, pass, rejoin

A walk around a loop can only count what the loop goes around, so the only way to change the count is to change what the loop goes around. Cut one strand, pass the other through the gap, and join the cut again. Physicists call it reconnection. Every pen in this part is that act, done by a fluid, a magnetic field, an enzyme or a weak spot in a superconductor.

Three things come out of this section. A count never changes by part of itself: the reader of READ already watched writes happen, and every one was a jump. The two writes that DNA's enzymes make can be done on a plain ribbon, and they change its count by exactly two and by exactly one. And the literature already names these writes, in fluids and magnetic fields, so this part walks them rather than claiming them.

What the reader already saw

READ watched three writes from the reader's side. In section 5 a whirl was moved across a loop of weather vanes: the count read 1 at every position inside, 1 with the whirl exactly on the walk, and 0 the moment it was outside, never a fraction. In section 7 one bare ring was slid through the other's wire: the count went from -1 to 0 to $+1$, and because a curve with no body has nothing to climb, the energy never spiked. And in section 8's read-only ground, where nature holds the pen, Cascadia's triangles almost never reverse in one step: of 13,569 departures from a charged week, only 129 went straight to the opposite sign, so a reversal almost always passes through zero first (Cascadia paper §10.6). A count moves in whole steps, and when it reverses it tends to pass through nothing on the way.

Two writes on one ribbon

A ribbon is a closed centre line with an edge running a short way beside it. Its linking number, how many times the edge winds around the centre line, is a whole number, and it is twist plus writhe, the two continuous halves of READ section 2 (Călugăreanu's theorem, named as Moffatt and Ricca argue it should be). Two writes can be made on it, and they are the two that DNA's enzymes make.

Two writes on one ribbon, the two that DNA's enzymes make




Figure W1.1. Two writes on one ribbon. (A) The centre line is a figure-eight lifted into space, so its two strands cross at the middle, one above the other; the upper strand is lowered through the lower one. The linking jumps from $+1$ to -1 as they pass. The writhe carries the jump, and the twist on either side shrinks toward zero at the passage. (B) The ribbon is cut, one end is turned through whole turns, and the ends are joined again: the linking changes by exactly the number of turns, the writhe not at all. The book's own pure-mathematics run, 1,600 points per curve, the edge 0.02 from the centre line.

Pass one strand through another. Lower the upper strand until it passes through the lower one. Before the passage the ribbon links $+1$, after it -1 : a jump of exactly 2, at every separation from 0.4 down to 0.03. The writhe carries the jump. Its change across the passage climbs from 0.93 at a separation of ± 0.4 to 1.91 at ± 0.03 , heading for 2, while the twist on either side shrinks toward zero, ± 0.0434 at the closest pair. The passage itself moves the writhe by 2 and the twist by nothing. This is the write DNA gyrase makes: it cuts both strands of the double helix, passes another stretch of helix through the gap and seals it, and it changes the linking number, in the words of the paper that found it, "in steps of two" (Brown and Cozzarelli 1979, read for the price table's DNA column).

Cut, turn the end, rejoin. Cut the ribbon, turn one free end through whole turns, and join it again. The centre line has not moved, so the writhe cannot change, and it does not: 0.5865 before and after, to six decimals. The linking changes by exactly the turns added: $+1$, $+2$, $+3$ and -1 turns give changes of $+1$, $+2$, $+3$ and -1 . This is the write of DNA's type IB enzymes, which cut one strand, let the cut end swivel and reseal it, changing the linking number by one or more whole turns at a time; their cousins of type IA pass the other strand through the break instead and change it by exactly one. A reference summary states both: type IA enzymes "change the linking number of a circular DNA strand by units of strictly 1", and type IB enzymes "by multiples of 1" (Wikipedia, "Type I topoisomerase", read 25 September, citing J. C. Wang, *Nat. Rev. Mol. Cell Biol.* 3, 430, 2002, not read; the enzymes' steps were checked against that summary in the WRITE checking pass of 25 September).

The literature's write heads

The same two writes are the fluid and magnetic literature's own. Moffatt and Ricca (1992) showed how the helicity of a tube is written by cutting: cut the tube, twist the free ends through a whole number of turns and reconnect, which is Dehn surgery and this section's cut-turn-rejoin; or switch a crossing, each switch adding or removing a fixed amount, which is this section's strand passage (as read for the Latch paper, version 2). Kleckner and Irvine (2013) tied linked and knotted vortices in water and watched them untie by reconnection (read at the level of the abstract). The Sun cuts and rejoins its own twisted magnetic ropes; that write belongs to the plasma chapter, and its sources are that chapter's to read.

What this section does not show

The ribbon's writes are made by hand, at a place and a moment we chose. A real medium decides for itself where and when to cut, and nothing here models that. The Latch paper's own open claim stands: "A link can be untied under a stated rule" remains open until a code carries a cut-and-rejoin rule of its own, and that rule would be the largest assumption any untying result carried. Nor does this section price anything. A ribbon of pure geometry, like the bare curve of READ section 7, writes for free; what a write costs in a medium with a body is section W3's question.

Sources: Moffatt and Ricca 1992 (read in full for the Latch paper); Brown and Cozzarelli 1979 (read for the price table's DNA column); Wang 2002, as a reference summary cites it (not read); Kleckner and Irvine 2013 (abstract); the Cascadia paper (its §10.6). Identifiers: the reference list.

Runs: WRITE_W1_FIGURES 20260924_174454 (write_w1_figures.py, pure mathematics); the READ runs named in the text; the Cascadia paper's WATERFALL_MAP 20260821_200450; the checking pass WRITE_CHECKS 20260925_234135. Also drawn on: B1_READ_CHECKS 20260923_233341, WALL_ESTIMATES 20260923_224526.

W2. Rungs you choose: the gearbox

Give a ring one deliberately weak spot and it becomes a ladder: a resting place for each whole number, a barrier between neighbours, and a knob that winds the ring until it slips exactly one rung. The whole number is the gear, the continuous turning between rungs is the spring wound between teeth, and the slip is the escapement.

Three things come out of this section. A superconducting ring closed by one weak spot, in its textbook model, has a well for each whole number near the knob's setting, and turning the knob moves it up or down one rung at a time, never two. How strong the weak spot is sets how far each rung holds and how many rungs survive once the knob is let go, so the pen and the guard are one setting. And the same gearbox has been built in a ring of cold atoms and is run every day by an enzyme, so this section walks it rather than claiming it.

A ring with one weak spot

The ring is superconducting, so the flux through its hole is held at a whole number of flux quanta ($h/2e$, the price table's superconducting column). Close the ring with one weak spot instead of solid wire, a junction that passes only a small current, and it becomes the device the field calls an rf SQUID. Its textbook energy has two parts. The ring's inductance acts as a spring pulling the flux through the ring toward the flux applied from outside, which is the knob. The weak spot adds a ripple with a dip at every whole number of flux quanta. Together they make a tilted washboard, and the ring rests in one of its dips. That is the standard textbook model of a superconducting ring closed by one Josephson junction, stated in full at the head of the run's tool, and the run is this book's own.

A ring with one weak spot: the textbook model, $\beta = 3$ in A and B




Figure W2.1. A ring with one weak spot, in the textbook model. (A) Energy against the flux through the ring at three settings of the knob; the dot is where the ring rests. At 0 there is one well; at 0.5 there are two, level with each other, and the ring stays in the first; at 0.74 the next rung is far lower, but the ring is still held on rung 0 by what is left of its well. (B) The flux through the ring as the knob is turned up to 2.5 flux quanta, down to -2.5 and back to 0: every slip is one rung, the way down is not the way up, and the dotted line is the rung with the lowest energy, which the ring does not follow. (C) Where rung 0 gives way against the strength of the weak spot; dots from the run, the line from the model. No measured constants and no random numbers; the knob turned in steps of 0.0005 flux quanta; heavy damping assumed.

With the weak spot of the figure (a strength of 3, defined below), the knob at zero leaves the ring one resting place, rung 0. At half a flux quantum there are two, rung 0 and rung 1, exactly level, with a wall between them. At 0.74 both wells still exist but rung 1 now lies far below rung 0, and the ring stays where it is. Its rung is always the nearest whole number to the flux through the ring: while the ring rests, that flux strays at most 0.2988 of a quantum from its whole number, never the half that would make the reading ambiguous.

Turning the knob: the escapement

Turn the knob slowly and let the ring settle after every step. The run assumes the ring settles with heavy damping, sliding straight downhill without overshooting, which is the textbook limit; a lightly damped ring can overshoot, and that is not modelled here. The ring holds rung 0 until its well is gone, at 0.7545 flux quanta, then slips to rung 1. It holds rung 1 until 1.7545 and slips to rung 2. Coming down it slips back at 1.2455, 0.2455, -0.7545 and -1.7545 , and going up again at -1.2455 and -0.2455 , and it ends where it started, on rung 0. Eight slips, every one exactly one rung.

It lands one rung up because the knob builds the next rung's well before it takes the old one away. When the old well vanishes the ring slides downhill into the nearest well, and the nearest is always the next rung. That is the escapement: the spring is wound between teeth, the flux through the ring creeping continuously within its rung, and the tooth releases exactly one step.

The ring does not follow the lowest rung. The lowest changes at each half-way point between whole numbers (-1.50 , -0.50 , $+0.51$ and $+1.51$ on the run's grid of 0.01), but the ring holds each rung a quarter of a flux quantum past that point, until its own well is gone. That lag is the latch: the ring remembers which way the knob came from. It also means the way down is not the way up, and the shaded loop between the two paths in panel B is energy spent on every round trip, about 0.36 of the energy one flux quantum stores in the ring's inductance per slip. Section W5 takes the loop up.

How hard the pen presses

The weak spot's strength, which the textbook calls β , is the most current it can pass times the ring's inductance, in flux quanta, times 2π . It is the one setting that decides what kind of object the ring is.

Below a strength of 1 there is no ladder. At 0.5 the ring has only one resting place at every setting of the knob, and its flux follows the knob smoothly: it crosses from rung 0 to rung 1 at exactly half a flux quantum, moving 0.0010 of a quantum in one knob step of 0.0005 , with no slip, no wall and no memory. The ring still reads a whole number, but nothing holds it; it is a follower, not a latch. Above 1 the slip appears: at 1.2 the flux jumps 0.2962 of a quantum in one step, at 3 it jumps 0.6386 , at 20 it jumps 0.8723 .

The stronger the weak spot, the longer each rung holds. Rung 0 gives way at 0.5445 flux quanta at a strength of 1.5 , 0.7545 at 3 , 1.0620 at 5 , 1.8500 at 10 and 3.4375 at 20 , each within a step of the model's own formula (panel C). With the knob returned to zero the ring keeps only the rungs inside that window: rung 0 alone at strengths of 1.5 and 3 , rungs -1 to $+1$ at 5 and 10 , rungs -3 to $+3$ at 20 . The wall between two level rungs grows the same way, from 0.0136 of the one-quantum store at a strength of 1.5 to 0.1193 at 3 , 0.2991 at 5 , 0.7861 at 10 and 1.7883 at 20 , and so does the energy each slip spends: 0.0415 , 0.3606 , 0.9092 and 2.4111 of the one-quantum store at 1.5 , 3 , 5 and 10 .

So the pen and the guard are one setting. A weak spot that writes easily keeps little, and a strong one keeps more rungs, walls them higher and spends more on every write. With no weak spot at all, a thin wire's wall is its whole cross-section, priced in the price table at about $13,000$ thermal jiggles for a niobium wire 20 nm across (an order of magnitude); a thick wire's is a vortex crossing it, not yet priced, and its current does not measurably decay. In the textbook picture, turning the knob then changes the current the ring carries but not its whole number.

A nick in the wire

An old wiring rule makes the same point from the other side; it is Bear's, from his own field practice, and the book gives it as experience, not as a citation. Communications were run in stranded wire, because a nick in a solid-core wire could compromise the circuit: the nick is where a solid conductor fails. In a superconducting ring the nick is where the count can change, and the engineer puts it there on purpose. The CERN lectures show the nicked ring directly: "In Fig. 34 the flux through a $\text{YBa}_2\text{Cu}_3\text{O}_7$ ring with a weak joint is shown. Flux jumps due to external field variations occur in multiples of $h/2e$ " (Schmüser, CERN-2004-008, the copy supplied on 23 September and searched for the passages quoted, Fig. 34 and the text beside it, from Gough et al., *Nature* 326, 855, 1987, as cited there, not read). And of the working instrument: "Technically one often uses superconducting rings with a single weak link which acts as a Josephson junction" (the same lectures, on the SQUID). The other half of the rule, strands and twisting as the guard against unwanted writes, is section W3's.

The same gearbox in other media

A ring of cold atoms. A ring-shaped gas of about 4×10^5 sodium atoms, 19.5 μm in mean radius, has been stirred up or down one whole unit of circulation at a time by a weak link rotated around it, with hysteresis: a loop of the same kind as panel B's, with the weak link's rotation as the knob (Eckel et al., *Nature* 506, 2014, arXiv 1402.2958, abstract and parameters read, as in the price table's cold-atom column).

DNA. Enzymes write DNA's linking number in whole steps. Gyrase cuts both strands, passes a segment through and reseals, and "both the introduction and the removal of supertwists by DNA gyrase change the linking number of DNA in steps of two" (Brown and Cozzarelli 1979, read for the price table's DNA column). Type I enzymes cut only one

strand: those of type IA change it in steps of exactly one, and those of type IB, which let the cut strand swivel, by one or more whole turns per cut (a reference summary citing Wang 2002; section W1). Section W1 made both writes on a ribbon; here they are a gearbox with two gear ratios. What each step costs is section W3's.

The first staircase ever seen. In 1961 two groups showed that the flux trapped in a superconducting ring comes in whole quanta. The CERN lectures describe the Doll and Näbauer experiment: without flux quantization the response "should be linear. Instead one observes a very pronounced stair-case structure which can be uniquely related to frozen-in fluxes of 0, 1 or 2 flux quanta" (Schmüser, CERN-2004-008; Doll and Näbauer, *Phys. Rev. Lett.* 7, 51, and Deaver and Fairbank, *Phys. Rev. Lett.* 7, 43, both 1961, as cited in the lectures, records located, not read). That staircase was written by freezing, cooling the ring with the field applied, not by a knob and a weak spot: it is section W4's pen.

What this section does not show

The ring here is the textbook model, not a device: no particular ring, junction or temperature is assumed, and every energy is given as a fraction of what one flux quantum stores in the ring's inductance, so no number here is a price in joules. The settling is heavily damped by assumption. Nothing thermal is in the run: at any real temperature the ring can hop a wall early, and a low wall lets it hop often; that is section W3's wall and section W4's freeze. The cold-atom ring was read only at the level of its abstract, so its slips are not claimed to match the model's in detail; the rf-SQUID hysteresis literature, Little and Parks, and the atom-ring experiments beyond Eckel et al. are still to check.

Sources: the textbook rf-SQUID model; Schmüser's lectures in CERN-2004-008 (the copy supplied on 23 September, searched for the passages quoted); Brown and Cozzarelli 1979; Eckel and co-workers 2014 (abstract and parameters); Doll and Näbauer 1961, Deaver and Fairbank 1961 and Gough and co-workers 1987, as the lectures cite them; Bear's field practice for the wiring rule. Identifiers: the reference list.

Runs: WRITE_W2_FIGURES 20260924_175837 (write_w2_figures.py, the textbook model, no measured constants, no random numbers); WALL_ESTIMATES 20260923_224526 for the thin solid wire's wall.

W3. What a write costs: the wall and the store

A write pays twice. It must climb the wall between whole numbers, and it must pay the difference in what the store holds on each rung. The wall is set by the body the count lives in, so a medium with no body writes for free, and a medium with a thick body may be impossible to write at all.

Three things come out of this section. The two costs are different quantities and can differ enormously: the price table's superconducting ring stores about one thermal jiggle on its first rung and holds it longer than anyone has watched, and even a wire 20 nm across walls its count at about 13,000 jiggles. The wall is a gap: a whole number can change only where something closes, which the Crossover Atlas shows in the simplest chain there is. And a wall guards only the path it stands on. Superconducting magnet wire is split into fine twisted strands to close a second path that thickness opens, which is the stranded wire of section W2's wiring rule.

Two costs

The **wall** is the highest point on the cheapest path from one rung to the next: what must be climbed for the whole number to change. The **store** is the energy held on each rung, and a write pays, or is paid, the difference between the two rungs. The prices are the price table's, of 23 September, and its three runs of record are named where each number is used. Both are measured here in thermal jiggles, the energy heat gives every part of a medium at its temperature (kT). The ruler matters because heat also climbs walls. In the textbook form the odds that heat carries a count over a wall on any one try fall off as e raised to minus the wall in jiggles, and the Field Guide names the wall for what it is: "A write barrier is a Morse height ... That is the ΔE in an Arrhenius retention law" (Field Guide, draft v4.6 of 9 August, §8, read in full for this part). A wall of a few jiggles is climbed constantly; a wall of thousands is never climbed by heat at all.




Figure W3.1. (A) The store (circle) and the wall (triangle) for a write in each medium the price table has priced, in thermal jiggles at the medium's own temperature. The bare filament has neither. DNA's store per step is priced; its wall, a covalent bond, is not. The superconducting ring stores about one jiggle per rung; no wall is drawn for its 0.5 mm wire, because the barrier formula prices only wires thinner than about one coherence length and this ring's wall, a vortex crossing, is not priced, while its current does not measurably decay (the dashed arrow); the wire of 10 nm radius is inside the formula and its wall is drawn. A weak spot of the strength used in W2, put in the same ring, lowers the wall to a tenth of a jiggle. (B) The SSH chain of the Crossover Atlas, Chapter 6, recomputed: the bulk energy gap against the ratio of its two hopping strengths. Dots are the paper's seven settings, the line is the exact formula, and the shading marks winding 1. Panel A is read line by line from the four ledgers named in the text, every plotted number printed back, and its one derived number, the nicked ring's wall, is the product of two of them; panel B is pure mathematics, the chain recomputed by the receipt's own method and by the exact formula. Drawn by write_w3_figures.py, version 2.

No body: nothing to pay. The bare filament of READ section 7 changes its count as one ring slides through the other's wire, and its energy drifts smoothly by 3.80 per cent along the whole path. Its departures from a smooth curve at the two crossings, 0.00314 and 0.00243, are smaller than the largest anywhere else, 0.00611. No wall, and nothing held on either rung: a pure geometry writes for free, which is why it is a label and not a latch.

DNA: a priced store behind an unpriced wall. A 5,000 base-pair plasmid at the typical bacterial density of extra turns needs 12.98 jiggles at room temperature for its next unit of linking and 26.40 for gyrase's next step of two (its constants are unverified or assumed). The wall is a covalent bond in the backbone, complete against heat and open to an enzyme; its height stays unpriced until the backbone's uncatalysed breaking rate is read (price table, DNA column). DNA also has the wiring rule in it. "A single nick lets the strands swivel, and the stored twist is gone at once" (the price table's LEAK row): a nick in one strand removes the wall entirely, and the type I enzymes of W1 make that nick on purpose and let the strand swivel under control.

A superconducting ring: one jiggle, held by a thick body. In the price table's example ring, a 1 cm loop of 0.5 mm wire at liquid-helium temperature, the first flux quantum stores 0.954 of a jiggle (geometry assumed). In a wire thinner than about one coherence length, changing the count means the superconducting order must vanish across a whole cross-section, and for a wire of 10 nm radius that wall is some 1.3×10^4 jiggles (an order of magnitude). The formula is Langer and Ambegaokar's (1967), read as printed and cited in Joshi and Bhattacharyay (2011), who give it for samples whose "width is smaller than the coherence length ξ ". The example ring's wire is 0.5 mm in radius against niobium's 39 nm coherence length, so the same formula's 3.1×10^{13} jiggles is only what a slip through its whole cross-section would cost. A thick wire changes its count by a vortex crossing it instead, a lower door that is not yet priced. The ring still holds: a thick superconducting ring's current does not measurably decay (the price table's LEAK row). This ring holds its number because its body is thick, not because the number is worth much.

The same ring with a weak spot. Section W2 measured the wall between two level rungs of a ring with one weak spot as a fraction of what one flux quantum stores: 0.1193 at a strength of 3. Put that weak spot into the example ring and the wall becomes $0.1193 \times 0.954 = 0.114$ of a jiggle, a derived number, not a device. That is below heat's own size, so heat would write this ring as freely as the knob does. The nick overshoots: it drops the ring's wall below heat. A weak spot is a usable pen only where its own wall still stands many jiggles high, which means a ring whose one-quantum store is many jiggles, a small ring of small inductance. The arithmetic is short: at liquid-helium temperature one flux quantum stores one jiggle in a ring of 36.87 nH, ten in a ring of 3.687 nH and a hundred in a ring of 0.3687 nH, and at liquid-nitrogen temperature one jiggle needs a ring of 2.011 nH (exact SI constants). That this is why working rf SQUIDS are built small is textbook and still to check against a device reference.

The gap is the wall

The cleanest picture of a wall is a chain, not a ring. The Crossover Atlas, Chapter 6, and its receipt `ssh_edge.py` walk the SSH chain: a row of sites joined by bonds that alternate between two strengths, v inside each pair and w between pairs. Its whole number is a winding, how many times a walk around the chain's own loop of wave numbers goes around zero: 0 when w is weaker than v , 1 when it is stronger. The paper tunes w across v and reads the energy gap of the chain's interior: "at $w = 0.3$ the gap is 1.40; at 0.7 it is 0.60; at 0.95 it has thinned to 0.10; at $w = v$ it collapses to 0.0031, numerically the closed abyss; and on the far side it reopens, 0.10, then 1.00, then 4.00. The winding holds at 0 all the way to the abyss and is 1 immediately after it." And: "The gap resists, then breaks over, then the integer has flipped — a triac, made literal."

This book's run recomputes the chain in its own code rather than copying the table. By the paper's own method the gaps come out 1.4000, 0.6000, 0.1000, 0.0031, 0.1001, 1.0000 and 4.0000. The exact formula, twice the difference of the two strengths, gives the same except at the meeting, where it gives exactly zero: the 0.0031 is the width of the method's grid, and at $w = v$ the walk passes within 10^{-16} of zero, so the winding has no value there at all. Scanning w in 800 steps, the winding changes only between 0.9991 and 1.0031, where the gap is below 0.008.

The same chapter states the rule every write in this part obeys: "a *topological* transition has a protected integer that cannot change without a gap or order parameter closing; a mere threshold has no such integer." The superconducting ring is the order-parameter case. In a thin wire its wall is the price of closing the superconducting order across one cross-section; in a thick one the order closes only along the core of a vortex as it crosses; and a weak spot is a place where the order is already thin. In the chain the knob closes the gap everywhere at once; in the ring the slip closes the order in one place for a moment. Either way something must close, and the cost of closing it is the wall.

A wall guards one path

A wall is only a wall against the path it stands on. The Field Guide says it from its own bench: "The escape path matters more than the barrier height ... An invariant is protected against smooth interior deformation and against nothing else. Any claim that a topological memory is stable must name the escape path it is stable against" (draft v4.6, §8, which gives magnetic skyrmions as the literature's example: their cheapest way out runs through the sample's edge, around the protection).

Superconducting magnet wire is the engineer's version. Thickness raises the phase-slip wall, but it opens other paths. One is a vortex crossing the wire, the door a thick ring's count takes. Another opens in a changing field. In a thick piece of superconductor the screening currents set up by a changing field "can become unstable and collapse. When this happens, magnetic flux rushes into the slab in a process known as a flux jump" (Wilson, "Practical superconductors for accelerators: filamentary conductors", in CERN-2004-008). So the conductor is split: "To avoid flux jumping, it is necessary to subdivide below a characteristic diameter, which for NbTi is $\sim 50 \mu\text{m}$ – and this is always done", and the accelerators' cable is "a flat twisted cable comprising ~ 30 wires, each of which contains some 5000 – 20,000 fine filaments of superconductor embedded in a matrix of copper." Splitting is not enough by itself. In a changing field, currents cross the copper between filaments and the bundle screens like one thick conductor again, decaying over "days or even years" in a long wire. The cure is the twist: "twist the wire so that the coupling currents are forced to reverse every twist pitch," which cuts their decay to "a few msec" (the same lecture).

That is the wiring rule's other half, Bear's rule from his own field practice, as in section W2. Stranded wire spreads the current so that one nick cannot open the circuit; a twisted pair reverses the loop the two wires enclose at every crossing, twice a twist, so a changing field induces opposite voltages in neighbouring loops and they cancel. The book checked it on ideal wires. Closed into one loop, a pair of whole twists links no flux at all from a uniform field across it, and none from a field that grows steadily along it. Whatever its length, it never links more than the area of one half-twist, 0.0318 for wires 0.1 apart twisting once per unit length. An untwisted pair of the same wires 20 twists long links 2.000, the whole strip between them. Magnet wire does both, for the same reason: to shut the paths a count could escape by that the builder did not choose.

A warning on words. The CERN lectures use "flux jump" for two different events. In section W2's nicked ring (Schmüser, Fig. 34) the flux jumps are the ring's rung slips, whole quanta at a time, the pen working as designed. In Wilson's magnet wire a flux jump is a thick conductor's screening currents collapsing, the fault the stranding exists to prevent. Same words, the pen in one lecture and the fault in the other.

What this section does not show

The walls in Figure W3.1's panel A are orders of magnitude. The superconducting wall's formula is now read in print (Joshi and Bhattacharyay 2011; the 1967 original is not read) and holds only for wires thinner than about one coherence length, so panel A draws no wall for the 0.5 mm wire; the thick wire's real wall, a vortex crossing, is held (Langer and Fisher 1967; Kogan's strip barrier). DNA's covalent wall is unpriced, and the cold-atom ring's wall (to be read from Eckel et al.) and water's cut (the reconnection timescale, from Kleckner and Irvine) are held. The nicked-ring row multiplies a textbook-model wall by the price table's assumed example geometry; it shows a scale, not a device. The odds that heat climbs a wall are given in their textbook form and not run here; temperature as the write-enable is section W4's. The twisted pair's cancellation is checked for ideal thin wires in a smooth field, not for real cable, and whether it matches what a field installation shows is a question for someone who has run the wire. The rf-SQUID arithmetic is exact; the reason working devices are small is still to check.

Sources: the price table of 23 September; Langer and Ambegaokar 1967, as Joshi and Bhattacharyay 2011 print and cite it (the 2011 paper read in full, the 1967 original not read); the Crossover Atlas, deposited version 3, Chapter 6 and its receipt `ssh_edge.py` (the markdown read in full); the Field Guide, draft v4.6, §8 (read in full); Wilson's and Schmüser's lectures in CERN-2004-008 (the copy supplied on 23 September, searched for the passages quoted); Bear's field practice for the wiring rule. Identifiers: the reference list.

Runs: PRICE_TABLE_ESTIMATES 20260923_212608, WALL_ESTIMATES 20260923_224526 and DNA_ESTIMATES 20260923_222155, the price table's runs of record; WRITE_W2_FIGURES 20260924_175837; WRITE_W3_FIGURES 20260924_214245 (`write_w3_figures.py` version 2, reproducing every number of the first run, 20260924_181503, and no longer drawing the 0.5 mm wire's whole-cross-section slip as a wall); WRITE_CHECKS 20260925_234135 (`write_checks.py` version 2, pure mathematics and exact SI constants, ledger LATCH_ATLAS_WRITE_CHECK_LEDGER_2026-09-25) for the twisted pair and the inductance arithmetic.

W4. Freeze and thaw: temperature as the write-enable

Heat a ring until its order melts and its count is free; cool it, and the count locks. Cool it in a field and it freezes, by odds, the whole number nearest the field, and keeps it after the field is gone. Cool it fast with no field and it freezes a whole number nobody chose.

Four things come out of this section, and every one was checked against the experiments that wrote rings this way, nine papers read in full in their arXiv versions, with Weiler and co-workers' *Nature* paper as a tenth and Kirtley and Tafuri's viewpoint on the niobium ring as an eleventh, both supplied by Bear. The reading behind every claim is recorded in LATCH ATLAS W4 DEEP PASS NOTES 2026-09-24. First, temperature is the write-enable, and it works at two temperatures, not one: a higher one where the rungs of the count come into being, and a lower one where they lock, with the count written in the gap between. Second, every freeze writes the same way: a spread of possible windings shaped like a bell curve, sorted into whole rungs. Heat is one route to that bell curve and a ring freezing in separate pieces is another, and the two give the same staircase so closely that no field scan can tell them apart. Third, in real rings the two routes write together; in a niobium ring they lock at almost the same moment, and the measured way the count falls with cooling speed is steeper than either route alone gives, which is the open question this section hands on. Fourth, a freeze writes only into a closed loop and keeps its number only around a hole, and a trap's shape alone can seed a winding as the order grows.

Thaw: the wall melts, at two temperatures

Section W3 priced a superconducting ring's wall as the cost of destroying the superconducting order across one cross-section of the wire. That order has a temperature. The CERN lectures state it plainly: "the energy gap in a superconductor is not a constant but depends on temperature. For $T \rightarrow T_c$ one gets $\Delta(T) \rightarrow 0$ ", and "The temperature dependence of the critical field is caused by the temperature dependence of the gap energy" (Schmüser, CERN-2004-008, the copy supplied on 23 September). The wall in W3 was built from that critical field (which uses niobium's 0.199 T), so as the ring warms the wall sinks with the gap and is gone at the transition itself. With no wall the count is not held; warm the ring through its transition and whatever it stored is released.

This is the Crossover Atlas's rule met from the other side: an integer "cannot change without a gap or order parameter closing" (the deposited version 3, Chapter 6). In W3 the knob closed the gap. Here heat closes it. It is the logic of a solder joint, heated to set and cooled to hold, with a whole number in place of a position. And it erases without walking back down the ladder: it removes the ladder, one of the ways section W5's erase differs from a write run backwards.

Two temperatures. Cooling does not lock the count the moment the order returns. Above the transition there are no rungs at all: the order that makes them is gone, and there is nothing to count. Just below it the rungs exist, but the walls between them are still lower than the thermal jiggle, so heat keeps moving the count from rung to rung. The count locks only lower down, at the freezing temperature, where the walls outgrow the heat. Kirtley, Tsuei and Tafuri write the rule for that moment (their Eq. 4). In words: the wall at the moment of freezing, measured in thermal jiggles, equals the logarithm of the number of tries the ring gets at crossing it while the wall climbs by one more jiggle. It is the Field Guide's lifetime law, $\tau = \tau_0 \exp(\Delta E/kT)$ (row 51; a stored bit lasts exponentially longer the higher its wall stands in jiggles), added up over a cooling ramp: the count freezes once fewer than one more crossing is expected before the ring is cold. Run on their own printed inputs, the rule puts the freeze of their Mo_3Si rings 1.57×10^{-3} of the transition temperature below it (printed 2×10^{-3}), with one flux quantum then worth 4.49 thermal jiggles (printed 4.3); in a niobium ring the gap is only 2.73×10^{-5} to 4.02×10^{-5} of the transition temperature (at 10^{10} attempts per second, an assumed rate). The narrower the gap, the less room heat has between the moment the rungs appear and the moment they lock.

Magnets have the same pair. A magnet's upper temperature is its Curie point, above which it is not magnetic at all. A grain small enough to be a single domain also has a lower one, a blocking temperature, above which heat flips its direction freely although it is still magnetic inside. The flipping "occurs below the Curie temperature of the material", in grains that are "single-domain", and the time between flips grows as e raised to the grain's wall in jiggles, $\tau = \tau_0 \exp(KV/kT)$, where K is the anisotropy energy per unit volume and V the grain's volume. The blocking temperature is where that time equals the time the grain is watched (Wikipedia, "Superparamagnetism", read 25 September, citing Néel 1949, not read). So the lower temperature is not the material's alone: it is set by how long the magnet is asked to hold, which is section W3's wall and odds in a magnet's words. The fence: ordinary resistance rising with heat is not this. In a plain metal heat adds scattering smoothly; there are no rungs, so none break.

Frozen with a field: one staircase, two routes

The first staircase of whole numbers ever seen was written this way. Doll and Näbauer's lead tube was warmed, placed in a "magnetising field" along its axis, cooled, and the field switched off; the flux left behind was then read. "Without flux quantisation the relation between resonant amplitude and magnetising field should be linear. Instead one observes a very pronounced stair-case structure which can be uniquely related to frozen-in fluxes of 0, 1 or 2 flux quanta" (Schmüser, CERN-2004-008, describing Doll and Näbauer 1961, which the book has not read). The price table's WRITE cell says the same in one line: cooled with flux applied, the ring "traps the whole number nearest the applied flux, keeping it when the field is removed."

A ring frozen two ways (stated models; random seed 20260924)




Figure W4.1. A ring frozen two ways, in stated models. (A) Frozen with a field applied: the mean whole number frozen in against the applied flux, for three values of what one flux quantum stores at the moment of freezing, in thermal jiggles (solid, the heat route); the curve for 1 jiggle lies under the one for 0.1. The dashed curves are the domain route, a bell curve centred on the applied flux and rounded into rungs, at the width that fits each solid curve best; the pale grey steps are the nearest whole number. (B) Frozen fast with no field, in pieces: the share of 20,000 freezes landing on each whole number, for 12 and for 192 independent domains around the ring. (C) The spread of that whole number against the number of domains: dots measured, line the square root of $N/12$. Stated models, no measured constants, the thermal weighting at freezing assumed; 20,000 freezes per ring size and 200,000 per comparison with Corman's rule, numpy's default generator with seed 20260924, so the run repeats exactly. Drawn by write_w4_figures.py, version 2.

"Nearest" hides a condition. At the moment the ring freezes, heat is still at work, and a rung that sits a little away from the applied flux costs a little more to hold. The ring takes each rung with odds that fall off as e raised to minus that extra cost in jiggles, the textbook thermal weighting; this is Kirtley, Tsuei and Tafuri's Eq. 5 term for term, and it is the heat route drawn solid in panel A. Where one flux quantum is worth 10 jiggles, a field a quarter of a quantum away from a whole number still writes that whole number 9,933 times in 10,000, and the mean frozen number stays within 0.0061 of the nearest whole number whenever the field is within a quarter of one: a sharp staircase. Where one quantum is worth 1 jiggle, even a field sitting exactly on a whole number writes it only 5,641 times in 10,000, and the mean simply follows the field. At 0.1 of a jiggle the nearest whole number wins only 1,784 times in 10,000. The laboratory has both ends of this panel. Kirtley and co-workers report that small aluminium rings (Davidović and co-workers) show the sharp staircase, while their own large MoSi rings give odds that barely follow the field, fitted at 1.75 jiggles per quantum, between the orange curve and the blue.

The domain route gives the same staircase. Weir, Monaco and Rivers write the ring frozen in pieces in a field as a smooth bell curve of windings centred on the applied flux and rounded to the nearest rung (their Eq. 6). Drawn dashed in panel A at its best-fitting width, it sits on heat's curve to within 0.00956 at 10 jiggles and to within 1.34×10^{-6} at 1. Across the range the experiments scanned, every heat store from half a jiggle to 30 has a width that brings the two within 0.0121 everywhere, and within 0.0043 at Kirtley's 1.75; the width Weir and co-workers fitted, 0.39, is a heat store of 3.11 jiggles in disguise. Those gaps are smaller than the counting error of a single point from 1,000 shots, 0.0158. The test proposed for whether a freeze is a bell curve at all, one fluxoid in fewer than about half the shots, has the same cap for both: 0.4839 for heat, 0.4843 for the rounded bell curve.

So the two routes are one operation seen from two ends, as Bear put it: heat counts a bell curve rung by rung, the pieces round a bell curve into rungs, and once the bell curve is broad, sampling it at the whole numbers and binning it into whole-number bins give the same centre and the same shape, with one exact difference: binning widens the bell

by a twelfth of a step squared in variance. For a bell one step wide centred at 0.3, the sampled odds have variance 1.00000 and the binned odds 1.08333, and the binned odds are the sampled odds of a bell of variance $1 + 1/12$ to within 0.0003 in every odd; the two routes' own odds differ by at most 0.014. Below half a step wide the routes part, the sampled route pulled harder to the nearest whole number. That check is pure mathematics. The twelfth is Sheppard's correction for grouped data (MathWorld, "Sheppard's Correction", read 25 September), and the sampled route's exactness for a broad bell is the Poisson summation behind the theta function. A field scan, however clean, yields one number, the width, and the twelfth means the two routes read the same width differently. Which route wrote the ring shows only in how that width moves with cooling speed and with the ring's shape.

The price table's 1 cm example ring stores 0.954 of a jiggle per quantum even at liquid-helium temperature (geometry assumed), and freezes warmer than that, so frozen in a field it would write only odds. A field of exactly half a quantum makes the two nearest rungs equally likely at any store: set the knob half-way, freeze, and the ring tosses a coin.

Frozen fast: a number nobody chose

Cool a ring quickly with no field and it does not order all at once. Different stretches order independently, each picking its own phase, and where two stretches meet, the phase joins the short way round. Walking once around the ring adds up the joins, and the total is a whole number of turns, which the ring then keeps. Nobody chose it.

The run cuts the ring into N domains, gives each a phase drawn evenly from the whole circle, joins neighbours the short way, and freezes 20,000 rings for each N . Every freeze produces an exact whole number, and the whole numbers centre on zero (a mean of +0.0069 at 384 domains), so no rung is favoured. Their spread grows as the square root of $N/12$: 0.9890 against 1 for 12 domains, 3.9910 against 4 for 192, 5.6333 against 5.6569 for 384, within 1.5 per cent at every N tried. With 12 domains, 61 per cent of freezes hold a nonzero count; with 384, 93 per cent do. READ section 5's half-turn limit holds here too: each join is less than half a turn, so a ring of N domains can never freeze $N/2$ turns or more, and the ring of 3 never froze more than 1.

The same formula is printed in three places in the literature read for this section. Aidelsburger and co-workers print it and measure it in a ring of gas cut into pieces by hand, where they set N directly: three pieces wind with odds of $1/8$ each way, a quarter in all, and the run's three-piece ring came out nonzero in 0.2465 of freezes. Weir, Monaco, Koshelets, Mygind and Rivers print the same width in their own notation. And Corman and co-workers' rule, which adds up random phases and takes the nearest whole number, gives the same average size of winding as the short-way joins: 0.2489 and 0.2494 at 3 domains, 0.6549 and 0.6554 at 10, from 200,000 freezes each. Their measured range, 0.2 to 0.6, is, in their words, "obtained for $3 \leq N \leq 10$ ": their ring froze into only a handful of pieces.

The loop closes because the pieces are binned. The three-piece case shows where the whole number comes from. The raw phase differences around a closed loop always add to exactly zero; without binning the loop would close with no winding at all. The winding appears only because each join takes its own step the short way, a bin between minus half a turn and plus half a turn, and the binned steps then add to a whole number of turns: none, or one either way, at odds of $1/8$ each way, counted exactly on a fine grid as 0.1249 and 0.1251. Each little join reports locally, and the whole number is what closing the loop makes of the local reports. Bear's reading, and the arithmetic bears it out: the frozen integer is the holonomy of the local bins.

How the number of pieces depends on cooling speed is the Kibble-Zurek law: the size of a domain grows as a power of the cooling time, in superconductors as its fourth root. That root is why every slope in this section is small and hard to measure. Over the 25-fold range of cooling times of the niobium experiment the domain size changes only 2.2361-fold. And the pieces do not stay fixed once formed: Aidelsburger's two-step joins give the first measured time for pieces to forget each other, 52 to 90 ms, so a ring's pieces keep merging until it closes.

Pieces that pull on each other while they settle narrow the draw. The short-way rule gives the spread of the frozen number when the pieces are joined all at once, as Aidelsburger's walled pieces were. If neighbouring pieces keep pulling on each other while the ring settles, some joins slip through half a turn and the winding changes. Groisman, De Vita, Fernández Bonder and Zhang (2025) work this out for a ring of phases each pulled by its two neighbours, the simplest such ring. The winding still comes out Gaussian, but its variance, $(n - 1)/12$ at the start, shows "the variance gradually decreasing during the early stage of the evolution but quickly converges to its final shape", and it settles quickly, at a time that grows only as the logarithm of the ring's size. They also prove where the losses come from: a winding can change only when a join passes through half a turn, and "once a phase difference enters $(-\pi/2, \pi/2)$, it will never leave". A join that starts within a quarter-turn can never slip, so every lost turn comes from a join that started strained toward half a turn. The size of the narrowing depends on the ring. In the book's run of that same ring, the settled spread was 0.585 of the starting one, and 0.487 when each phase is also pulled by its

next-nearest neighbours; joins that started above 0.8 of a half-turn were 20.03% of all joins but carried 89.30% of the slips. So when a measured spread is read back to a number of pieces, as with Corman’s 3 to 10, pieces that settled together would mean more pieces than the short-way rule says. The paper gives no number for how many more.

Three clouds merged. Scherer’s dissertation (University of Arizona, 2007, read in full), the fuller account of the Scherer, Weiler, Neely and Anderson letter (*Phys. Rev. Lett.* 98, 110402, 2007, record checked), merged three condensates held apart by a laser barrier and counted the vortices left behind. The fraction “displays a maximum of ~ 0.6 for the smaller τ values, and decreases to ~ 0.25 for long ramp-down times”, which the author reads as the geodesic rule at the slow end and extra vortices from interference fringes at the fast end. The points rest on “5, 11, 10, 10, 5, and 5 images”. Five images carry a counting error of 0.194 at a true fraction of 0.25, and a true 0.25 would give exactly one vortex image in five with odds 0.396, a true 0.6 with odds 0.077, so each slow point leans toward the quarter without settling it. Aidelsburger’s rings, with more than 200 shots per point, carry the precise test of the quarter.

Heat or pieces: which wrote the ring?

Seven experiments are read for this section. The table sets out what each froze, how it was read and what it measured; the paragraphs after it take the question the table leaves, which route did the writing.

experiment	what froze	how it was read	what was measured	which route wrote it
Kirtley, Tsuei and Tafuri (2003)	thin-film Mo_3Si rings, cooled through 7.81 K at set rates, in and out of a field	in place, by a scanning SQUID microscope	odds that barely follow the field, fitted at 1.75 jiggles per quantum	heat: vortices carried across the ring’s wall
Monaco, Mygind and Rivers (2002)	an annular Josephson junction, niobium	by motion: driven round the loop, N fluxons give a voltage $V = N\Phi_0 v/C$	slope 0.27 ± 0.05 against the predicted 1/4; size within 1.18 of prediction	pieces, on thin statistics; which freeze writes a junction is open
Monaco, Mygind, Rivers and Koshelets (2009)	one niobium ring, radii 30 and 50 μm , 200 nm thick	in place, by a junction along the inner edge	one fluxoid in 0.18 of shots at 1 ms, falling as the cooling time to the power -0.62 ± 0.15	both, locking together (model below); the slope is open
Corman and co-workers (2014)	a flat ring of rubidium gas, radii 9 and 15 μm	by release: a spiral interference pattern with a central disk	mean size of the winding 0.6 at 0.025 s, 0.2 at 2 s; slope 0.19 ± 0.06	either; heat alone fits all but the fastest point (model below)
Aidelsburger and co-workers (2017)	a ring of gas cut into N pieces by hand, then joined	by release: a spiral with a second ring as reference	three-piece odds of a quarter; the spread as the square root of $N/12$	pieces, set by hand; heat 1.5 (8) per cent in the one-piece ring
Scherer’s dissertation (2007)	three condensates, merged	by release, after the cloud expands about five times	about 0.6 fast, about 0.25 slow, on 5 to 11 images a point	the geodesic rule, with extra vortices from fringes when fast
Weiler and co-workers (2008)	a condensate grown by evaporative cooling	by release	spontaneous vortices, more with faster cooling; a ring keeps its winding	regions and heat together, in their simulation

Kirtley’s route: heat carries the number across the wall. Kirtley, Tsuei and Tafuri account for their rings by a “freezout of thermally activated fluxoids, mediated by the transport of bulk vortices across the ring walls”, which they call “complementary to a mechanism proposed by Kibble and Zurek”. Their escape path is not the one section W3 priced. W3 priced a phase slip, the cost of destroying the order across a whole cross-section of wire. In a wide thin film the count changes by a cheaper door: a vortex made inside the strip and carried across it, with a wall equal to one quantum’s store multiplied by the logarithm of the strip’s width measured in vortex cores (their Eq. 7, from Kogan). A claim of stability has to name the escape path it is stable against, and a wide film has a lower door than a thick wire.

Corman’s ring: heat alone could have done it. Carried into Corman’s flat ring of gas, heat’s route gives the moment of freezing and the odds at that moment from the ring’s own sizes. At the instant the ring turns superfluid, a quantum of winding is worth 1.022 jiggles, so heat picks the number. Heat alone then lands on Corman’s slow end (0.1938 against their 0.2) and on their slope (0.190 against 0.19 ± 0.06), and falls short only at the fastest cooling, where it tops out at 0.4492 against their 0.6; 36 shots from heat’s ceiling would average 0.6 or more with odds 0.0502. So Corman’s data do not separate the routes, except weakly at the fastest cooling, which is also where the pieces should be strongest.




Figure W4.2. Heat or pieces, in a stated model of Corman's ring. (A) Heat alone: the mean size of the frozen winding against how superfluid the ring is when it freezes; the band is Corman's measured range, the dashed line the moment the ring turns superfluid, the dotted line Corman's own check at the end of cooling. (B) Heat frozen at the moment of freezing, against cooling time, for three vortex-core sizes, with Corman's two measured ends. (C) The fork at the fastest cooling: how the windings change when the ring is made twice as large, or fatter at the same middle radius, by the heat route and by the domain route. Every number is printed in that run's ledger.

The fork that does separate them. The two routes read different things off the ring. The pieces are counted around the circumference, so they answer to the ring's length. Heat prices a quantum by the ring's fatness, so it answers to the ring's shape. Fatten the ring at the same middle radius and cool it fast: heat's windings fall to 0.761 of Corman's at an outer-to-inner ratio of 2 and to 0.406 at 3, while the pieces' stay where they are, for as long as the strip stays narrower than a piece (Figure W4.2, panel C, from the same run). That is a test a cold-atom group can run with rings of one middle radius and different fatness.

Monaco's niobium ring: priced cold, and priced at the freeze. Monaco and co-workers checked heat by weighing one quantum's energy against the thermal jiggle, which is this section's store. Recomputed from their printed inputs, their ring's cold store is 227.4 jiggles per quantum (the printed energy is 4.90×10^{11} times too large, and the correction does not change their conclusion). But that store is named at no temperature. Near the transition the film's penetration depth, and with it the ring's inductance, grows without limit, and at the moment of freezing a quantum is worth only 2.380 jiggles at 1 ms. This is the Field Guide's row 172 worked in a real ring: "a constant stated without naming its temperature has no address".

At the freeze, heat alone gives a share of shots at one fluxoid close to the measured size, 0.1562 at 1 ms against their 0.18. That agreement is not a finding. It rests on the one input no niobium paper gives, how often the ring tries to cross its wall, and across a hundredfold range of that rate the share moves from 0.0381 to 0.3311. Bear's paper *A Coincidence Has No Single Price* (Schulz 2026, cited for its method only) gives the rule: a claim can only be tested where it stops depending on the quantity nobody has measured well. Here that is the slope, which barely moves with the rate. Heat alone falls as the cooling time to the power -0.309 at 10^{10} tries a second, and between -0.250 and -0.452 across the measured ranges of penetration depth and coherence length. The measured -0.62 ± 0.15 sits 2.08 standard errors from heat's -0.309 and 2.47 from the pieces' $-1/4$. Each route alone is too shallow.

In niobium the two routes lock together. In Kirtley's films heat locked about ten thousand times further below the transition than the pieces did, 10^{-3} of the transition temperature against 10^{-7} (his own estimate), so heat had the last word. In Monaco's niobium ring heat locks only 2.13 to 15.64 times further below it than the pieces (at 10^{10} tries a second). At that point the ring's wall is only 3.75 to 6.85 jiggles high, so crossings are not rare events: Kirtley's picture of a wall climbed now and then turns into the steady wandering of phase that the pieces describe. Weiler and co-workers reached the same place from a condensate: the simulation that matches their counts carries the forming regions and the heat of the surrounding cloud together, and heat "can have an important role".

The slope, and a smothered bend. What steepens Monaco's slope to 0.62 is still open, and there are two stories. Kirtley and Tafuri's viewpoint on the ring (2009) granted the suppression without a number, on the ground that the rings are narrow enough, which is not Monaco's own reason; neither paper priced heat's share. Weir and Rivers judge a ring's regime by what it measured: the simple power law breaks when fewer than about 0.2 of the shots hold a single fluxoid, and Monaco's fit runs from 0.1800 at 1 ms down to 0.0245 at 25 ms, at or below that line throughout. So the ring sits in the bend toward exponential suppression, the authors' own "shoulder". A square root in that bend

keeps it looking straight: the suppression forced through the two ends of Monaco's fit reads as a slope of 0.611 on five points, but with a straightness ($R^2 = 0.9486$) poorer than Monaco's measured 0.987, so his points are part power law and part bend. Heat may also be writing alongside. Where to point a specialist: a noisy Ginzburg-Landau simulation mapped to Monaco's niobium ring should say whether heat and pieces together steepen the slope to 0.62. That is beyond this book's tools, and the book says so.

The junction ring: clean numbers, an open writer. The oldest experiment is the cleanest for the pieces. Its printed numbers reproduce from its printed inputs (domain size 5.92 mm at 1 s, prefactor 0.0845 against a fitted 0.1), and its slope sits 0.40 standard errors from 1/4. But its count is a twist in the phase difference between two films, and by topology the difference's twists equal the top film's twists minus the bottom film's. The book checked this on 980 wiggled pairs of loops with windings from -3 to 3 each, and the difference's winding was $n_1 - n_2$ in every one. It is the additivity of winding numbers, and it holds whenever each loop is read finely enough that no step reaches half a turn. Each film freezes first, at 9.1 K, and the junction later, at 8.95 K, so some of what was counted may have been written by the films' freeze. Where to point a specialist: in a closed annular junction, is the net count set at the films' transition or at the junction's?

Keeping what the freeze writes: the loop, the hole and the shape

Cut the loop and nothing is written. Aidsburger and co-workers put a wall across their reference ring so that cooling could not write into it. With the loop cut, there is no path to wind round, and no count ("no loop, no count", READ section 4). Cutting the loop is a write-protect.

Close it around a hole and the number is kept. In Weiler and co-workers' plain trap the spontaneous vortices can drift out and the freeze's writing fades. With a barrier through the middle, making a ring, a share of the runs keeps its winding "despite the presence of dissipation". A ring keeps what a freeze writes; a disk does not. Scherer's simulations show the step in between: "a single vortex core may migrate to and be pinned at the center of the barrier where the energy cost of displacing fluid is low", the hole taking the number in.

The trap's shape can seed the winding. With a weak barrier, too low to keep three clouds apart as they grow, Scherer's condensates merged by themselves while forming, and vortices appeared in 0.56 ± 0.06 of 16 images; "vortices are formed during the BEC creation process rather than during the ramp down". The dissertation concludes that "the shape of a static confining potential may be sufficient to induce vortex formation during condensate growth", which "challenges the notion that a condensate necessarily forms with no orbital angular momentum", and it points the idea at superconducting films whose flux is held by a regular array of built-in defects (its references 100 and 101). The barrier is also the handle: it "breaks the trap's cylindrical symmetry, and thus allows for the exchange of orbital angular momentum between the atoms and the trap." A write needs something to push against. Sorted by where the angular momentum comes from, there are three kinds: brought in by light, as in the laser writes of the price table's WRITE cell, and chosen; borrowed from the container, as here, and steered by the pieces' phases; or none, the plain freeze, a draw that averages to zero.

A ring thawed and refrozen on purpose: the ratchet. Gurtovoi, Exarchos, Antonov, Nikulov and Tulin (2015) strung 667 aluminium rings, each about a micrometre across, into one chain; one arm of each ring is 100 nm wide and the other 125 nm. Knocked back and forth across the transition by an alternating current or by electrical noise, each ring refreezes over and over, and each refreeze draws a whole number of turns, weighted by the field toward the state of least energy, "the predominant probability $P_n \propto \exp -E_n / k_B T$ ". The unequal arms turn the ring's circulating current into a voltage, so the draws add up to a steady voltage that swings with the field, one period per flux quantum, and is zero at whole and at half quanta. In the authors' words, "The asymmetric superconducting ring is a ratchet because of the predominate probability of one of the directions of the persistent current." A symmetric ring gives nothing, and at zero field neither does the asymmetric one. So the jobs split three ways: the field writes, the shape reads, and the kick rerolls, which it can do only when it is strong enough to melt the ring: "The non-equilibrium noise can not switch the rings in the normal state when the minimal value of the critical current (4) exceeds its amplitude". Here the shape seeds no sign. It is the needle, and the field is the handle.

A chain reads in step only when its rings are one size. The rings' voltages add "as well in series of any dc power sources", so one common field makes the chain's signal grow in proportion to the number of rings, while the coin-toss noise at half a quantum grows only as its square root (slopes 1.0000 and 0.5051 in the book's stated model of the chain). The authors ask for rings of one area, "because $\Phi = BS$ ". In the same model, 667 rings whose areas differ by 1% drift out of step as the field grows. The swing fades with the square of the field times the square of the spread, while its period holds at the mean ring's, until the chain is no louder than independent rings would be. So the fade

measures the spread: the swing halves after about 0.187 divided by the spread, in flux quanta (18.5 by that rule and 19.5 in the run, for a spread of 1.011%). The shape of the spread acts as a window does. A flat spread of sizes, or two batches, brings the swing back with its sign reversed, and a reader there would read every turn backwards. This is the mirror of pieces that settle together: a common field lines the draws up, and a spread of sizes lets them drift apart again.

Heat runs over the walls. In Scherer's trap the thermal cloud sits near 22 nK, the strong barrier at $k_B \times 26$ nK and the weak one at $k_B \times 7$ nK, and the author assumes "the energy of the central barrier is low enough that it has negligible effect on the thermal atom cloud". The walls sort only the ordered part, the condensate, and heat washes across them. The chapter's numbers agree with one another: ideal-gas scaling in that trap gives 21.5 nK for the printed cloud.

A latch, not a label. The dissertation draws the book's distinction in one sentence: "vortices are long-lived topological states in superfluids, whereas interference fringes are subject to decay and instability."

The bin is also history. Section W2's ring, closed by one weak spot, holds its rung past the half-way point and turns over only further on, so which rung it sits in depends on where the knob came from: a relay with a dead band. A bin can record history in two ways. The dead band records the path the knob took, and does it deterministically; the frozen bin records the last draw heat made before the walls locked, and does it at random. The thaw erases both, because when the walls come down the bin forgets.

Reading a frozen number

The superconducting rings are read in place: Kirtley's by a scanning SQUID microscope, Monaco's by a junction along the ring, and the junction ring by making its fluxons move, a record silent until something traverses it (Field Guide, row 443). The atom rings are read by letting go. Vortex cores are "smaller than the optical diffraction limit", so the cloud must "expand in size by a factor of ~ 5 " before it is photographed, and Corman's and Aidelsburger's spirals appear only after the gas has expanded. So in the atoms one image is one freeze, read once, which is why atom statistics are counted in shots. In the niobium experiments the reader and the heater are disconnected during every write.

A crystal's loop

Crystals hold the same kind of number. Walk a closed loop of lattice steps around a dislocation and you come back one lattice step short; that missing step, the Burgers vector, is a whole lattice vector, and the Field Guide lists it beside the winding number, the flux quantum and the circulation quantum as one of the whole numbers a loop returns (Field Guide, draft v4.6, appendix B.2). It is written by stress: working a metal drives existing dislocations through it and makes "a great number of new dislocations". It is kept at ordinary temperatures, where heat alone is too weak to clear them. And it is erased by annealing: "the effects of cold working may be reversed by annealing the material at high temperatures where recovery and recrystallization reduce the dislocation density" (Wikipedia, "Work hardening", read 25 September, citing Degarmo, Black and Kohser 2003 and two later sources, not read). The Field Guide says the last part in its own words: "ordinary annealing repairs a crystal by erasing defects — which erases the memory" (v6, row 66). Work-hardening and annealing are freeze and thaw with a crystal's own whole numbers.

What this section does not show

The runs are stated models, not devices. The field-cooled odds assume the textbook thermal weighting at the moment of freezing. The heat route's share of shots in niobium rests on an attempt rate no niobium paper gives, so only its slope is a result. The heat model also ASSUMES a ring-shape factor (the logarithm of outer over inner radius) and the ring's inductances added in series, carries Kirtley's temperature law over to niobium, and uses a thin-strip wall where the vortex core is a quarter of the strip. The pieces' route is modelled only from pieces to a whole number, with the pieces joined at once; the narrowing by pieces that settle together is shown in a stated ring model and in the literature, not in the experiments' own rings. The Kibble-Zurek law from cooling speed to piece size is quoted, not run. The ratchet chain's fade with a spread of ring areas is a stated model; the ratchet paper reports no spread for its rings, and the model takes each ring's voltage to follow its mean current and leaves out the ring's own weakening in a strong field. The ratchet paper closes by asking whether noise at equilibrium could do the same; that would take work from a single heat bath, which the second law of thermodynamics forbids, and the book takes nothing from the question. Why Monaco's slope is 0.62 is not settled, and neither is which freeze writes a junction's count. The identity joining the two staircases and the junction's topology were checked on 25 September by this book's own computation. The magnets' two temperatures and the dislocation paragraph were checked against reference summaries whose own sources were not read (LATCH_ATLAS_WRITE_CHECK_LEDGER_2026-09-25). The experiments were read in their arXiv versions. The journal records of the Scherer letter and of Kirtley, Tsuei and Tafuri are checked; the rest are still to check. The niobium slope comes from one ring whose slowest point rests on few counts.

Sources: Schmüser's lectures in CERN-2004-008 (the copy supplied on 23 September; Doll and Näbauer 1961 as cited there, not read); the price table's superconducting WRITE cell; the Crossover Atlas, deposited version 3, Chapter 6; the Field Guide, draft v4.6 appendix B.2 and v6 rows 51, 66, 172 and 443. Read in full, in their arXiv versions: Kirtley, Tsuei and Tafuri 2003 (journal record checked 25 September); Corman and co-workers 2014; Aidelsburger and co-workers 2017; Monaco, Mygind, Rivers and Koshelets 2009; Weir, Monaco and Rivers 2012; Weir and co-workers 2013; Weir and Rivers 2011; Monaco, Mygind and Rivers 2001; Scherer's dissertation 2007 (the letter it expands, Scherer, Weiler, Neely and Anderson 2007, record and abstract checked 25 September). Also read in full: Groisman, De Vita, Fernández Bonder and Zhang 2025 and Gurtovoi and co-workers 2015 (journal records not checked); Weiler and co-workers 2008 and Kirtley and Tafuri 2009, both supplied by Bear on 25 September. Reference summaries for the magnets' two temperatures and the dislocation paragraph (Wikipedia, "Superparamagnetism" and "Work hardening"; MathWorld, "Sheppard's Correction"), their own sources not read. Bear's *A Coincidence Has No Single Price*, for its method only. Identifiers: the reference list; the reading, line by line, in LATCH_ATLAS_W4_DEEP_PASS_NOTES_2026-09-24.

Runs: WRITE_W4_FIGURES 20260925_071957 (write_w4_figures.py version 2, seed 20260924; it reprints every number of version 1's run, WRITE_W4_FIGURES 20260924_182551, which it supersedes); FREEZE_ROUTE_CHECK 20260924_192909 for Figure W4.2; and, each named where its number appears, FIELD_STAIRCASE_CHECK 20260925_052324, GAUSSIANITY_CHECK 20260925_055240, MONACO_NUMBERS_CHECK 20260925_050903, NIOBIUM_HEAT_ROUTE 20260925_053945, JUNCTION_RING_CHECK 20260925_061323, SQRT_BEND_CHECK 20260925_062608, SCHERER_CHECK 20260925_063927, SCHERER_TRAP_CHECK 20260925_070048, RING_FACTOR_CHECK 20260925_183151 (ring_factor_check.py version 2), RING_ARRAY_SPREAD_CHECK 20260925_194622 (ring_array_spread_check.py version 1) and WRITE_CHECKS 20260925_234135 (write_checks.py version 2; ledger LATCH_ATLAS_WRITE_CHECK_LEDGER_2026-09-25); the price table's PRICE_TABLE_ESTIMATES 20260923_212608 and WALL_ESTIMATES 20260923_224526.

W5. Erasing is not writing backwards

Wind a ring up one rung at a time and then back down, and the path down is not the path up. The loop between the two paths is the latch's memory, and its area is the energy each round trip spends.

Three things come out of this section. The ring of section W2, a superconducting ring closed by one weak spot, spends energy on every slip, up or down, and gives none of it back, so a round trip has a price and the price is the loop's area. Erasing a rung is not the write run in reverse: it is a second write, to the rung below, made from the other side of the knob, and it costs what the write cost. And a write that does not keep still pays: where the ring cannot hold a rung, it climbs and falls back and spends twice.

The way down is not the way up

Section W2 turned the knob up to 2.5 flux quanta, down to -2.5 and back to zero, with a weak spot of strength 3. The ring slipped eight times, one rung each time, and ended where it started. But it slipped up at 0.7545 and 1.7545 and back down at 1.2455 and 0.2455, so the path down lies below the path up, and the two paths enclose a loop. The loop encloses 1.4406 flux quanta squared, and the energy spent per round trip is 2.8854 times what one flux quantum stores in the ring's inductance: the same by the work the knob did, by the energy the ring gave up as it settled, and by the loop's area, to 0.15 per cent, the difference set by the knob's step.

That loop is the memory. Inside it the ring's rung depends on where the knob came from, which is what a latch is for; outside it there is one answer whichever way the knob came. And the loop is paid for. Every slip spends energy: 0.0415, 0.3606, 0.9092 and 2.4111 of the one-quantum store at strengths of 1.5, 3, 5 and 10. That is more than the wall between two level rungs at the same strengths, 0.0136, 0.1193, 0.2991 and 0.7861, because a slip does not climb the wall and roll down the other side: the well it sat in vanishes under it, and it falls from where the well used to be to the bottom of the next. That run recomputes the four loops of section W2's run and reprints them line for line, checked ledger against ledger.

Figure W5.1. Erasing is not writing backwards (textbook rf-SQUID model; run WRITE_W5_FIGURES)




Figure W5.1. Erasing is not writing backwards, in the textbook model of a ring closed by one weak spot. (A) At a strength of 5: the knob is turned up to 1.11 flux quanta and back, and the ring slips once and keeps rung 1 (blue); then it is turned down to -0.11 and back, and the ring slips once and returns to rung 0 (orange). The write and the erase take different roads. (B) The knob swept up and back down at four strengths: the stronger the weak spot, the wider the loop, and the loop is the memory. (C) What one slip spends, against the wall between two level rungs, in units of what one flux quantum stores in the ring. The textbook rf-SQUID model of section W2, its functions copied unchanged from that section's tool; no measured constants and no random numbers; the knob turned in steps of 0.0005 flux quanta; heavy damping assumed. Drawn by write_w5_figures.py, version 1.

Erasing is a second write

At a strength of 5, a ring on rung 0 keeps it until the knob reaches 1.0617 flux quanta, where rung 0's well ends. Turn the knob up to 1.1117 and back to zero, and the ring slips once, at 1.0620, and stays on rung 1: with the knob at zero, rung 1 still has a well, so the write keeps. Now erase it. Turning the knob back to zero does nothing, because the knob is already there and the ring is resting. Rung 1's well reaches from -0.0617 to 2.0617 on the knob, so the only way off it is past one of its ends. The near end is on the other side of zero: turn the knob down to -0.1117 and back, and the ring slips once, at -0.0620 , onto rung 0 and stays there.

So the write turned the knob 1.1117 quanta up, and the erase 0.1117 quanta down. The erase is not the write's road taken backwards; it is a second write, to the rung below, made from the other side. And it costs the same: the write spent 0.9094 of the one-quantum store and the erase 0.9091. At a strength of 10 the story repeats with wider wells: up to 1.8995 to write, one slip at 1.8500; down to -0.8995 to erase, one slip at -0.8500; 2.4112 spent on the way in and 2.4111 on the way out. Nothing the write spent comes back when the rung is erased. Both spends are lost as heat, through the textbook model's heavy damping.

A write that does not keep still pays

At strengths of 1.5 and 3, only rung 0 has a well when the knob is at zero. Turn the knob up just past where rung 0's well ends and back, and the ring slips up to rung 1 and then, on the way back, slips down again: at 0.5445 and 0.4555 at a strength of 1.5, at 0.7545 and 0.2455 at 3. It ends on rung 0, having written nothing, and it has spent 0.0829 and 0.7212 of the one-quantum store, two slips' worth. This is section W2's rule from the other end: the pen and the guard are one setting, and a weak spot too weak to keep a rung will still charge for the attempt.

The same loop in other media

A ring of cold atoms. Eckel and co-workers stirred a ring of superfluid gas with a weak link rotated around it and saw the loop directly: "hysteresis in a quantized atomtronic circuit: a ring of superfluid BEC obstructed by a rotating weak link" (Eckel et al. 2014, abstract read through a web listing on 25 September). Their knob is the weak link's rotation, and their rungs are whole units of circulation, but the object is this section's.

A bell. The book's register records the Jellyfish paper's mirror stroke (*Jellyfish: the Holonomist of the Sea*, Schulz 2026) as the same object as this loop (from the register; the paper was not re-read for this section).

The ground, as contrast. Cascadia's windings are read, never written. In the WRITE outline's words, they "keep going the same way with a median gap of two weeks and reverse only after a median of eighteen" (Cascadia paper §10.1, as the outline records). That is a read-only medium with a long memory of its own direction, not a write; it is here to show that the direction a count last moved is a quantity nature also keeps.

What this section does not show

The ring is the textbook model of section W2, not a device: no particular ring, junction or temperature, and every energy a fraction of what one flux quantum stores in the ring's inductance. The settling is heavily damped by assumption; a lightly damped ring can overshoot a rung, and that is not modelled. There is no temperature in the run, so the section does not compare what a slip spends with the thermodynamic least cost of erasing one bit (Landauer's bound of $kT \ln 2$: " $k \log_e 2 = 0.6931$ k per bit" of entropy removed, paid as "0.6931 kT per restored bit"; R. Landauer, "Irreversibility and heat generation in the computing process", *IBM J. Res. Dev.* 5, 183, 1961, read in part through a web copy of its 2000 reprint on 25 September): the ring here pays what its own store sets, at any temperature, and a warm ring can also hop a wall early, which is section W3's and W4's. The cold-atom ring was read at the level of its abstract; the bell's stroke is taken from the book's register, not re-read; and the Cascadia numbers are the WRITE outline's record of the Cascadia paper, not re-read here.

Sources: the textbook rf-SQUID model of section W2, stated at the head of both tools; Eckel and co-workers 2014 (abstract, read through a web listing on 25 September); the Jellyfish paper, through the book's register, not re-read; the Cascadia paper's §10.1, through the WRITE outline; Landauer 1961 (for the bound only, read in part through a web copy of its 2000 reprint). Identifiers: the reference list.

Runs: WRITE_W5_FIGURES 20260925_211314 (write_w5_figures.py version 1, the textbook model, no measured constants, no random numbers; it reprints the four closed loops of WRITE_W2_FIGURES 20260924_175837 line for line, checked by a line-by-line comparison of the two ledgers); WRITE_W2_FIGURES 20260924_175837 for the round trip at strength 3; the Cascadia paper's WATERFALL_MAP 20260821_200450, as the WRITE outline records it; the checking ledger LATCH_ATLAS_WRITE_CHECK_LEDGER_2026-09-25 for the Landauer reading.

W6. The same number, written two ways

Two strands can share one linking number in many ways: wound tightly around each other as twist, or thrown into loops as writhe. The count stays fixed while its two continuous halves trade, so the same whole number can be written in more than one hand, and a reader that sees only the total cannot tell which.

Three things come out of this section. A band's linking number splits into twist, how much the band turns about its own spine, and writhe, how much the spine itself coils, and bending the band moves count from one to the other without changing the total at all. Only a cut changes the total: bending trades, cutting writes. And a reader of the total is blind to the split, so the split is where a second, continuous record can hide inside a whole number.

Twist and writhe trade; the count holds

Take a thin band whose spine winds three times round a ring as it goes round, and let the band lie flat along the direction from the ring to the spine. Its two edges link -3 times. Now let the spine coil out further from the ring, from almost nothing to nearly the ring's own radius. With the spine coiled out by 0.02 of the ring's radius, the count is almost all twist: twist -2.9952 , writhe -0.0048 . At 0.40 it is split: twist -1.9265 , writhe -1.0735 . At 0.95 most of it is writhe: twist -0.9769 , writhe -2.0230 . At every step the sum is -2.9999 on the run's 2000 points, and the edges' linking, read on its own by the Gauss integral without reference to twist or writhe, stays within 7.85×10^{-5} of -3 ; the share of the count carried as writhe runs from 0.002 to 0.674. READ section 3 showed the middle of this family, coil radii 0.15 to 0.55, at 3,000 points, and this run reprints those five lines exactly, checked ledger against ledger, with the linking and writhe functions copied unchanged from that section's tool.

Figure W6.1. One linking number, written in different hands (pure mathematics; run WRITE_W6_FIGURES)

A One count, three hands, and one cut: each bar is twist plus writhe; the line is the linking




Figure W6.1. One linking number, written in different hands. (A) Four bands, each sketched above its bar: the spine coiled out by 0.05, 0.40 and 0.90 of the ring's radius, and the 0.40 band again after one cut. Each bar stacks the band's twist (blue) on its writhe (orange); the black line is the edges' linking, read on its own. The first three bars reach the same line at -3 however they split; the cut moves the line to -2 . (B) Twist and writhe against how far the spine coils out: they slide in opposite directions while the linking stays on -3 . Pure mathematics, no physical constants, no random numbers. Drawn by write_w6_figures.py, version 2.

The theorem behind the picture says that for any closed band the linking number of its edges equals its twist plus its writhe. Moffatt and Ricca argue that it should carry Călugăreanu's name (their 1992 paper, read in full for the book on 23 September), and the book uses that name. Dennis and Hannay's account of the theorem names it the same way, citing them: "Călugăreanu's theorem (Călugăreanu 1959, 1961; Moffatt & Ricca 1992) (also referred to as White's formula ...)" (M. R. Dennis and J. H. Hannay, "Geometry of Călugăreanu's theorem", arXiv math-ph/0503012, read in part 25 September). Twist and writhe are each continuous and each can take any value; only their sum is forced to be whole.

Bending trades; cutting writes

To change the total, the band has to be cut. Cut it, turn one end once and rejoin it, section W1's write, and the linking number moves from -3 to -2 whatever the coil: -2.0000 at coil radii of 0.05 and 0.40 and -1.9999 at 0.90 , with the extra turn landing in the twist each time (twist -1.9703 , -0.9265 and -0.0257) while the writhe stays where the coil put it. So two different operations act on a band. Bending, which never breaks it, moves count between twist and writhe and writes nothing. Cutting, the only pen, changes the total by a whole number. A reader who sees only the total sees the second and never the first.

The living example

DNA is the band this section is about. Its two strands are its two edges, its linking number is set by enzymes that cut (section W1: gyrase in steps of two, the type I enzymes in steps of one or a few whole turns; Brown and Cozzarelli 1979, read for the price table, and Wikipedia, "Type I topoisomerase", read 25 September). Between cuts the molecule is free to trade twist for writhe, coiling round itself into an interwound braid, a plectoneme, or wrapping into a toroid: "Supercoiled DNA forms two structures; a plectoneme or a toroid, or a combination of both", and "twist and writhe are interconvertible" (Wikipedia, "DNA supercoil", read 25 September; its sources not read). The book's register records that DNA's store lives in that shape, in the bending and twisting (its row "DNA's shape is where the store lives", from Bear's intuition of 23 September), and that cells hold the total at a setpoint with two opposing pens. What a cell does with the count and its split is the living use of this section. Hold the linking number below its relaxed value, and the strands open more easily where a gene must be read: opening a stretch takes up the deficit as lost twist and lets the stored writhe relax. The same summary puts it as "Negative supercoils favor local unwinding of the DNA, allowing processes such as transcription, DNA replication, and recombination" (Wikipedia, "DNA supercoil", read 25 September). The mechanism in the middle clause is the theorem's arithmetic, not a quotation. The count is the latch; the split is a working variable carried inside it.

Does anything read the split?

The WRITE outline left this as a question. One answer the book can point to, recalled, with the paper's record checked but its text not read: a single DNA molecule held stretched between a surface and a small magnet, and turned by turning the magnet, reports its writhe through its length. The first turns go into twist and the length barely changes; past a buckling point each further turn goes into a loop of writhe and shortens the molecule by a fixed step (magnetic tweezers; T. R. Strick, J.-F. Allemand, D. Bensimon, A. Bensimon and V. Croquette, "The elasticity of a single supercoiled DNA molecule", *Science* 271, 1835, 1996, record checked 25 September, abstract not reached). There the total is counted by the magnet's turns and the split is read by a ruler. None of the book's own readers does this: the ring read of READ and of the vortex runs reads the total, and the lean of its spiral is the sign of the total, not of either half.

What this section does not show

The band is mathematics: a thin band with no stiffness, no thickness and no energy, so the section says nothing about which split a real band chooses. That choice is set by how stiff the band is to bending and to twisting, and it is the physics of supercoiled DNA and of twisted cables, not run here. The DNA statements were checked on 25 September against a reference summary whose own sources were not read. The theorem's name rests on the book's reading of Moffatt and Ricca. The magnetic-tweezers result is recalled; its paper's record is checked and its text is not read (LATCH_ATLAS_WRITE_CHECK_LEDGER_2026-09-25). The run's writhe and twist are computed on 2000 points, and the sum sits at -2.9999 rather than -3 at that resolution; READ's 3,000 points give -3.0000 .

Sources: the band, its twist, its writhe and the Gauss linking integral, this book's own computation; section W1 for the cut and the enzymes' steps (Brown and Cozzarelli 1979, read for the price table); the register rows "DNA's shape is where the store lives" and the cell's two opposing pens; Moffatt and Ricca 1992 (read in full on 23 September) for Călugăreanu's name; Dennis and Hannay 2005 (read in part); Wikipedia, "Type I topoisomerase" and "DNA supercoil" (read 25 September, their sources not read); Strick and co-workers 1996 (record checked, text not read; the checking ledger LATCH_ATLAS_WRITE_CHECK_LEDGER_2026-09-25). Identifiers: the reference list.

Runs: WRITE_W6_FIGURES 20260925_220239 (write_w6_figures.py version 2, pure mathematics, the linking and writhe functions copied unchanged from read_r3_figures.py; it draws a simpler Figure W6.1 and reprints version 1's run, WRITE_W6_FIGURES 20260925_214049, line for line); READ_R3_FIGURES 20260924_103544, whose five coiled-band lines this run reprints exactly, checked by a line-by-line comparison of the two ledgers.

W7. Waves write labels, and a label can write a latch

A ring of emitters, each fed a little later than its neighbour, writes a twist into a wave: a whole number that flies away and holds nothing. But when that wave lands on matter whose phase must close on itself, it can hand its whole number over, and the matter keeps it. The wave is the ink; matter is the page.

Three things come out of this section. A ring of N emitters writes a twist into its beam, but only the feed folded into a window of about $N/2$ either side of zero, and at exactly half a turn per step an even ring writes no twist at all, only a pattern with dark spokes. The same fold is the READ part's half-turn limit, met from the write side. And a twist written into light has been handed to a ring of superfluid gas and kept there as a latch, which is how a label becomes a latch.

A ring of emitters writes a twist, folded

Take eight emitters on a ring one wavelength in radius, each fed a little later than the one before, so that the feed winds once, twice or more round the ring. Far away, read the beam round a cone about the ring's axis with the book's loop read. Fed one, two and three turns, the beam carries one, two and three: the ring writes what it is fed. Fed four, half a turn from each emitter to the next, the beam carries none, and its field has eight dark spokes, where its weakest point is 1.4×10^{-16} of its strongest. Fed five, six and seven, the beam carries -3 , -2 and -1 , and fed eight, a whole turn per step, it carries nothing: every emitter is back in step with its neighbours. So the ring writes its feed folded into the window from -3 to $+3$; the feed and its fold are the same feed to the emitters, which see only the step between neighbours.

Figure W7.1. A ring of emitters writes a twist into a wave, folded at half a turn (pure wave arithmetic; run WRITE_W7_FIGURES)




Figure W7.1. A ring of emitters writes a twist into a wave, folded at half a turn. (A) The winding the beam carries against the winding fed round a ring of 8 emitters (blue) and of 7 (orange); the crosses mark the two feeds for which the 8-emitter beam has dark spokes and no twist. (B) Purity, the beam's weakest point over its strongest on the brightest cone: a clean twist is 1, and purity falls as the feed step nears half a turn. (C) The beam of 8 emitters fed three turns: a bright ring, rippled by its folded twin. (D) Fed four turns: eight dark spokes, the twist and its mirror in equal parts. Pure wave arithmetic: ideal point emitters on a ring one wavelength in radius, the far field summed exactly, each beam read on its brightest cone at 720 points; no measured constants, no random numbers. Drawn by write_w7_figures.py, version 1.

The beam's purity tells how clean a label the ring has written. Fed one turn, the 8-emitter beam is 0.9996 pure; fed two, 0.9529; fed three, 0.3630, rippled by the folded twist it also carries. A ring of seven emitters has no feed of exactly half a turn per step, so it never writes dark spokes, but its top twist is barely a twist: fed three turns, purity 0.0605, and fed four, it writes -3 . An engineer who wants a clean twist of 1 keeps well below $N/2$. The 2008 twisted-sound experiment did exactly that, with four loudspeakers for a twist of one and eight for a twist of two, a quarter turn between neighbours in both (Volke-Sepúlveda, Santillán and Boullosa, Phys. Rev. Lett. 100, 024302, 2008, read at summary level). The price table's radio column, and the register's row on the ring of radio antennas, meet the same wall at four with a ring of eight antennas.

The same fold, read and written

This is the READ part's half-turn limit from the other side. A ring of N readers reads a count only up to $N/2$ before the turns fold over, and at exactly half a turn per step it cannot tell a count from its mirror; a ring of N writers writes a twist only up to $N/2$, and at half a turn it writes the twist and its mirror together. The book's run of read heads this week found the reader's version of the dark spokes: a 4-point head cannot read a count that leaves half a turn per step, and odd heads have no such blind spot. The writer's version is here: an even ring of emitters has one feed that writes no twist, and an odd ring has none. On both sides, the step between neighbours is all the ring knows.

Other pens for twisted waves

A ring of emitters is one pen. A hologram is another: Daria, Palima and Glückstad write the twist on a spatial light modulator, up to 20 turns of phase around the beam (Optics Express 19, 476, 2011, read in full). A Dove prism turns a twist of l into its mirror, $-l$, with nothing to climb, and for tightly focused beams the flip smears into a spread of twists centred on $-l$ (González, Molina-Terriza and Torres, Optics Express 14, 9093, 2006, read in part). Every one of these writes a label. The twist rides in the beam and leaves with it; nothing holds it back, and there is no wall to climb in writing it or in turning it into its mirror.

A label can write a latch

The path from label to latch needs a receiver whose phase must close on itself. An ordinary object will not do: the same twisted light drives silica beads along a spiral, a particle revolving "with continuous transfer of OAM" (Daria et al. 2011), and a turning bead keeps a speed, not a whole number. A superfluid will. Twisted light has written one unit of angular momentum per atom into a ring-shaped superfluid of atoms, which then held it as a latch for up to 10 seconds (Ryu et al., Phys. Rev. Lett. 99, 260401, 2007, abstract read, as the price table's superfluid and light columns record). A year earlier Andersen and co-workers had shown the transfer itself in a condensate of sodium atoms: "the coherent transfer of the orbital angular momentum of a photon to an atom in quantized units of $\hbar$ ", and "vortices of charge 2 by transferring to each atom the orbital angular momentum of two photons" (M. F. Andersen, C. Ryu, P. Cladé, V. Natarajan, A. Vaziri, K. Helmerson and W. D. Phillips, "Quantized rotation of atoms from photons with orbital angular momentum", Phys. Rev. Lett. 97, 170406, 2006, abstract read 25 September). The light writes one whole number per photon's twist, and two photons write two.

That is the section's point, and the book's rule, a register row of its own, that every label is asked whether it is a path to a latch, a wave writing a latch into matter. The wave's whole number is the ink, written by a ring of emitters or a hologram and folded by the same half-turn rule that limits every ring of readers. The matter's whole number is the page: it keeps what it is given, because its phase must close, and it keeps it against the walls of sections W3 and W4.

What this section does not show

The emitters are ideal points of one strength in one plane, their far field computed by simple wave arithmetic; no antenna, loudspeaker or hologram is modelled, and each beam is read on its brightest cone only. The twisted-sound experiment was read at summary level. Ryu et al. 2007 and Andersen et al. 2006 were read at the level of their abstracts. Whether a twisted microwave or sound beam has ever written a winding into a superconducting ring or a ring superfluid has not been searched; it is held in the book's verification queue.

Sources: the price table's light, radio and sound columns and the literature they read: Daria, Palima and Glückstad 2011 (read in full); González, Molina-Terriza and Torres 2006 (read in part); Volke-Sepúlveda, Santillán and Boullosa 2008 (summary level); Ryu and co-workers 2007 and Andersen and co-workers 2006 (abstracts, the second read 25 September); the register rows on the ring of radio antennas and on a wave writing a latch into matter; the READ part's half-turn limit. Identifiers: the reference list.

Runs: WRITE_W7 FIGURES 20260925_220409 (write_w7_figures.py version 1, pure wave arithmetic); the price table's RADIO_ESTIMATES 20260924_094021 and SOUND_ESTIMATES 20260924_094237; the read-head run RING_REMAINDER_HEAD_CHECK 20260925_213837.

W8. Engines, regulators and locks: when the pen never stops

Some whole numbers are written once and kept. Others are rewritten every cycle, by a drive that never stops, and they last only while it runs. Between the two sit the regulator, which shuttles energy and turns nothing, and the engine, which turns a load because its two layers are out of step.

Four things come out of this section. Two layers of the same pattern turn a load in proportion to the sine of how far out of step they are: in step they only shuttle energy, a regulator; a quarter cycle apart they turn, an engine. A lock between two rates holds while the pull is stronger than their difference, and its lag reaches a quarter turn exactly at the edge of the lock, the stall, beyond which it slips one whole turn at a time; a crowd of rates under one pull locks all at once past a threshold. A lock held by a drive is not a latch: its whole number lasts only while the drive runs. And a driven ring can write a sign it takes from its own shape: an engine that writes.

Regulator or engine

Take two layers of the same three-lobed pattern round a ring, a quarter of a lobe apart in space, each breathing in time, the second offset from the first by a time δ . Sort the combined motion into the part that runs forward round the ring and the part that runs back. The forward share is $(1 + \sin \delta)/2$: 0.5000 in step, 0.7500 at 30° , 0.8536 at 45° and 1.0000 at 90° . In words, the net turning goes as the sine of the offset. In step, half runs each way and nothing turns: energy shuttles back and forth, as it does in a regulator. A quarter cycle apart, everything runs one way: that is the two-phase motor, an engine. It is the same sine as the current through a Josephson junction and the power a generator sends down a line against its angle, and it is the register's row REGULATOR OR ENGINE, RINGS THAT RING, AND LOCK run through.

The lock, and the stall

Two dials with their own rates, each pulling the other with the sine of the angle between them, is the textbook lock. With the pull set to 1, dials whose rates differ by 0.25, 0.5, 0.75, 0.9 and 0.99 settle into step with the one lagging the other by 0.2527, 0.5236, 0.8481, 1.1198 and 1.4293 radians, each equal to the textbook arcsin of the difference over the pull. The lag grows as the dials are pushed apart and reaches a quarter turn, 1.5708 radians, only when the difference equals the pull. Past that the lock breaks and the lag runs, one whole turn at a time, at the textbook rate of $\sqrt{(\text{difference}^2 - \text{pull}^2)}/2\pi$ turns per unit time: 0.02256 at a difference of 1.01 and 0.27562 at 2.0, against the textbook's 0.02256 and 0.27566 (in a longer run, over 4,000 time units).

That quarter turn is the stall. The pull between two phases is a sine, strongest at a quarter turn; a lock can hold any difference up to the strongest pull it has, and holding the largest one uses the whole of it, at a quarter turn. Push a little harder and there is no stronger pull to answer with. It is the edge of the lock that the literature calls the edge of an Arnold tongue, the point where a Josephson junction's current reaches its critical value and its phase starts to run (both textbook), and the stall that Bear drew on an S rotor on 25 September ("the quarter phase is the stall break").

Many dials, one pull

Give 2000 dials their own rates, drawn from a bell curve of width 1, and pull each toward the phase of the whole group with a strength K times the group's own coherence, the standard model of synchronisation (Kuramoto's). Below a threshold the group stays scattered: at $K = 1.5$ its coherence is 0.0968 and 0.0220 of the dials are locked. Just past the textbook threshold of 1.5958 it snaps together: at $K = 1.75$ the coherence is 0.4924 and 0.6160 of the dials are locked, at 2 it is 0.7176 and 0.8535, and at 3 it is 0.9266 and 0.9955. At $K = 3$, 1991 of the 2000 dials are locked, each lagging the group by the arcsin of its own rate over K times the coherence, to within 0.0045 radians, and the largest lag is 0.2238 of a turn: every locked dial sits inside the stall. The nine that are not locked are exactly those whose own rates lie beyond the strongest pull the group can give, 2.7799. The crowd's lock is the two dials' lock, shared.

Figure W8.1. Locks: two dials, then a crowd, sharing one pull (pure mathematics; run WRITE_W8_FIGURES)




Figure W8.1. Locks. (A) Two dials under one pull: the lag while locked (blue, turns) grows toward a quarter turn and reaches it only at the edge; past the edge the lag slips whole turns at a steady rate (orange). Points from the run, lines from the textbook. (B) 2000 dials with their own rates, one shared pull: the group's coherence and the share of dials locked, against the pull's strength; the dotted line is the textbook threshold. (C) At a pull of 3, the lag of every locked dial behind the group: all inside a quarter turn either way. Pure mathematics: the Adler lock integrated by RK4 in steps of 0.01; the Kuramoto model with 2000 dials, their own rates from a bell curve of width 1 drawn with seed 20261007, Euler steps of 0.05 to 300 time units, averages over the last 100. Drawn by write_w8_figures.py, version 1.

A powered lock is not a latch

A Josephson junction held at a voltage rings at exactly 483.5978 GHz per millivolt. Drive it with a microwave tone and it locks to the tone, and its voltage sits on whole-number steps, 144.75 microvolts apart at an assumed 70 GHz. NIST's voltage standard locks about 300,000 junctions to one tone to give volts (NIST, "A Primary Voltage Standard for the Whole World", news release, 2013, read, as the price table records). A series array of junctions, each with its own pitch, is also a crowd of dials under one pull: Wiesenfeld, Colet and Strogatz mapped it onto the synchronisation model of the last section (Phys. Rev. Lett. 76, 404, 1996; the mapping recalled, the record checked, the abstract not read, because the record carries none). The whole number on a voltage standard is real and exact, but it is written by the tone's frequency and held by the tone: turn the tone off and the steps are gone. That is the powered lock page one promised to show, and it is the example of what a latch is not. A latch keeps its number with the drive switched off.

An engine that writes

Give the book's ring of phases the two things a Savonius rotor has, a handedness (each phase pulling harder on its forward neighbour than on its backward one) and a drive (a lag in the pull, which keeps the ring out of equilibrium). With either alone, plus and minus windings stay equally likely, as symmetry requires; with both, rings of 358 and 359 phases all freeze with the sign the handedness and the drive choose, at a mean winding of 13.567 and 13.577 turns. And a twisted state of such a ring keeps turning, every phase advancing at the steady rate $(1 + a) \sin(k - \alpha) - (1 - a) \sin(k + \alpha)$, which is not zero when there is a lag (derived, not measured). So it is an engine, and one that writes: it keeps running while it holds a count whose sign came from its shape. The same model shows the stall moved by the shape: the window of stable twists, a quarter turn either side of its centre, slides to a centre at $\arctan(a \tan \alpha)$, where a is the handedness and α the lag.

Two machines that hold all of it

A mechanical watch. The mainspring is the store: wound once a day, it holds the energy. The balance wheel and its hairspring are the regulator: they swing back and forth at their own rate, shuttling energy between the wheel's motion and the spring's coil, and turn nothing by themselves. The escapement is the latch that counts: at each swing it unlocks the gear train for one tooth and locks it again, and on the way it gives the balance a small push to make up what friction took. Between swings the train is locked; the hands' position is a count of teeth released; and the rate belongs to the balance, not to the mainspring, so a fuller spring swings the balance wider rather than faster. That holds for the ideal: "the ideal sprung balance is inherently isochronous". A real escapement and hairspring leave a small dependence on the swing, and as the spring runs down "the amplitude tends to decrease". That is why watchmakers built the "fusee and chain, the remontoire, and the simple stopwork" to even out the spring's push (SJX, "In-Depth: Understanding Isochronism and Oscillators", watchesbysjx.com, December 2025, read 25 September). Bear, reading this picture on 25 September, set a hand tool beside it: a ratchet wrench's tooth count, and a socket's

six points against twelve, are chosen to keep a grip without slipping or rounding the bolt head. More teeth give a finer step and a shorter swing; fewer and bigger teeth hold more before they slip. That is Bear's experience, and a tool maker's buyer's guides say the same. The swing is 360° over the tooth count, so "a 72-tooth ratchet has a 5° swing arc", and "more teeth means each individual tooth is smaller, which could theoretically reduce strength", though alloy steels narrow the gap. A six-point socket bears on a bolt head's flats, not its corners, with "significantly less chance of rounding off the fastener"; a twelve-point one catches every 30° instead of every 60° and bears on the corners (Capri Tools' buyer's guides "How Many Teeth in a Ratchet? 72 vs 90 Tooth Count Explained" and "6-Point vs 12-Point Sockets and Wrenches: When to Use Each", read 25 September). The twelve-point socket trades grip for a finer step, the same trade as the ratchet. It is the balance this book keeps meeting: section W2's weak spot, where a pen that writes easily keeps little and a strong one keeps more and costs more, and the read heads of the READ part, where a head with more points reads further but has a thinner margin at its largest step.

A cell's DNA setpoint. A cell holds its DNA's supercoiling at a setpoint with two opposing pens: gyrase writes the linking number down, "converting circular DNA to a negatively supercoiled form" (Menzel and Gellert, below), and topoisomerase I relaxes it back up, since the bacterial enzyme "can only relax negative supercoiled DNA" (Wikipedia, "Type I topoisomerase", read 25 September, citing Champoux 2001, not read). Each pen's supply answers to the supercoiling itself. Menzel and Gellert showed it for gyrase: "the synthesis of DNA gyrase is itself controlled by DNA supercoiling; synthesis is highest when the DNA template is relaxed", and they concluded that "DNA supercoiling in *E. coli* is controlled by a homeostatic mechanism" (Cell 34, 105, 1983, abstract read). Later work reports the same for both enzymes (Snoep and co-workers, "DNA supercoiling in *Escherichia coli* is under tight and subtle homeostatic control, involving gene-expression and metabolic regulation of both topoisomerase I and DNA gyrase", *Eur. J. Biochem.*, 2002, title seen). In a control engineer's terms this is a push-pull loop: two actuators working against each other on one variable, each driven harder the further the variable strays to its side. The book's register (its row THE CELL HOLDS DNA AT A SETPOINT WITH TWO OPPOSING PENS) adds a picture: the supercoiling density acts like a voltage and the turns like charge, a bus held at a fixed voltage by a charger and a load.

How equipment runs such a loop. Real control loops, Bear notes from his working life in control systems (25 September; his field experience, not a citation), keep a deadband: a band round the setpoint inside which nothing switches, because a loop that hunts wears its equipment and burns energy on every cycle. The band is set from above, and widened when the stakes change. Two of this book's sections meet there in practice. The deadband is section W5's loop, set on purpose: the way up is not the way down, and a wider band means fewer paid trips and looser control. The supervisor who widens it is section W4's write-enable, set from outside the loop. Whether a cell keeps such a band between its two enzymes, or runs both at once and pays for the stiffness, the book does not know; it is a clean question for someone who studies supercoiling.

The ground, as contrast

Cascadia's windings, read and never written, cycle through three states without circulating: the net turning is -14.0 inside a null band of -67.0 to +63.3 (as the WRITE outline records from the Cascadia paper). In this channel the ground is not an engine.

What this section does not show

The dials are mathematics: the textbook lock and the textbook synchronisation model, with no noise, no mass and no particular device. In the lock run the slip rates just past the edge are counted over only a few slips and sit a little off the textbook values (0.02430 against 0.02256 at a difference of 1.01), which is why the section quotes the longer run of check C for them. The S-rotor ring's drive is a stand-in for a real one-way process, and whether any ring of atoms or superconductor carries both a handedness and such a lag is a question for a specialist. The watch and the DNA setpoint are pictures in words; Bear judged both on 25 September from how equipment behaves and added the ratchet and the deadband, which are his field experience, not citations; the watch's rate, the enzymes' roles and the tooth counts were checked on 25 September against a watch journal, a reference summary and a tool maker's guides, not against primary sources (LATCH_ATLAS_WRITE_CHECK_LEDGER_2026-09-25), Menzel and Gellert were read at the level of their abstract, and Snoep and co-workers and Wiesenfeld and co-workers at the level of their titles and records.

Sources: NIST 2013, read for the price table; Wiesenfeld, Colet and Strogatz 1996 (record checked, not read); Menzel and Gellert 1983 (abstract); Snoep and co-workers 2002 (title seen); the register rows REGULATOR OR ENGINE, RINGS THAT RING, AND LOCK and THE CELL HOLDS DNA AT A SETPOINT WITH TWO OPPOSING PENS; the WRITE outline for the Cascadia contrast; Bear's field experience of 25 September for the ratchet, the deadband

and supervisory control; checked on 25 September against SJX 2025, Capri Tools' two buyer's guides and Wikipedia, "Type I topoisomerase" (the checking ledger LATCH_ATLAS_WRITE_CHECK_LEDGER_2026-09-25). Identifiers: the reference list.

Runs: WRITE_W8_FIGURES 20260925_225517 (write_w8_figures.py version 1, pure mathematics, seed 20261007) for the lock and the crowd; ROTOR_RING_CHECKS 20260923_232715 for the two layers (check B) and the long lock run (check C); SC_PHASE_ESTIMATES 20260924_093814 for the junction's pitch and steps; RING_S_ROTOR_CHECK 20260925_202830 for the engine that writes; the Cascadia paper's WATERFALL_MAP 20260821_200450, as the WRITE outline records it.

Part Four: the new media, first runs

The Latch Atlas. The first runs for the media built after the reading draft: liquid crystals, the power grid, plasma, magnetic skyrmions, and the Sun, read from real maps of its magnetic field through two great flares. Their cards are in Part Two; here are the runs, their figures and the walls they met, with the people who met the same walls first.

Liquid crystals: a count in halves

The molecules of a nematic liquid crystal line up along a direction that has no head and no tail. A measurement gives that direction only up to half a turn, and round a defect it can turn by half a turn in one lap. The run asked how the book's loop read copes (run NEMATIC_HALF_READ_CHECK 20260926_053148).

A count in halves needs its own read, and the rule is simple. READ's loop walk takes each step the short way, within half a turn either side. For a field of lines, each step must be taken within a quarter turn either side, because a line turned half a turn is the same line. With that one change the read was right in every case: seven strengths from $-3/2$ to $+3/2$, loops of 8, 16 and 64 points, and 1,000 random choices of which end is the head for each. READ's own arrow read never returns a half. With the heads chosen at random it scatters over whole numbers: for a $+1/2$ defect on a loop of 16 points it returned everything from -6 to -1 .

The price of the headless read is half the margin. Put noise on every sample and the headless read first goes wrong more often than 1 in 100 at 0.3 radians, on every loop size tried. The arrow read, on a true arrow field, holds to 0.6 or 0.7 radians. That is READ section 5's half-turn limit, halved, because a headless field's whole turn is half a turn.

Halves add like whole numbers. A loop round a $+1/2$ and a $-1/2$ reads 0; round two $+1/2$ defects it reads $+1$; round a $+1/2$ and a $+1$ it reads $+3/2$. A loop round neither reads 0.

An arrow carried round a half needs two trips. Carry an arrow round a $+1/2$ defect, keeping at each step whichever end of the line is nearer the arrow's last direction. After one lap the arrow points the opposite way, turned half a turn; after two laps it is home. The same happens round $+3/2$, while round $+1$ it is home after every lap. This is the same two-trip pattern as the spinor on page 51 of Berry's paper. The half never breaks the count: a second trip pays it back.

Figure N1. Reading a count in halves: a field of lines with no heads (pure mathematics; run NEMATIC_HALF_READ_CHECK)




Figure N1. (A, B) The lines of a nematic round defects of strength $+1/2$ and $-1/2$: one lap of the blue loop turns them by half a turn. (C) Noise on every sample: the headless read (blue) goes wrong at about half the noise that an arrow read of a true arrow field (orange) can take, on a loop of 16 points. (D) An arrow carried round the defect: round $+1/2$ and $+3/2$ it is reversed after one lap and home after two; round $+1$ it is home every lap. Pure mathematics; run NEMATIC_HALF_READ_CHECK 20260926_053148.

What the colloquium changes: halves in a thin layer, a parity in the bulk. The review of record (Alexander, Chen, Matsumoto and Kamien, *Rev. Mod. Phys.* 84, 497, 2012; read in full on 26 September) gives this run's result as the classification of a flat layer. There the strength is a quarter of the number of dark brushes between crossed polarisers, it comes in halves, and halves add. So the run reproduces the textbook by computer; it does not find anything new. In the bulk, where the molecules can tip out of the plane, the count changes kind. "No distinction can be made any longer between positive and negative strength defects": a line defect is either there or not, two halves make nothing, and a strength-one line escapes into the third dimension. "Halves add like whole numbers" therefore holds in a thin layer only. A whole number does survive in the bulk, at points. A point defect, the hedgehog, carries the solid angle its surroundings sweep, in whole spheres, the same count as the fifth domain's skyrmion. But its sign

reverses if it is carried once round a line defect, so reading it needs the path it took. The colloquium names three pens: cooling, which throws down many defects that then coarsen; a laser; and shape, where a small particle whose surface holds the molecules pointing outward forces a matching defect into the liquid round it.

The pen: light, one site at a time. Tkalec, Ravnik, Čopar, Žumer and Mušević (*Science* 333, 62, 2011; read in full on 26 September) tie the defect lines round small spheres into knots and links with laser tweezers. The laser melts the order at one spot and lets it re-form, so the lines there are cut and rejoined. That is W1's cut, pass and rejoin, done by W4's freeze and thaw at a single site. Where two lines meet, each crossing site can be switched between exactly three states, and on an array of spheres every site can be switched on its own, reversibly. On a 4 by 4 array the choices give almost 40 kinds of knots and links, the Borromean rings among them, built to order. The lines are three-sided strips, so their twist comes in thirds of a turn. The paper counts it with a self-linking number that it ties to Berry's phase: halves from a line with no head, thirds from a strip with three sides. One more point bears on the tilt thread in the register. The cell's walls are rubbed at right angles to twist the liquid crystal, and "Using twisted director structure is essential for the stability of nematic knots and links". The twist also picks the hand: "left-handed crossings (1) are favored in a left-twisted nematic profile".

For the card, still to come. The count, the read and the pen are in hand. The store, wall and leak rows are not priced in either paper; the card will say so.

Sources: Alexander, Chen, Matsumoto and Kamien 2012 and Tkalec and co-workers 2011, both read in full on 26 September, their quotations from their texts (LATCH_ATLAS_DOMAIN_READING_NOTES_2026-09-26); Berry 1984 for the spinor's two trips. Identifiers: the reference list.

Runs: NEMATIC_HALF_READ_CHECK 20260926_053148 (tools/nematic_half_read_check.py version 1 with RUN_NEMATIC_HALF_READ_CHECK.bat; pure mathematics; seed 20261009). Every number in the section is printed in its ledger.

The power grid: a whole number round one loop

The model is the book's own ring of phases, used as the textbook picture of buses joined by lines. It is a ring of N buses joined by equal lines, with nothing injected or drawn at any bus, no losses, no inertia, and heavy damping. It is a stated model, not a grid (run GRID_LOOP_FLOW_CHECK 20260926_053413).

A ring holds every winding with less than a quarter turn per line, and no other. This was checked by the stability of each steady state for every ring from 4 to 24 buses. A ring of 12 buses holds five windings, from -2 to $+2$, and a ring of 24 holds eleven, from -5 to $+5$. Each winding sends its own flow round the loop with nothing injected anywhere. In the ring of 12, winding 1 carries half of a line's greatest power in every line, and winding 2 carries 0.866 of it.

A weak line is the pen, as W2's weak spot is. Settle a ring of 12 buses at winding 2 and slowly weaken one line. The count slips by exactly one turn at a time: from 2 to 1 when the line's strength reaches 0.835, and from 1 to 0 at 0.405, never two at once. It is W2's gearbox, in a grid.

The stall belongs to the loop, not to the line. The weak line reaches its own quarter turn at a strength of 0.8413 for winding 2, but the ring holds a little further, to 0.8398, because the other eleven lines, in series, carry it. The slip comes at the first sweep point below that exact edge. For winding 1 the numbers are 0.4154 and 0.4069. So inside a loop, one line can sit a little past its quarter turn while the rest of the loop holds it. That is W8's stall refined for a loop, and in the model it is exact.

Higher windings sit nearer the stall. Give every bus a random knock of 0.6 radians and let the ring settle. Winding 2 changes in 0.341 of the knocks, winding 1 in 0.006, and winding 0 in none. Winding 2 carries the most flow and has the least margin.

Figure G1. A whole number of turns round one loop of a grid, in a stated model (run GRID_LOOP_FLOW_CHECK)




Figure G1. (A) The windings a ring can hold, against the number of buses: every winding under a quarter turn per line (orange), and no other. (B) The flow each winding sends round a ring of 12 buses, as a share of a line's greatest power. (C) One line weakened slowly: the count slips one turn at a time, just past each exact edge (dashed). (D) Random knocks on every bus: the share that change the winding, for windings 0, 1 and 2. Stated model; run GRID_LOOP_FLOW_CHECK 20260926_053413.

Where to point a specialist. The model says a loop of a grid can hold a whole number, circulate power with nothing injected, and slip one turn at a time at a weak line. Three questions go to a power-systems engineer or grid operator. Do real loop flows stay put like this? Does a trip that clears one look like a slip of a whole turn? And is the winding a latch, or a powered lock that lasts only while the grid is energised?

What the published paper changes: in a real grid the pen is the switch. Coletta, Delabays, Adagideli and Jacquod (*New J. Phys.* 18, 103042, 2016; read in full on 26 September) count the same whole number, the voltage's turns round a loop, and their 18-bus ring carries at most nine windings, which is what this run finds for a ring of 18. They find three ways to write one. The first is the run's own: load a line until the state with no winding gives way. It is the case of their supplement's Figure S1, and they judge it "of little relevance for power grid operation", because it needs a line to carry close to a quarter turn, beyond its thermal limit. The pens that matter are switching operations. A tripped line can leave a winding behind, and reclosing a line writes one more turn each time the angle across the open gap has passed an odd number of half turns. So the grid is the book's plainest case of W1's cut, pass and rejoin. The switch cuts the loop, and the reclose writes the count. The rule it follows is READ section 5's half-turn limit, used to write: the rejoined line takes the short way, and the loop keeps the rounding. Once written, a winding held in their models with losses included and when conditions went back. It raised the losses by almost half in one model and doubled them in a model with the UK grid's layout. If the extra supply is cut back, the count still holds and the whole grid shifts its frequency instead, which they say "would be intolerable in a real power grid". That answers the three questions in part. The winding stays put in the models. A trip or a reclose can write it, and a clearing trip is not itself a whole-turn slip. And a winding that lives only in a driven, synchronised grid and must be paid for in losses points to a powered lock, whose count is protected while the drive runs. The authors fence their own words: "there is no quantumness in the power flow problem", and the barrier they call an "energy" is a Lyapunov function, not a store. They also say the usual definition of the recorded loop flows, up to a gigawatt in North America, is "ambiguous". So whether real loop flows are windings at all stays with a power-systems engineer, together with the latch-or-lock reading.

Sources: Coletta, Delabays, Adagideli and Jacquod 2016, its quotations from the arXiv text, read in full on 26 September, and its journal details from its DOAJ record, seen the same day. Identifiers: the reference list.

Runs: GRID_LOOP_FLOW_CHECK 20260926_053413 (tools/grid_loop_flow_check.py version 1 with RUN_GRID_LOOP_FLOW_CHECK.bat; a stated model; seed 20261010). Every number in the section is printed in its ledger.

Plasma and the Sun: a count of linked flux, weighted by flux

In a hot plasma, such as the Sun's outer layers, the magnetic field is carried along with the gas, so a tube of magnetic flux keeps its identity the way a vortex line does in water (textbook; to be cited when plasma's papers are read). Two such tubes can be linked, and a tube can be twisted. The quantity that keeps track of both is the magnetic helicity; readers who know it by another name will recognise the helicity integral as the abelian Chern-Simons form, a name the book uses nowhere else (textbook, recalled). The run asked what whole number sits inside it (run FLUX_TUBE_HELICITY_CHECK 20260926_145656).

The model. Each tube is drawn as a bundle of fine strands laid round its centreline, each strand carrying an equal share of the tube's flux. For such a bundle the helicity is a sum over every pair of different strands: the flux of one, times the flux of the other, times the number of times the two strands link, with each pair counted both ways. So the helicity is built entirely out of whole numbers, and the fluxes only weight them. This is Berger and Field's own way of writing it (1984, their equation 16, read in full on 26 September), so the run reproduces their result by computer; it does not find it. Every linking number was read with READ's Gauss integral, summed exactly over the straight pieces of each strand, and every strand pair read in the run's first and third parts landed within 1.5×10^{-8} of a whole number.

The helicity is the count, weighted by flux. Two tubes that carry no helicity of their own, linked m times, carry twice m times the product of their two fluxes. That is the formula Moffatt and Ricca give for two tubes (1992, their equation 1.1, read by the book on 23 September) and Berger and Field give after Moffatt 1969 (their equation 2 and the lines after their equation 15), and the run matched it exactly for all four pairs and all three pairs of fluxes tried (1 and 1, 1 and 2, 0.5 and 3). The pair threaded three times, with fluxes 1 and 2, carries exactly -12 . Every strand of one tube links every strand of the other just as the centrelines do, and strands of the same untwisted tube link each other 0 times. The signs are the book's own: READ section 6's reader gives these pairs -1 , -2 and -3 , and the helicity carries the same sign. "No helicity of their own" is said with care, because "untwisted" means different things: in the run it means a tube whose strands do not link each other, while Berger and Field call a tube untwisted when a field line can be drawn along its top in one flat view. The two agree for the flat Hopf circles and not for the coiled torus links. Berger and Field also warn that as the linking rises the tubes must kink more and more, so in a real pair the tubes' own helicity cannot be ignored. The run sets it to zero by construction, so its answer is the helicity of the link alone. The book's midpoint walk reads the three linked centrelines as -1.00011 , -2.00058 and -3.00160 ; the exact form reads them as -1 , -2 and -3 .

Shape does not enter. Take the pair threaded twice, give one tube one whole twist, and give the tubes fluxes 1 and 2. The run shrank it to half size, enlarged it three times, stretched it 2.5 times one way, bent it with a smooth wobble, and slid one tube a sixth of a turn round the other. The tubes' lengths ran from 9.03 to 54.17, and the closest approach of any two strands from 0.046 to 0.276. The helicity was -7.125 every time. Only a cut can change it.

Each whole twist adds the tube's flux squared: this is the pen. Cut a tube, twist the ends a whole number of turns and join them again: this is Moffatt and Ricca's write head, which the book already credits in WRITE. Every pair of that tube's own strands then links k times, where k is the number of twists, and the helicity moves in whole steps of the flux squared, on top of the link's -2 . With 16 strands per tube, twists of -2 to $+2$ gave -3.875 , -2.9375 , -2 , -1.0625 and -0.125 . The small shortfall from whole steps is the model's own. A bundle of n strands shows only $(1 - 1/n)$ of a twist, because each strand's own twist is left out. Measured, that share was 0.75, 0.875, 0.9375 and 0.96875 for 4, 8, 16 and 32 strands, reaching the whole twist as the strands get finer; Berger and Field note the same, that the strands' own part vanishes as one over their number. Whole steps come from the cut. In a real flux rope the field lines can wind round the axis any number of times, whole or not (Berger and Field: a torus twisted T times carries T times its flux squared, and T "can be any real number"), so a rope's own twist is a measure, not a count. Only the linking of closed tubes is a whole number.

Figure P1. Linked tubes of magnetic flux: the helicity is a whole-number count weighted by flux (run FLUX_TUBE_HELICITY_CHECK)




Figure P1. (A) Two tubes of flux, 8 strands each, threaded twice; the blue tube carries one whole twist. (B) Helicity divided by the two fluxes, against the linking as read: twice the linking, for all three pairs of fluxes. (C) A Hopf pair of fluxes 1 and 1, with whole twists given to one tube: the link alone carries -2 , and each twist adds one flux squared (shown as 0.9375 with 16 strands). (D) The share of one twist that n strands can show: $1 - 1/n$, reaching the whole twist as the strands get finer. Pure mathematics; run FLUX_TUBE_HELICITY_CHECK 20260926_145656.

What the run shows, and what it does not. Plasma's whole number is the linking of its flux tubes, as water's is the linking of its vortex lines. The helicity is that count weighted by flux, so it is not itself a whole number, but every piece of it is. The run is geometry only: it computes no field, no energy and no plasma. It adds twist to a flat ring only; the trade of twist for writhe as a twisted tube coils is the arithmetic of W6 and was not rerun here. Berger and Field write that trade for a flux tube: its own helicity is its twist plus its writhe, times the flux squared (their equations 25 and 26, with White's theorem, their equation 23).

One point the reading suggests, for a specialist to confirm. The energy floor of Freedman and He, which the book read in full on 23 September, bounds the energy of a field whose own lines are linked. In a plasma the linked lines are magnetic field lines, so the floor stands under the magnetic energy itself. In water the linked lines are vortex lines, which is why the flag of 23 September says water's floor may sit on the wrong energy. If this holds, plasma is the medium where the floor stands on the right energy. A plasma physicist should say whether it does.

For the card, still to come. Berger and Field and Taylor's 1986 review (both read in full) give most rows a first sentence. The store is the least magnetic energy linked tubes can relax to, not yet sourced. The wall is reconnection. In a perfectly conducting plasma every linking is kept, because the field lines move with the gas. Reconnection can change the links while keeping the total helicity approximately, which is Taylor's conjecture as Berger and Field report it (Taylor 1974, *Phys. Rev. Lett.* 33, 1139, in their references). Their example: two linked tubes carrying twice the flux squared between them can reconnect into one tube that carries the same amount inside itself, as twist and kink. So in a hot plasma the wall guards the total, not the separate counts, and reconnection moves helicity from the link into the twist. The leak is slow decay through the plasma's resistance; Berger's 1984 limits for the Sun's corona are cited there, not read. The write, on the Sun, is twisting from below: Berger and Field show that a boundary turning in place pours helicity into the space above it (a model on a flat boundary, not an observation of the Sun). Two fences come with it. Helicity has a meaning only inside surfaces the field does not cross, so the Sun's corona can only be read against a reference field, the least-energy field with the same field at the surface, which is READ's reference disk in plasma. And where field lines wander without ever closing, helicity is an average linking, not a count.

What Taylor's review changes: the count is lost, the floor is kept. The book expected plasma to be a latch that leaks slowly. Taylor's review (1986, read in full) supports that only for a quiet plasma that conducts almost perfectly: there, "if one closed field line initially links another n times", the two "must remain linked n times during any subsequent plasma motion". In a turbulent plasma, field lines break and reconnect everywhere, and "flux tubes have no continuous independent existence". The separate whole numbers cannot even be read, "not because the magnetic flux changes significantly, but because it is no longer possible to identify the field line to which the flux belongs". What survives is the total helicity. It is not a whole number, but it still holds energy back: the plasma settles into the least-energy state that total allows, and in the Beta-II experiment the field kept its shape while its energy fell to about an eighth. The guard is not an energy wall above heat's jiggle. It is a race between timescales: relaxation is fast, and resistive decay, which takes the helicity too, is slow. So in a turbulent plasma the count is lost and the floor is kept. That is the book's reading of Taylor. A plasma physicist should confirm it before the card gives its verdict, and the verdict is Bear's to settle with the book's other media. Taylor also gives the pen in the laboratory: a voltage round the

plasma pours helicity in, and a spheromak held that way is the book's powered lock. A lead worth one look, recalled and not checked: S-shaped structures seen on the Sun, called sigmoids, are said to lean one way in the northern hemisphere and the other way in the southern, following the sign of the helicity. If that holds, it is Bear's S rotor on the Sun: a shape that shows a sign. It holds, with a fence (the review of Pevtsov and co-workers 2014, read in full; LATCH_ATLAS_DOMAIN_READING_NOTES_2026-09-26, section 6): S-shaped for positive and reverse-S for negative, obeying the hemispheric rule in 64 % to 87 % of sigmoids depending on the sample, though a view from above can show the same S for writhes of either sign. It is to be checked before it goes anywhere.

Sources: Berger and Field 1984 (its points from the paper) and Taylor 1986 (its points and quotations from the review), both read in full on 26 September; Taylor 1974, cited as both of them cite it; Moffatt and Ricca 1992's two-tube formula and cut-twist-reconnect write head, and the Freedman and He point, as the book's reading notes of 23 September record them, the Freedman and He point carried to plasma here and not yet checked by a specialist; the sigmoids, from memory when first written, checked in the review of Pevtsov and co-workers 2014 (reading notes, section 6). Identifiers: the reference list.

Runs: FLUX_TUBE_HELICITY_CHECK 20260926_145656 (tools/flux_tube_helicity_check.py version 1 with RUN_FLUX_TUBE_HELICITY_CHECK.bat; pure mathematics; seed 20261011, used only for the wobble's phases). Every number in the section is printed in its ledger.

Magnetic skyrmions: the sphere of directions, covered by a whirl

A magnet is made of atoms, each carrying a tiny arrow, its magnetic moment. In some magnets the arrows form small whirls called skyrmions: down at the core, up at the edge, turning round in between. Walk over one whirl and its arrows, between them, point in every direction there is: they cover the whole sphere of directions once. How many times they cover it is the skyrmion's whole number (run SKYRMION_COUNT_CHECK 20260926_192130).

The model. Unit arrows sit on the sites of a square grid, wrapped round at its edges so the grid has no boundary. A whirl of radius R tips its arrows from down at the centre to up at radius R , and turns them round the centre as you go round it: once forwards for a skyrmion, once backwards for an antiskyrmion. A handedness angle sets whether the arrows point outward, like a hedgehog's spines, or swirl round, and the core can point down or up. Pure mathematics.

The read: add up solid angle, triangle by triangle. Cut each square of the grid into two triangles. The three arrows at a triangle's corners mark out a small patch of the sphere of directions, and the patch's signed area is its solid angle. Add up every patch and divide by a whole sphere, 4π . On a grid with no edge that sum is a whole number of spheres for any arrows at all. For skyrmions of radius 3, 6 and 12 grid spacings the read gave -1 exactly, with no distance from a whole number at twelve decimal places. The textbook's smooth formula, applied to the same grid with differences between neighbours, gave -0.669294 , -0.905854 and -0.975672 : near -1 only for large whirls, and never whole. This is READ section 3's integer diagnostic, since the right read of a count returns whole numbers, and it is READ section 2's Berry reading made into a count: the solid angle swept, in whole spheres.

What sets the sign: the turning and the core, never the handedness. The run tried 32 whirls: turnings -2 , -1 , $+1$ and $+2$, handedness 0 , $\pi/2$, π and $3\pi/2$, and the core down and up. Every read was whole, the handedness never changed it, and in every case the count was minus the turning times the core's direction. A skyrmion with its core down reads -1 ; flip the core and it reads $+1$; an antiskyrmion reverses the sign again. So a minus sign can be written two ways, by turning the whirl the other way or by flipping its core. The handedness chooses among whirls with the same count. Bear's cell demo calls its sign the swirl's handedness; in this model that is not what sets it.

The grain: a whirl the size of the atoms slips out. In a smooth plane the count cannot change, and a whirl's exchange energy does not depend on its size, because stretching a pattern in two dimensions leaves that energy unchanged. This whirl's shape costs 19.3322 in units of the exchange J at any radius, above $4\pi = 12.5664$, the least a whirl of count 1 can cost in a smooth plane (Belavin and Polyakov's bound, recalled, to check). On a grid, the run shrank the whirl step by step from radius 6 to 0.2. Its count held at -1 down to about 1.4 grid spacings, then slipped to 0 in one step: at radius 1.40 with the centre on a site or in the middle of a square, and at 1.35 with the centre at a general point. The energy fell nearly all the way down. Just before the slip it was 11.185, 12.304 and 11.130, below the smooth-plane floor (0.890, 0.979 and 0.886 of 4π), and there was no barrier at the slip; the largest energy of the sweep came at or near the starting size (with the centre mid-square, 18.423 at radius 5.70, a ripple above the 18.089 it started at). With the centre on a site, one flipped arrow was left behind after the slip, costing 8.000 and counting nothing.

So in this model, a magnet with exchange alone, there is no wall. The whirl shrinks downhill until it is about the size of the grain, and there the count goes. The book expected a latch whose wall is finite because a magnet is made of atoms. The run shows the atoms are where the count can be lost, and that something besides exchange must hold a real skyrmion's size and put a wall in its way.

Noise makes a wrong count, never a fraction. Tilt every arrow of a radius-6 skyrmion by a random amount of about s radians, 400 times for each s . No read was ever further from a whole number than 1.3×10^{-13} . The reads first went wrong more than 1 time in 100 at $s = 0.6$ (0.0200 of them), and 0.2900 were wrong at 0.8. This is READ section 5's rule in two dimensions. READ section 5's ring of 16 vanes went wrong 1.1 per cent of the time at the same 0.6 radians (run B1_READ_CHECKS 20260923_233341, as READ prints it): a match noted, not explained. The Field Guide met this failure on a hedgehog of degree one, the same sphere count (register row 305): at high noise the whole number was right only 8 per cent of the time, because "white noise manufactures spurious wrappings", so the count is fabricated, not merely lost; smoothing to the texture's own scale before reading restored it to 100 per cent. This run shows the same failure, noise past about 0.6 radians making the read wrong; it recorded only whether each read was wrong, not which wrong values it gave, so whether its wrong reads were fabricated counts or misses is not yet measured.

Figure S1. A whirl of arrows on a grid: its count is the swept solid angle in whole spheres (run SKYRMION_COUNT_CHECK)




Figure S1. (A) A skyrmion on a grid: arrows in the plane, shading for up (blue) and down (red); its count as read, -1 . (B) What sets the sign: the count for each core direction and turning, at four handedness angles; the handedness never changes it. (C) The whirl shrunk on a grid of 32 by 32, with its centre on a site, mid-square and at a general point: the exchange energy (solid), the smooth-plane value for this shape at any size (dashed), the smooth-plane floor 4π (dotted), and the count (dotted steps), which holds at -1 until the whirl is about a grid spacing and then slips to 0. (D) Random tilts on every arrow: the share of 400 reads that come out wrong. Pure mathematics; run SKYRMION_COUNT_CHECK 20260926_192130.

Scattered stations, joined into triangles. Real measurements rarely sit on a neat grid. Bear's Cascadia GPS work (the Cascadia paper, deposited; described here from his project notes) read its scattered stations by joining them into triangles, a Delaunay triangulation with every triangle turned the same way, and the same read works here. Scatter N stations at random over a square 40 grid spacings on a side round a skyrmion of radius 6, add the square's four corners, join them into triangles, and add up the solid angles. The arrows here are clean, with no noise. No sum was ever further from a whole number than 3.8×10^{-15} . With 1600 or 800 stations the read was -1 in every one of 200 layouts, and with 400 stations, a mean spacing of 0.33 of the whirl's radius, it was right in 0.990 of them. Thinner than that, the count is missed, not bent: 0.275 of layouts read 0 at 200 stations, 0.665 at 100, 0.975 at 25, and no layout at any density read anything but -1 or 0. A sparse survey of clean arrows under-counts; it does not invent a count, and it never returns a fraction. Scattered stations need to be denser than an even grid, which held its count until the whirl's radius was about 1.4 grid spacings: random stations leave gaps, and a whirl that falls in a gap is read by a few large triangles whose solid angles cancel (Figure S2, B). One difference from the Cascadia read: there, each triangle read the turning of a flat phase, a whole number for each triangle; here each triangle holds a piece of a sphere, and only the sum over all of them is whole. This measures once a worry Bear raised on 26 September, that the way data are sampled can run against a count: with clean data, coarse sampling costs the count, not its honesty. Noise is different, as the paragraph above and the Field Guide's row 305 show: it can fabricate a count. So the scale of a read has a window. Too coarse and the whirl falls in a gap and is missed; too fine and unsmoothed, noise fakes a count; smoothing to the texture's own scale sits between.

Figure S2. A skyrmion read from scattered stations joined into triangles (run SKYRMION_COUNT_CHECK)




Figure S2. A skyrmion of radius 6 (dashed circle) read from scattered stations joined into triangles, each triangle shaded by its solid angle. (A) 400 stations: many small triangles, each carrying a share of the same sign; the sum is -1 , right. (B) 100 stations: a few large triangles near the core, carrying large shares of opposite sign that cancel; the sum is 0, missed. (C) The share of 200 layouts read right, missed or other, against the stations' mean spacing in whirl radii. Pure mathematics; run SKYRMION_COUNT_CHECK 20260926_192130.

Grouping before reading: by nearness, not by likeness. Bear asked whether arrows that behave alike could be grouped into a voting block, whatever their distance, to suppress noise before the count is read. The run put noise on every arrow of the radius-6 skyrmion and read it five ways, 200 noisy textures per noise level (run GROUPING_READ_CHECK 20260926_195247). All five reads gave -1 for the clean whirl. Read raw, noise mostly fabricates counts rather than missing them: at 1.0 radian 0.445 of reads came out as a whole number other than -1 or 0, 0.205 read 0 and 0.350 were right, and at 2.0 radians 0.825 were fabricated. That settles the question the noise paragraph left open, and it agrees with the Field Guide's row 305.

Smoothing by nearness before reading rescues the count, as row 305's repair did. Averaging each arrow with its neighbours within about one grid spacing kept the read right in at least 99 of 100 textures up to 1.2 radians, 1.000 right at 1.0 radian, and it fabricated nothing below 2.0 radians (0.030 there). Averaging over a wider patch, a third of the whirl's radius, held to 0.8 radians and then began to miss, 0.265 at 2.0 radians: the other side of the window.

Smoothing by likeness, whatever the distance, did the opposite. Averaging each arrow with every arrow on the grid whose 3 by 3 neighbourhood looked like its own erased the whirl. Its core, the few arrows pointing down, was outvoted by the far larger crowd of up-pointing background arrows, and the read missed at every noise level from 0.6 radians; the lowest arrow left after smoothing pointed nearly straight up ($+0.997$ to $+0.999$, where -1 is a full core). A stricter test of likeness, added as a check after the first likeness read erased the whirl and labelled so in the tool, held to 0.6 radians (0.990 right) and then failed the same way at 0.8, once noise made a core patch and a background patch look alike. So in this model a voting block chosen by likeness alone lets the majority outvote the rare feature that carries the count. The count lives in the rare part. Whether likeness can be made to guard a rare feature, for instance by choosing blocks by likeness and then smoothing only by nearness inside each block, is untested, and whether image processing's own likeness averaging meets the same trap is for the prior-art pass.

Figure S3. Grouping before reading: a noisy skyrmion read raw, smoothed by nearness, and smoothed by likeness (run GROUPING_READ_CHECK)




Figure S3. (A) A skyrmion with noise of 1.0 radian on every arrow, read raw: a fabricated count. (B) The same arrows averaged with their near neighbours: read right. (C) The same arrows averaged with every arrow that looks alike, anywhere: the core is washed out and the count is missed. (D) The share of 200 reads right against the noise, for the five reads. (E) At 1.0 radian, each read's results split into right, missed and fabricated. The lowest-arrow measure used in the text is the lowest arrow anywhere on the grid, so for raw noisy arrows it says only that noise tipped some arrow down somewhere. Pure mathematics; run GROUPING_READ_CHECK 20260926_195247.

Where to point a specialist, and what comes next. Three questions go to a magnetician. What holds a real skyrmion's size against the downhill slide the run shows? The book expects the twist the crystal itself favours and an applied field, to check. How high is the barrier to collapse in real materials? That is the medium's wall, and its price. And what does a microscope tip do when it writes or deletes one skyrmion Romming and co-workers 2013 answer it for one film: the tip's electrons flip a pinned whirl by their energy, not by heat, and the current's direction chooses the way. Bear's own deposited work on this medium, Paper A (concept DOI 10.5281/zenodo.20587337), the first sighting of a skyrmion lattice (Mühlbauer and co-workers 2009) and the first single skyrmions written and deleted (Romming and co-workers 2013) are read in full (LATCH_ATLAS_DOMAIN_READING_NOTES_2026-09-26, sections 7 to 9); the card comes next.

Sources: Bear's Paper A, Mühlbauer and co-workers 2009 and Romming and co-workers 2013, read in full (reading notes, sections 7 to 9); Belavin and Polyakov 1975, recalled, not read; the Field Guide's row 305; the Cascadia paper's triangle read, described from Bear's project notes. Identifiers: the reference list.

Runs: SKYRMION_COUNT_CHECK 20260926_192130 (tools/skyrmion_count_check.py version 2 with RUN_SKYRMION_COUNT_CHECK.bat; pure mathematics; seed 20261012 for Part D's noise and 20261013 for Part E's stations; Parts A to D and Figure S1 identical to version 1's run of the same day, which it replaces), which prints every skyrmion number in the section except the grouping paragraphs', printed by GROUPING_READ_CHECK 20260926_195247 (tools/grouping_read_check.py version 1 with RUN_GROUPING_READ_CHECK.bat; pure mathematics; seed 20261014); READ's B1_READ_CHECKS 20260923_233341 for the 1.1 per cent.

The Sun, read from a real magnetogram

The sections above read models. This one reads the Sun itself (run SUN_WINDING_READ 20260926_194226). Stanford's Joint Science Operations Center publishes maps of the magnetic field over each of the Sun's active regions, every twelve minutes since May 2010. The maps come from the Helioseismic and Magnetic Imager on NASA's Solar Dynamics Observatory (the SHARP series; Bobra and co-workers, arXiv 1404.1879, 2014). Bear downloaded one: active region 11166 at 23:24 on 9 March 2011, the region and time of the SHARP paper's own Figure 2. The map gives the field in three parts at each of 280,006 points (the count each file's header gives, DATAVALS), each with its error. Two more maps say which points belong to the region, and where the pipeline is sure of the field's direction. The data are marked CC0, and NASA asks that use carry "Courtesy of NASA/SDO and the AIA, EVE, and HMI science teams."

What is walked. The part of the field that lies along the surface, the sideways field, is an arrow at every point, like the book's weather vanes. READ section 1's walk applies unchanged: go round a closed loop, add up how far the arrow turns from each point to the next, and the total is a whole number of turns. The run walks square loops 8, 16, 32 and 64 pixels on a side (2.9 to 23.3 Mm; a pixel is 0.364 Mm). It reads a loop only if every point on its rim passes three tests: it is inside the region, the pipeline is confident of it, and its sideways field is at least three times its error. Of the region's 156,196 points, 44,984 pass. The inside of a loop may hold anything, because the count is read from the rim alone. No 64-pixel loop found a rim that passes all the way round.

Two checks came first. The pipeline builds its own summary numbers from the region’s confident points only, and records how many there are in the keyword CMASK. Counted from the maps: 51,073, equal to CMASK. Then the arrows’ directions. Stanford’s table gives one part of the field as westward and another as southward, and the file’s own coordinates show the map’s x growing westward. So the sideways arrow on the map is the westward part across and minus the southward part up. Part D of the run shows why this matters: leave that one sign unflipped, and every loop reads minus its true count.

What the loops read. Every total landed within 5×10^{-15} of a whole number. Most loops read 0; of the 221 readable 32-pixel loops, 85 read +1 and 7 read –1. The nonzero loops enclose points where the sideways field vanishes; the ledger groups them by the weakest point inside each loop, into nine groups at eight places (one place is read +1 by some loops and –1 by one, because different loops take in different neighbours).

- Five of those groups sit at sunspot centres, where the field stands straight up (at least 1,000 gauss), and all five read +1. Physics expects exactly that: a spot’s sideways field fans straight out from a positive spot and straight in to a negative one, and a fan either way is +1 turn.
- The run checks the fan a second way, by the field’s flow across each loop’s rim: outward round every positive spot and inward round every negative one, 5 of 5.
- The –1 counts are saddles, crossings where arrows come in from two sides and leave by the other two, and they sit between spots.




Figure H1. Active region 11166 at 23:24 TAI on 9 March 2011. Grey: the upward field (white up, black down, to $\pm 2,500$ gauss). Tint: the points a loop’s rim may use. Arrows: the sideways field. Marks: the zeros the loops count, red circles +1 and blue squares –1. Data: SDO/HMI, SHARP series hmi.sharp_cea_720s, courtesy of NASA/SDO and the AIA, EVE, and HMI science teams. Run SUN_WINDING_READ 20260926_194226.

Noise. Fifty copies of the map were made, each point’s field moved by a random draw of its own error, and the same loops read again. Round the three big spots standing on their own, no loop’s count changed in any copy: 48 of 48 loops, 27 of 27 and 17 of 17. The loops at the region’s crowded centre did change. The saddle there kept its count in only 3 of its 13 loops.

Where the fragile counts sit. The crowded centre is where the region flared. Vemareddy and Wiegmann, in “Quasi-Static 3D-Magnetic Field Evolution in Solar Active Region NOAA 11166 Associated with X1.5 Flare” (*Astrophysical Journal* 792, 40, 2014, DOI 10.1088/0004-637X/792/1/40, as its arXiv record 1406.7823 gives it; abstract read), report an X1.5 flare at 23:13 UT that day, about ten minutes before this map. It sat “over a positive

flux region surrounded by negative flux at the center of the bipolar AR”, with “a single magnetic null point associated with a fan-spine topology” in the corona above, and “the footpoints of the fan separatrix surface agree with the inner edge of the quasi-elliptical ribbon”. The run’s only saddles, and its least stable counts, sit in the same place: a positive patch (881 gauss upward at its zero) among negative spots. Whether a saddle in the sideways field at the surface marks the foot of that null above is a question for a solar physicist. The book does not claim it.

Who came first. Counting a field’s zeros with a topological index round closed loops or surfaces is standard. Zhao, Wang, Zhang and Xiao find null points in three-dimensional models of the corona that way (IAU Symposium 233, 369–372, 2006, abstract read), and Longcope’s review sets out the topology of the Sun’s field: null points, separatrices and bald patches (Living Reviews in Solar Physics 2, 7, 2005, abstract read). The Sun’s twist has also been read on these same maps: Xiao, Prior and Yeates compute winding-based helicity and winding density for SHARP magnetograms, from single maps, with no velocities (Solar Physics 298, 116, 2023, read in full). So this run is a calibration: the book’s walk, run on the real Sun, gives the answers physics expects. The book credits the people who read the Sun’s topology and winding before it.

What the read does not do. It counts the zeros of the sideways field, which is a label: read-only, holding nothing back. How far a spot’s fan is turned from straight out, its twist, is a measure, not a count. The linking the plasma card is about lives in the field above the surface. Xiao, Prior and Yeates (read in full) show how far a single map goes: it gives how fast field lines twist round each other at the surface, a measure, and the whole winding needs the field above too. Their paper also shows the sign trap in print: one sentence gives the SHARP directions turned half a turn from the book’s, though their results show their computation was consistent (prior-art ledger, S1).

Sources: the SHARP maps (Bobra and co-workers 2014 for the series; Stanford’s table for the field’s parts), the record for HARP 401 at 2011.03.09_23:24:00_TAI; Vemareddy and Wiegmann 2014, Zhao, Wang, Zhang and Xiao 2006 and Longcope 2005, read at abstract level through a page-summary tool on 26 September, evening, the flare paper’s abstract from its arXiv record; Xiao, Prior and Yeates 2023, read in full (sittings 123 and 124). Identifiers: the reference list.

Runs: SUN_WINDING_READ 20260926_194226 (tools/sun_winding_read.py version 1 with RUN_SUN_WINDING_READ.bat; real data: eight files whose sha256 heads the ledger lists; seed 20261016 for the noise copies). Every number in the section is printed in its ledger except 280,006, from the files’ headers. Earlier runs of the same evening were redone after changes to labels and wording before anything shipped, and are not kept.

The Sun through a flare

One map is a snapshot. Bear then downloaded fifteen: the same region every twelve minutes, from 22:00 on 9 March 2011 to 00:48 the next morning, with the X1.5 flare at 23:13 UT in the middle (run SUN_SERIES_READ 20260927_041443). The map times are in TAI, which in 2011 ran 34 seconds ahead of UT (astropy’s tables). Each map is a weighted average of data gathered over 1,215 seconds, about twenty minutes (as Stanford’s page on the instrument’s field data gives it), so the 23:12 map straddles the flare’s start and the 23:24 map follows it. The walk is the one above, unchanged: square loops of 8, 16 and 32 pixels, stepped 4 pixels at a time, read only where the whole rim passes the three tests. One thing is new: a fixed window round the flare site, the crowded centre where the first read’s fragile counts sat, pixels 370 to 480 across and 120 to 185 up, fixed from that first read before the series was read.

Checks first. The 23:24 map, read again inside the series, gives exactly the first read’s numbers (1,464, 844 and 221 readable loops, with the same counts). Across all fifteen maps, every loop’s total landed within 7.1×10^{-15} of a whole number. The run found four walls.

Wall one: a strong rim holds its count. 1,589 loops kept a readable rim in all fifteen maps (1,023 of side 8, 489 of side 16 and 77 of side 32). Not one changed its count: not before the flare, not during it, not after. Fifty-two of them read +1 the whole time, and the rest read 0 the whole time. This is the book’s rule met on the real Sun, through a large flare: a count can change only where its rim goes weak.

Wall two: at the flare site, no rim holds. Of those 1,589 loops, 186 are centred inside the site, and all 186 read 0. So no loop round any zero in the site kept a readable rim for the three hours. There the zeros come and go. Following each zero from map to map (it is taken to be the same zero if one of the same count sits within 8 pixels in the next map), they make 29 tracks, and 17 of those are seen in one map only. The site’s net count, its +1 zeros minus its –1 zeros, wanders between +1 and +5 (Figure H2, A). None of this breaks wall one: every change happens where a rim is weak or unreadable. The flare site’s counts are read in the gaps between strong rims, not protected by them. The

noise test agrees. Fifty copies of each map were made, as before, and the share of the site's loops whose count never changed ran from 0.934 to 1.000. The lowest share is the 23:24 map's, the map gathered in the thick of the flare, whose X-ray output ran from 23:13 to 23:29 UT.

Three zeros last the whole series. Two are the +1 zeros of the two strong negative spots inside the site, near pixels (391, 155) and (459, 155). At both, the upward field points down with more than 1,900 gauss in every map. The third is the -1 saddle near (405, 166). The upward field at that saddle stays between -108 and +2 gauss in every map, against about 2,000 gauss in the spots. So it sits on the polarity inversion line, the line where the upward field changes sign. At that one place, the sideways field vanishes right where the upward field changes sign, and it stays there for three hours, through the flare.

Wall three: the instrument, during the flare. The positive patch at the flare site has a +1 zero of its own, followed for 12 maps. The upward field at that zero reads 1,743 gauss at 22:48, then 1,348, 1,184 and 881 in the next three maps, and 1,686 at 23:36. The fall starts in the 23:00 map, one map before the one that straddles the flare's start. The zero also shifts by a few pixels from map to map (from (397, 141) at 23:00 to (394, 143) at 23:24 and (398, 142) at 23:36), over a patch whose field changes steeply. So the run cannot say how much of the dip is the Sun and how much is the zero moving. There is a third possibility too. The field is worked out from the light the instrument sees, and during a flare the light itself changes. And a fourth, found later ("Are the zeros real?", below): in the 22:00, 22:36 and 22:48 maps, the loops round this patch counted +1 with no weak field anywhere inside (625, 672 and 570 gauss at the weakest point, each more than three times its error), so whatever made the count there then, it was not a place where the field vanishes. The upward field "at the zero" therefore compares different kinds of points, starting with the 22:48 value, and the dip means less than it seems.

Wall four: an average can hide a change. Over the site's confident pixels, the mean sideways field is 776.5 gauss at 22:00 and 762.4 at 23:12, and it stays between 761.4 and 767.9 afterwards (Figure H2, B). That is a fall of 1.8 per cent, most of it before the flare. The upward flux through the site stays between 6.770 and 6.854×10^{21} maxwell. Read this way, the flare hardly shows. Yet for this very flare, Lu and co-workers measured the sideways field rising by 258 ± 73 gauss along the polarity inversion line (below). An average over the whole window, spots and all, cannot see a change confined to a few narrow strips.




Figure H2. Active region 11166 through the X1.5 flare of 9 March 2011: fifteen SHARP maps, twelve minutes apart. (A) Zeros of the sideways field inside the flare site, +1 (red) and -1 (blue); dotted, the share of the site's loops whose count survived all fifty noise copies, times 10. (B) The mean sideways field inside the site. The axis spans only 761 to 777 gauss, so the whole change is about 2 per cent. (C) and (D) The site before (22:48) and after (23:36) the flare. Grey: the upward field. Arrows: the sideways field where readable. Red circles +1 and blue squares -1: the zeros. Dashed box: the site. Times are T_{REC} , in TAI; the dashed line marks the flare's start, 23:13 UT. Data: SDO/HMI, SHARP series hmi.sharp_cea_720s, courtesy of NASA/SDO and the AIA, EVE, and HMI science teams. Run SUN_SERIES_READ 20260927_041443.

Who met these walls before, and how they got in. People have been reading these same maps for what a flare does to the surface since 1911, and each of the book's walls has been met before, from another side.

- *The average that hides a change (wall four).* Wang Shuo and co-workers found the sideways field at the polarity inversion line of region 11158 rising by 70 per cent in an X2.2 flare (arXiv 1103.0027, 2011). Liu Chang and co-workers found a 28 per cent rise at the centre of a sigmoid in an M6.6 flare in the same region (*Astrophysical Journal Letters* 745, L4, 2012). Lu and co-workers surveyed 20 X-class flares (*Astrophysical Journal*, 2019, DOI 10.3847/1538-4357/ab16d4; read in part, its introduction, method and the table rows for region 11166). Their way in is the one the book's average lacks: a difference map of the sideways field, from the flare's start to its end, keeping only the places where it rose by more than 100 gauss. They then make the same difference three hours later, to throw out changes the flare did not cause. For this region's X1.5 flare they list a rise of 258 ± 73 gauss and a change of inclination of 2.31 ± 1.13 degrees. They call it the flare's "magnetic imprint".
- *The instrument during the flare (wall three).* Lu and co-workers write that "inversions of Stokes profiles are generally not reliable during the flare", citing Hong and co-workers 2018. They get round it by comparing the maps before and after, where the inversions are reliable. Martínez Oliveros, Lindsey, Hudson and Buitrago Casas met the same wall in the instrument's velocity signal. A blueshift at a flare's footpoints proved to be an artifact once checked against a second instrument, GONG++ (*Solar Physics* 289, 809, 2014). Ambastha and Maurya review how flares change the spectral lines the field is read from (arXiv 1203.3326, 2012).

- *A surface label that changes at a flare (wall two)*. Lee, Sun and Kazachenko read a topological feature off these maps too: the bald patch, a stretch of the polarity inversion line where U-shaped field lines “turn tangent to the photosphere”. In the X9.3 flare of region 12673, the central bald patch “largely ‘disintegrated’ within 35 minutes”: the part of the sideways field across the inversion line changed sign, so the field now crossed the line the other way (*Astrophysical Journal Letters* 921, L23, 2021). That label changed at the polarity inversion line, where the book finds no rim that holds. The book’s persistent saddle sits on that line as well.
- *The measure beside the count*. Raphaldini, Prior and MacTaggart follow the flux of magnetic winding into five regions over time. For that they need the surface velocities, worked out from the maps. In the flaring regions, “sharp inputs” of winding and helicity came “several hours” before individual flares (*Astrophysical Journal* 927, 156, 2022). Their open tool, ARTop, downloads SHARP maps and makes these time series (Alielden, MacTaggart, Ming, Prior and Raphaldini, *RAS Techniques and Instruments* 2, 398–407, 2023). The same group then built these time series for 144 regions and found extreme values often 1 to 8 hours before a flare (Williams, Prior and MacTaggart, arXiv 2412.04335, 2024). They then trained a forecaster on them, together with each region’s flare history; flare history and the accumulated current-carrying winding and helicity were its most important inputs (Williams, Prior, MacTaggart and Bloomfield, arXiv 2512.14840, 2025). Sorriso-Valvo and co-workers found that the scaling of current helicity inside flaring regions changes abruptly at big flares (arXiv astro-ph/0207244, 2002).

What the flare series teaches. A count held by a strong rim does not move, even through an X-class flare. The flare happens where no rim holds. The people who study flares at the surface read measures there: difference maps, winding fluxes, the sign of the field across the inversion line. And the measures move, sometimes hours ahead. The book’s count does something else. It maps where the surface’s topology is locked, and so, by what is left over, where a flare is free to change it. Whether that map is useful beside the measures is a question for a solar physicist. So are two more. Is a zero of the sideways field that sits on the inversion line for three hours the foot of the null Vemareddy and Wiegmann find above, the end of a bald patch, or neither? And is the dip at the positive patch the inversion trouble Hong and co-workers describe? The next run makes Lu and co-workers’ difference map on these same maps, and asks whether the book’s zeros and the saddle sit in the imprint’s strips. The book does not claim any of this is new; the prior-art ledger (S5) keeps what was read and at what depth.

Sources: the fifteen SHARP records for HARP 401 from 2011.03.09_22:00:00_TAI to 2011.03.10_00:48:00_TAI; Stanford’s page on the instrument’s field data for the averaging window, and astropy for TAI minus UT; the flare papers, found by web search on 26 September, late evening (Pacific): Lu and co-workers 2019 read in part from its arXiv text; Wang Shuo and co-workers 2011, Liu Chang and co-workers 2012, Ambastha and Maurya 2012, Sorriso-Valvo and co-workers, Martínez Oliveros and co-workers 2014 and Lee, Sun and Kazachenko 2021 at abstract level from their arXiv records, and Williams and co-workers’ two forecasting papers the same way; Raphaldini, Prior and MacTaggart 2022’s abstract from the University of Glasgow’s record; ARTop from its README; Hong and co-workers 2018 as Lu and co-workers cite it, not read; journal details as those records, Lu and co-workers’ reference list, search listings or the DOIs give them; the prior-art ledger, S5. Identifiers: the reference list.

Runs: SUN_SERIES_READ 20260927_041443 (tools/sun_series_read.py version 1 with RUN_SUN_SERIES_READ.bat; real data: 120 files whose combined sha256 heads the ledger; seed 20261017 plus the map’s index for the noise copies). Every number in the section is printed in its ledger except the first read’s loop numbers (the SUN_WINDING_READ ledger), the 34 seconds, the 1,215 seconds, and the bounds (more than 1,900 gauss, about 2,000 gauss, 1.8 and about 2 per cent), worked by script from the ledger’s numbers. An earlier run the same day printed an identical ledger and was redone only to move Figure H2’s legend off the data, before anything shipped; it is not kept.

The flare’s imprint, read three ways

The flare series showed that an average over the whole site cannot see the flare. The next run reads the flare’s imprint the way Lu and co-workers do, on purpose, so that their published method can be run on the book’s own maps as a check of the tool, and then two more ways that came from Bear (run SUN_FLARE_DIFFERENCE 20260927_045554). It uses three pairs of maps, each pair 24 minutes apart. FLARE is 23:12 to 23:36, the maps either side of the flare; the 23:24 map, gathered in the thick of it, is left out. BEFORE is 22:24 to 22:48. AFTER is 00:24 to 00:48, about an hour after the flare ended. A difference map is the later map’s sideways field minus the earlier one’s, pixel by pixel, over the pixels inside the region and confident in both maps.

Lu and co-workers' way. Keep the pixels whose sideways field rose by more than 100 gauss. Inside the flare site the FLARE pair has 1,258 of them, against 576 and 650 in the two quiet pairs. Their region of interest, made automatic here, is the site pixels that rose by more than 100 gauss across the flare and did not rise that much across AFTER: 1,102 pixels, 146 square megametres. Its mean sideways field, weighted by field strength as theirs is, went from 854.9 gauss at 23:12 to 1,072.6 at 23:36, a rise of 217.7 gauss. Lu and co-workers list 258 ± 73 gauss for this flare.

A wall: pick pixels because they rose, and they rose. The same recipe run on a quiet pair (AFTER as the event, BEFORE as its check) finds 531 pixels and a "rise" of 153.1 gauss, where nothing flared. So a rise between two maps, measured over pixels chosen from those same two maps, mostly measures the floor that chose them. The way past the wall is time: follow the chosen pixels through maps the choice never used (Figure H3, D). The flare region sits flat before the flare, between 829.6 and 854.9 gauss from 22:00 to 23:12. It climbs through 987.3 at 23:24 to 1,072.6 at 23:36, and it stays up, between 1,064.8 and 1,083.1 from 23:36 to 00:36. Its mean over the six maps from 22:00 to 23:00 is 840.2 gauss, and over the three maps from 23:48 to 00:12 it is 1,079.9: a step of 239.6 gauss, measured only in maps the selection never looked at. The quiet region does the opposite. It drifts down for two hours, from 801.1 gauss at 22:00 to 702.7 at 00:00, and then climbs, most steeply across the pair that picked it (742.1 at 00:24 to 895.1 at 00:48). Between the same two windows its step is -67.3 gauss. So the flare's rise is real and it lasts, which is what Lu and co-workers mean by an imprint, "rapid and irreversible". The book's 239.6 gauss and their 258 ± 73 are close, though the two regions of interest are chosen differently.

Three hours later. Bear then pulled two maps from three hours after the flare, 02:12 and 02:24 TAI on 10 March (run SUN_FLARE_DIFFERENCE 20260927_053821: the same tool, unchanged, with those two maps added; every number of the run of record came out the same, and Part E gained two rows). The flare region is still up: its mean sideways field reads 993.2 and 986.6 gauss, 149.7 above its level before the flare. But part of the step has faded: the region kept 0.62 of its 239.6-gauss step, and the noise zones' region 0.68 of its 354.5. The quiet regions went the other way. Picked for rising between 00:24 and 00:48, they kept climbing, to 974.0 and 966.2 gauss, 191.8 above their early level. So slow changes as big as the flare's happen over a few hours with no flare at all. What marks the flare's change is its suddenness: flat for an hour before, a step inside 24 minutes, flat for the hour after. That is why Lu and co-workers check against the same change three hours later, and why others fit a step to each pixel's history.

Bear's noise zones. Bear's picture was to break the substrate into zones of common noise. The site's pixels were sorted by the pipeline's own error of the difference into four zones of equal size, and each zone's quiet scatter was measured on the BEFORE pair. The zone's floor is three times that scatter. The floors came out at 140.4, 182.9, 192.5 and 261.8 gauss, so in every zone the maps' quiet scatter is larger than Lu and co-workers' flat 100 gauss allows for. Where the pipeline says the field is best known (zone 1, errors of the difference from 16.2 to 44.0 gauss), the quiet scatter is 46.8 gauss, larger than the stated errors: there the maps change more from one to the next than their error bars say. In the noisiest zone it is the other way round (a scatter of 87.3 gauss, against errors from 87.9 up). A floor set by each zone's quiet scatter follows what the maps actually do. With the zone floors, the FLARE pair passes 382 site pixels as rising and the quiet AFTER pair 115, a ratio of 3.3. The flat floor passes 1,155 and 597, a ratio of 1.9, and each pixel's own error, three times over, passes 361 and 148, a ratio of 2.4 (Figure H3, C). The zones' flare region, 402 pixels, steps up by 354.5 gauss between the same two windows and stays up; its quiet twin, 113 pixels, steps by -191.4 . So Bear's zones give a smaller, purer imprint: fewer pixels, a bigger step, and fewer false alarms. (The zone floors were set on the BEFORE pair, so BEFORE is not a fair test of them; AFTER is.)

Where the book's zeros sit. Of the six zeros in the site at 23:12, two sit on the imprint: the -1 at (393, 144) and the positive patch's $+1$ at (397, 140). Five of the six are within one pixel of it. For a site pixel picked at random those shares are 16 and 30 per cent. At 23:36 none sits on it, and three of five are within a pixel. The persistent saddle is 2.2 pixels (0.81 Mm) from the imprint at 23:12 and 1 pixel at 23:36. Six zeros are too few to say more than this: they sit at the imprint's edge more often than chance would put them there.

Bear's rim. Bear asked whether it helps to black out a loop's inside and read only its rim. From the rim alone come three things. One is the count, as before. Another is the circulation of the sideways field round the rim, which by Ampère's law is the upward electric current through the loop: the current equals the circulation divided by the magnetic constant. The third is the flow of the sideways field out across the rim. Nothing inside is needed. The check adds up the circulation round every small cell inside each loop and matches the rim's to within 5.8×10^{-11} , where the largest circulation is 4.7×10^4 in the same units: what crosses the inside shows up on the rim, exactly (Stokes' theorem). Across the FLARE pair, 722 loops of side 16 had rims readable in both maps, and none changed its count: wall one again. Near the flare site the currents through these loops ran from -4.29×10^{11} to 1.36×10^{12} amperes at 23:12. Across the flare their typical change, 1.33×10^{10} amperes, was no bigger than across the quiet pairs (1.32 and

1.37×10^{10}). But in the site, 10 of 80 loops changed their current by more than 8.68×10^{10} amperes, the quiet pairs' 99th percentile, against 4 of 103 before and 1 of 89 after (Figure H3, E). They cluster round the positive patch. The loop holding the patch's +1 zero (x 396 to 412, y 132 to 148) carried 7.04×10^{11} amperes at 23:12 and 1.07×10^{12} at 23:36. So there is an advantage. The rim gives the current through the loop exactly, whatever the inside holds, and it shows when a flare has changed the current somewhere inside. What it cannot show is where inside, or whether currents of both signs hide in there. Like the count, it gives the net.

HARP 401 (NOAA 11166), X1.5 flare of 9 March 2011: the imprint read three ways. Data: SDO/HMI SHARP; courtesy of NASA/SDO and the AIA, EVE, and HMI science teams




Figure H3. The X1.5 flare's imprint on active region 11166, 9 March 2011. (A) The change in the sideways field from 23:12 to 23:36 TAI, across the flare: red rises, blue falls; red contour, a rise of more than 100 gauss; black line, where the upward field is zero; circles, the book's zeros at 23:12 and squares at 23:36 (green +1, purple -1); dashed box, the site. (B) The same for the quiet AFTER pair, 00:24 to 00:48. (C) Site pixels passed as rising by three floors: Lu and co-workers' flat 100 gauss, Bear's noise zones, and three times each pixel's own error; red the FLARE pair, grey the quiet AFTER pair. (D) The field-strength-weighted mean sideways field of four regions of interest through all fifteen maps: solid, chosen by their rise across the flare; dotted, chosen the same way on the quiet AFTER pair; red, flat floor; blue, noise zones. Dashed line: the flare's start, 23:13 UT. (E) Bear's rim: the share of loops of side 16 whose current, read from the rim alone, changed by more than the quiet pairs' 99th percentile, for loops centred in the site (red) and elsewhere (grey). Data: SDO/HMI, SHARP series hmi.sharp_cea_720s, courtesy of NASA/SDO and the AIA, EVE, and HMI science teams. Run SUN_FLARE_DIFFERENCE 20260927_045554.

What it teaches. The imprint is real, it lasts, and it is close to Lu and co-workers' size. A rise between two maps is not evidence on its own, because choosing the pixels makes the rise; the step that stays up in maps nobody chose from is. Bear's noise zones sharpen the imprint. And the rim, read alone, sees a flare's change in current without looking inside. Who met these walls before, and how they got past them, is the next search.

Sources: Lu and co-workers 2019, whose method is run here on purpose; the six SHARP records the tool names, and all fifteen of the series for Part E; Bear's two post-flare maps of 02:12 and 02:24 TAI. Identifiers: the reference list.

Runs: SUN_FLARE_DIFFERENCE 20260927_045554 (tools/sun_flare_difference.py version 1 with RUN_SUN_FLARE_DIFFERENCE.bat; real data: 48 files whose combined sha256 heads the ledger; no random numbers), which prints every number of the section except the ratios 3.3, 1.9 and 2.4, worked by script from its counts; SUN_FLARE_DIFFERENCE 20260927_053821 (the same tool with the two post-flare maps added) for the three-hours-later paragraph, except 149.7, 191.8, 0.62 and 0.68, worked by script from its rows. Three earlier runs, minutes before, were redone after the selection check, the distance baseline, the through-time check and a clearer last panel were added, before anything shipped; they are not kept.

Along the polarity line: across or along

Bear's "near 90 degree off axis" points at the polarity inversion line, the line where the upward field changes sign. There the sideways field can do one of three things. It can cross the line straight, from the positive side to the negative, the way a relaxed field would: call that angle 0. It can run along the line, 90 degrees off: the field solar physicists call sheared. Or it can cross the wrong way, from the negative side to the positive, past 90 degrees: the direction of a bald patch, where field lines dip down to touch the surface. The run (SUN_PIL_SHEAR 20260927_051041) finds the line in each of the fifteen maps inside the flare site. A pixel is on the line where the upward field, lightly smoothed (2 pixels), changes sign. The field must change by at least 50 gauss per megametre there, and the sideways field must be readable. That gives 203 to 242 line pixels per map.

The line was already sheared, and the flare sheared it a little more. In the six maps from 22:00 to 23:00 the mean angle ran from 68.6 to 74.7 degrees, with 80 to 84 per cent of the line past 45 degrees. In the seven maps from 23:36 to 00:48 (leaving out the two that straddle the flare) the mean angle was 73.5 degrees against 70.8 before. The share past 45 degrees was 0.88 against 0.82, and the share within 10 degrees of running along the line 0.38 against 0.31. The field along the line rose by 33.4 gauss, to 1,027.0, and the field across it fell by 42.0, to 240.5 (Figure H4). Every change points the same way, toward more shear. That is the way Wang 1992 and Wang and co-workers 1994 first saw flares change the surface: "A strong and permanent increase of the magnetic shear at the polarity inversion line", as Lu and co-workers retell it. But here the changes are small beside the scatter from map to map, and the before and after ranges overlap. The field along the line had been rising for an hour before the flare (945.1 gauss at 22:00, 1,025.2 at 23:00), and the mean angle kept rising for the hour after (75.2 degrees at 00:48). Read over the whole line, the flare's change is diluted again, the same way as the site average in a new form. The share crossing the wrong way did not move (0.209 before, 0.216 after): no bald patch disintegrated along this line that the run can see.

Where the saddle sits on the line. The persistent saddle is within 1 to 5.1 pixels of the line in every map. Summed over all the maps, 67 of the 128 line pixels within 6 pixels of it cross the wrong way: 52 per cent, against 22 per cent for the line as a whole. That is unusual but not rare. Taking every line pixel of every map in turn as a centre, the wrong-way share within 6 pixels averages 0.221, and one centre in five (0.203) has a share of a half or more. So the saddle tends to sit where the crossing flips between the right way and the wrong way. Where the crossing reverses, the part of the field across the line passes through zero, and if the part along it is weak there too, the whole sideways field vanishes: a zero, on the line. That makes the specialist's question sharper: is the saddle the end of a bald patch? In the map gathered in the thick of the flare (23:24), all 14 line pixels near the saddle crossed the wrong way. That is also the map the instrument is least sure of, so the book notes it and does not lean on it.




Figure H4. The sideways field on the polarity inversion line of active region 11166 through the X1.5 flare of 9 March 2011. (A) and (B) The line at 22:48 and 23:36 TAI over the upward field (grey), each pixel coloured by the angle between the sideways field and the straight crossing: blue crosses the right way (0), white runs along the line (90), red crosses the wrong way (180). Square: the persistent saddle. (C) The mean angle along the line in all fifteen maps. (D) The mean field along the line (red) and across it (grey, right way positive). (E) The share of line pixels past 45 degrees (blue) and past 90 degrees (purple). Dashed line: the flare's start, 23:13 UT. Data: SDO/HMI, SHARP series hmi.sharp_cea_720s, courtesy of NASA/SDO and the AIA, EVE, and HMI science teams. Run SUN_PIL_SHEAR 20260927_051041.

Who met these walls before. All of them, each from another side (prior-art ledger, S6). The picking wall is passed by reading each pixel's history rather than one difference: step-like, permanent changes in maps taken every 90 or 135 seconds (Sun and co-workers 2017), and pixels grouped by the shape of that history with machine learning (Alanazi and Fletcher 2026), a cousin of Bear's noise zones. Bear's rim is Ampère's law round a contour, which Venkatakrishnan and Tiwari (2009) drew as circles round sunspots to show that a spot carries no net current; currents in flare ribbons have been seen to double during a flare (Janvier and co-workers 2014). Along the line, Petrie (2012) found in six flares, this region's among them, that the field parallel to the line strengthened, as the book's run shows, and that the field's tilt from vertical changed most clearly of all, which the book has not read. The book's "wrong way" across the line is the standard bald-patch test (Seehafer 1986; Titov and co-workers 1993). Not found in what was read: floors set per zone of common noise on a quiet pair, and rim currents used to flag where a flare changed the current. Those are doors to show a specialist, not claims.

Sources: Lu and co-workers 2019, retelling Wang 1992 and Wang and co-workers 1994; the papers in "Who met these walls before", found by web search on 26 September, late evening (Pacific), and read at abstract level except where the prior-art ledger (S6) says otherwise: Sun and co-workers 2017, Alanazi and Fletcher 2026, Venkatakrishnan and Tiwari 2009, Janvier and co-workers 2014, Petrie 2012, Seehafer 1986 and Titov and co-workers 1993. Identifiers: the reference list.

Runs: SUN_PIL_SHEAR 20260927_051041 (tools/sun_pil_shear.py version 1 with RUN_SUN_PIL_SHEAR.bat; real data: the same fifteen SHARP records as the flare series, whose 120 files hash the same, 6cc858e1ec8c26b0; no random numbers). Every number in the section is printed in its ledger; 52 per cent is the ledger's 0.523. Two earlier runs, minutes before, were redone before anything shipped, one to add the baseline round every line pixel and one to replace a colour scale that showed 0 and 180 degrees alike; they are not kept.

A second flare: region 11158

Everything above rests on one region and one flare. The second is the X2.2 flare of 15 February 2011 in region 11158 (HARP 377), the archetype many of the papers above study, from 01:44 to 02:06 UT by Lu and co-workers' table. Bear fetched fifteen maps, every 12 minutes from 00:36 to 03:24 TAI, with the book's own download script (tools/sun_fetch_sharp.py): one double-click in place of 120 files by hand. All fifteen records were complete, region 11158 and quality 0, and each one's count of confident pixels matched its CMASK.

The same tool read them (run SUN_SERIES_READ 20260927_060906; sun_series_read.py version 2, which reads any listed region and repeats region 401's run of record number for number). This time the site was chosen by a rule, on the first map only, before the flare: the strong polarity line (gradient at least 300 gauss per megametre), its pieces of at least 20 pixels, boxed and widened by 20 pixels. That gives x 321 to 460 and y 200 to 282, the line between the region's two central spots.

Wall one holds again. 1,352 loops kept a readable rim in all fifteen maps, and not one changed its count. Every loop total landed within 7.2×10^{-15} of a whole number.

But here, rims inside the site held. In region 11166 no loop centred in the site kept a readable rim for the three hours. In region 11158, 486 did, and 15 of them read +1 throughout: they circle the two big spots either side of the line, whose zeros both last all fifteen maps. The zeros that come and go sit on the line between the spots. In the 02:12 map, the first after the flare's end, the loop walk finds seven zeros in the site where there were three, four of them new, at and beside the polarity line (x 387 to 403, y 222 to 235). The check below ("Are the zeros real?") shows that three of the four are not zeros: there is no weak field anywhere inside their loops (426 to 536 gauss at the weakest point), only a jump of 173 degrees between neighbouring arrows. The fourth, a saddle where the upward field is 36 gauss, is a real zero. So the spots' counts are locked, and what changes between them at the flare includes the instrument's own reading of which way the arrows point.

The site average sees this flare. Because the rule centred the site on the flaring line, its mean sideways field steps up at the flare, from between 765.6 and 772.9 gauss in the six maps before (768.1 on average) to between 790.1 and 802.9 in the seven after (795.7), a rise of 27.6 gauss, or 3.6 per cent (Figure H2b). Most of it comes in one step, from 773.9 gauss at 01:48 to 789.9 at 02:00. The region's upward flux grew 3.7 per cent over the three hours, so some slow growth rides under the step. For this flare Lu and co-workers list 286 ± 59 gauss in their region of interest, about ten times the box's rise, because their region is only the strips that changed.




Figure H2b. Region 11158 (HARP 377) through the X2.2 flare of 15 February 2011, as Figure H2. (A) Zeros of the sideways field inside the site, and the share of site loops whose count survived fifty noise copies (dotted, times 10). (B) The mean sideways field inside the site (narrow axis). (C) and (D) The site at 01:24 and 02:12 TAI: grey, the upward field; arrows, the sideways field where readable; red circles +1 and blue squares -1, the zeros; dashed box, the site. Dashed line: the flare's start, 01:44 UT. Data: SDO/HMI, SHARP series `hmi.sharp_cea_720s`, courtesy of NASA/SDO and the AIA, EVE, and HMI science teams. Run `SUN_SERIES_READ 20260927_060906`.

Read by Lu and co-workers' method, on purpose, the tool reproduces their step. Their method, a difference map across the flare, a rise of more than 100 gauss to pick the pixels, and a quiet pair as the check, was read in their paper first and then run here on the same maps, as a calibration of the book's tool against a published result; the method and the finding are theirs, and the book's two other reads of the imprint, the noise zones and the rim-only current, are its own. The imprint tool, now reading any listed region (run `SUN_FLARE_DIFFERENCE 20260927_061515; sun_flare_difference.py` version 2, which repeats both of region 401's runs number for number), compares 01:36 with 02:12, the maps either side of the flare, against two quiet pairs of the same 36-minute span, 00:36 to 01:12 and 02:48 to 03:24. Inside the site, 2,353 pixels rose by more than 100 gauss across the flare, against 1,091 and 1,332 in the quiet pairs. Followed through all fifteen maps, the chosen region sits flat before the flare (894.0 to 905.1 gauss), climbs through 951.1 at 01:48 and 1,129.4 at 02:00, and stays up (1,171.7 to 1,186.1 from 02:12 to 03:24; Figure H3b, D). Between the two windows the selection never used, 00:48 to 01:24 and 02:24 to 02:36, its step is 287.1 gauss. Lu and co-workers list 286 ± 59 . The quiet region picked the same way steps 60.2 gauss over the same windows, and the recipe's built-in rise on a quiet pair is 152.6 gauss, as in region 11166 (153.1): the wall of picking pixels is the same wall in both regions.

Bear's noise zones, again the cleanest. The four zones' quiet scatters came out at 64.5, 64.5, 65.1 and 90.3 gauss, so the zone floors (193.4 to 270.9 gauss) are again all above the flat 100. In the best-measured zone the pipeline's errors (13.2 to 39.4 gauss) again fall short of the quiet scatter. With the zone floors the FLARE pair passes 1,151 site pixels as rising and the quiet AFTER pair 239, a ratio of 4.8; the flat floor passes 2,280 and 1,217 (1.9), each pixel's own error 1,470 and 546 (2.7). The zones' region steps up 374.5 gauss and stays up; its quiet twin steps 27.3.

Bear's rim, much louder here. Across the flare, 60 of the 279 loops of side 16 centred in the site changed their current by more than the quiet pairs' 99th percentile (1.20×10^{11} amperes), against 6 of 252 and 7 of 302 in the quiet pairs (Figure H3b, E). They line the polarity line and the two spots either side of it. One count changed across the flare pair among 706 loops readable in both maps; the check below finds its cause.

HARP 377 (NOAA 11158), X2.2 flare of 15 February 2011: the imprint read three ways. Data: SDO/HMI SHARP; courtesy of NASA/SDO and the AIA, EVE, and HMI science teams




Figure H3b. Region 11158's X2.2 flare read three ways, as Figure H3. (A) The change in the sideways field from 01:36 to 02:12 TAI: red rises, blue falls; red contour, a rise of more than 100 gauss; black line, where the upward field is zero; circles, the loop walk's zeros at 01:36 and squares at 02:12 (green +1, purple -1). (B) The same for the quiet AFTER pair, 02:48 to 03:24. (C) Site pixels passed as rising by three floors. (D) Four regions of interest through all fifteen maps: solid, chosen by their rise across the flare; dotted, chosen the same way on the quiet AFTER pair; red, flat floor; blue, noise zones. Dashed line: the flare's start, 01:44 UT. (E) The share of loops of side 16 whose current, read from the rim alone, changed by more than the quiet pairs' 99th percentile. Data: SDO/HMI, SHARP series `hmi.sharp_cea_720s`, courtesy of NASA/SDO and the AIA, EVE, and HMI science teams. Run `SUN_FLARE_DIFFERENCE 20260927_061515`.

Sources: the fifteen SHARP records for HARP 377 from 00:36 to 03:24 TAI on 15 February 2011, 120 files fetched by `tools/sun_fetch_sharp.py` version 1 on Bear's laptop; Lu and co-workers 2019's table for the flare's times and its 286 ± 59 gauss. Identifiers: the reference list.

Runs: `SUN_SERIES_READ 20260927_060906` (`tools/sun_series_read.py` version 2 with `RUN_SUN_SERIES_READ.bat`; real data: the 120 files' combined sha256 heads the ledger; seed 20261017 plus the map's index), which prints every number of the series paragraphs except 1,352, 486 and 15, sums of its Part E lines, and the means 3.6 and 3.7 per cent, worked by script from its Part D; version 2, run on region 401's maps, printed a ledger identical to the run of record `SUN_SERIES_READ 20260927_041443` except for its version and two header lines (the region and the site), and a byte-identical Figure H2. `SUN_FLARE_DIFFERENCE 20260927_061515` (`tools/sun_flare_difference.py` version 2; the six maps it names, 48 files whose combined sha256 heads it) for the imprint paragraphs, except the ratios 4.8, 1.9 and 2.7, the shares of site loops and the percentage, worked by script from its counts; version 2, run on region 401, printed ledgers identical to both of that region's runs of record (20260927_045554 and 20260927_053821) except for its version and two header lines, and a byte-identical Figure H3.

Are the zeros real? The instrument's half-turn

The instrument measures the line along which the sideways field lies, but not which way along that line it points. The pipeline picks a direction at every pixel (the “180-degree ambiguity”; the SHARP maps resolve it by a minimum-energy method, as Lu and co-workers describe). Where it picks the other way on a small patch, the arrows jump by half a turn at the patch's edge. A loop crossed by that edge can then count a turn with no zero inside. So every zero the loop walk reports was tested (run `SUN_ZERO_CHECK 20260927_061935` for region 11166 and `20260927_061940` for region 11158; tools/sun_zero_check.py version 1). A true zero needs the field to fall to nothing somewhere inside its loops: the weakest sideways field there must be below three times its own error. A phantom has no weak field at all, 300 gauss or more at the weakest point, and a jump of more than 150 degrees between neighbouring arrows inside.

Most zeros are real. In region 11166's fifteen maps, of 103 sightings of zeros inside the site, 92 pass as true zeros, 9 are unsure and 2 have no weak field; outside the site 62 of 63 pass. The persistent saddle on the polarity line passes in every map (its weakest field is 3 to 99 gauss). In region 11158, of 59 sightings inside the site, 54 pass, 1 is unsure and 4 have no weak field; outside the site all 58 pass. The four phantoms are the three new “zeros” on the polarity line at 02:12 and one near them at 02:36, on the patch where the flare struck.

The one count that changed. In region 11158, the loop of side 16 at x 388 to 404, y 228 to 244 read 0 at 01:36 and at 02:00 and +1 at 02:12, with its whole rim readable in both the 02:00 and the 02:12 maps, the weakest rim field 435 and 382 gauss. Between those two maps its rim pixels hardly turned (the median turn was 1.2 degrees), and only one, at its corner, turned by more than 90 degrees: by 156. That pixel sits on the edge of a patch on the polarity line where, in the 7-by-7 block round it, 14 pixels turned by 155 to 180 degrees between the two maps: the pipeline's choice of direction flipped there. The check across the flare pairs agrees. The one count change in region 11158's flare pair had a rim pixel touched by such a flip, and no count changed on a rim that no flip touched. So wall one still holds for the Sun's field: its counts change only where a rim goes weak. But the instrument can move a count by itself, by flipping its half-turn choice on a patch the rim crosses, and on the polarity line, where the upward field is near zero, the choice is hardest to make.

The rim currents are not flips. Of the 79 loops whose current changed by more than the quiet 99th percentile across region 11158's flare, none had a rim pixel touched by a flip; in region 11166, 1 of 15. So the currents Bear's rim reads through a flare are the field's own.

What it adds. A loop count on a vector magnetogram carries one more wall than the book's other media: the instrument's half-turn choice, a pen in the reader's own hand. The test above separates true zeros from phantoms on a single map, and a flip between maps shows as a single rim pixel turning by nearly half a turn while its neighbours hardly move. Where to point a specialist: the people who resolve the 180-degree ambiguity (the SHARP pipeline uses a minimum-energy method; Metcalf 1994 and Leka and co-workers 2009, as Lu and co-workers cite them), with the question whether a loop's count, run across a map, could itself flag where the choice is unstable.

Sources: Lu and co-workers 2019 for the pipeline's minimum-energy method; Metcalf 1994 and Leka and co-workers 2009 as Lu and co-workers cite them, not read. Identifiers: the reference list.

Runs: `SUN_ZERO_CHECK 20260927_061935` (region 11166) and `SUN_ZERO_CHECK 20260927_061940` (region 11158), tools/sun_zero_check.py version 1. Every number in the section is printed in their ledgers except the history of the one loop whose count changed (its rim through 01:24 to 02:24, its pixels' turns, and the flipped patch round its corner), printed by a separate script while the cause was being found, whose numbers the run's Part C then confirmed.

Region 11158 along its line, and the tilt

The polarity-line tool, now reading any listed region (run `SUN_PIL_SHEAR 20260927_062644`; sun_pil_shear.py version 2, which repeats region 11166's run of record in its first three parts and adds a fourth), finds 231 to 299 line pixels per map in region 11158's site.

Here the field along the line steps up plainly. It runs from 804.5 to 853.9 gauss in the five maps before the flare (828.7 on average) and from 942.1 to 1,027.0 in the seven after (983.5): up 154.9 gauss, with the two ranges apart (Figure H4b, D). In region 11166 the same measure rose 33 gauss inside the scatter. This is Petrie's finding, that the change comes “almost entirely” from the part of the field parallel to the line, seen plainly. The mean angle barely moves (73.4 to 73.8 degrees), because the field across the line also rose a little, by 23.6 gauss. The share crossing the

wrong way fell from 0.238 to 0.187 (ranges 0.208 to 0.268 before, 0.165 to 0.213 after). Fewer bald-patch crossings after the flare is the direction Lee, Sun and Kazachenko saw in their flare, though here the change is small, and it is on the line, where the instrument's half-turn choice is hardest.

The tilt, against a potential field. The new fourth part computes, for each map, the current-free field whose upward part at the surface matches the map (by Fourier transform, treating the patch as flat), and compares it with the observed field in a band three pixels either side of the line: the tilt from vertical, and the shear, the angle between the observed and the potential sideways field. As a check, the potential field crosses the line nearly straight: on average within 15.7 to 28.5 degrees of straight across, in every map of either region. In region 11158 the observed field near the line stands a little more upright than the potential field before the flare (72.37 degrees from vertical against 73.55) and a little less so after (73.25 against 74.10), so the gap narrows from 1.18 degrees to 0.85. The shear against the potential field moves by 0.70 degrees, no pattern (Figure H5b). In region 11166 the observed field near the line is already more horizontal than the potential field (73.20 against 70.18 before the flare), its tilt moves by 0.73 degrees and the potential's by 1.11, and the shear drifts up by 1.99 degrees through the whole series, flare or no flare. So the band read agrees with Petrie that the shear against a potential field shows no flare pattern. But the tilt collapse he saw most clearly comes to less than a degree here, inside the scatter. His regions along the line and his potential fields are chosen differently, and whether a band three pixels wide is the wrong place to look is a question for a solar physicist.

HARP 377 (NOAA 11158) round the X2.2 flare of 15 February 2011: the sideways field on the polarity inversion line (dashed: flare, 01:44 UT). Data: SDO/HMI SHARP; courtesy of NASA/SDO and the AIA, EVE, and HMI science teams




Figure H4b. The sideways field on region 11158's polarity inversion line through the X2.2 flare, as Figure H4: (A) and (B) the line at 01:24 and 02:12 TAI coloured by angle (blue crosses the right way, white runs along the line, red crosses the wrong way); (C) the mean angle; (D) the mean field along the line (red) and across it (grey); (E) the shares past 45 and past 90 degrees. Dashed line: the flare's start, 01:44 UT. Data: SDO/HMI, SHARP series `hmi.sharp_cea_720s`, courtesy of NASA/SDO and the AIA, EVE, and HMI science teams. Run `SUN_PIL_SHEAR 20260927_062644`.




Figure H5b. Region 11158: (A) the tilt from vertical of the observed field (red) and of the potential field (blue, dashed) in a band three pixels either side of the polarity line; (B) the shear, the angle between the observed and the potential sideways field, in the same band. Dashed line: the flare’s start, 01:44 UT. Data as Figure H4b. Run `SUN_PIL_SHEAR 20260927_062644`; region 11166’s tilt figure is `fig_H5_tilt_401.png` of run `SUN_PIL_SHEAR 20260927_062635`.

Sources: Petrie 2012 and Lee, Sun and Kazachenko 2021, as in the sections above. Identifiers: the reference list.

Runs: `SUN_PIL_SHEAR 20260927_062644` (tools/sun_pil_shear.py version 2 with `RUN_SUN_PIL_SHEAR.bat`) for region 11158, and `SUN_PIL_SHEAR 20260927_062635` of the same version for region 11166’s tilt, whose first three parts are identical to the run of record `SUN_PIL_SHEAR 20260927_051041` and whose Figure H4 is byte-identical to it. Every number in the section is printed in their ledgers except the gaps 1.18 and 0.85 degrees, worked from the ledger’s means.

Rims grown round a sunspot: does the current come back?

Venkatakrishnan and Tiwari (2009) found, in twelve sunspots seen by the Hinode satellite, that the part of the field circling a spot fades faster than one over the distance in the spot’s outer ring. By Ampère’s law that means the current through a circle round the spot rises, then falls back toward zero: the current that flows one way through the spot’s middle flows back the other way round its edge, and the spot as a whole carries none. Bear’s rim reads this directly. Centre a square loop on a spot, read the current through it from its rim alone, and grow it.

The run (tools/sun_rim_growth.py version 1; run `SUN_RIM_GROWTH 20260927_064149` for region 11166 and `20260927_064153` for region 11158) takes the spots that lie away from each flare site: three in region 11166, the ones the book’s first Sun read found, and the two outer spots of region 11158. In each of the fifteen maps it centres the loops on the spot’s strongest field and grows them two pixels at a time, stopping if the rim leaves the region or reaches a half-side of 60 pixels (21.9 megametres). In every map the loops reached the full 60. The test is one number: the current through the largest loop as a share of the largest current through any loop, the spot’s own. Near 0, the current came back; near 1, it did not.

One spot in each region comes back. Region 11166’s spot 1 (downward field, $-2,821$ gauss at its centre in the first map) peaks at a median 3.57×10^{11} amperes at a half-side of 24 pixels (8.7 megametres). The largest loop keeps a median 0.12 of that (0.01 to 0.45 over the fifteen maps), and in 5 of the 15 maps its current has turned the other way. Region 11158’s spot 1 (upward field, 2,380 gauss) peaks at 7.28×10^{11} amperes at 22 pixels (8.0 megametres) and keeps 0.17 (0.00 to 0.26), turned the other way in 1 map. This is Venkatakrishnan and Tiwari’s finding, read with a square rim on another instrument. In the figures, the flux inside these loops (grey) levels off past the peak: the loop has taken in the spot, and beyond it adds only weaker field.




Figure H6. Bear's rim grown round the three spots of region 11166 that lie away from the flare site, in all fifteen maps from 22:00 to 00:48 TAI. Red: the current through a square loop centred on the spot, read from its rim alone, against the loop's half-side. Grey (right axis): the unsigned upward flux inside the loop. Spot 1's current rises, peaks and falls back toward zero; spots 2 and 3, neighbours of the same sign, keep much of theirs. Data: SDO/HMI, SHARP series hmi.sharp_cea_720s, courtesy of NASA/SDO and the AIA, EVE, and HMI science teams. Run SUN_RIM_GROWTH 20260927_064149.

The other three keep a good part. Region 11166's spots 2 and 3 keep a median 0.45 and 0.67 of their peak currents (2.32×10^{12} amperes at 15.3 megametres, and 9.09×10^{11} at 12.4), and region 11158's spot 2 keeps 0.35 (1.89×10^{12} amperes at 13.8 megametres). None of the three turns the other way in any map, and the flux inside their loops keeps climbing to the largest loop. For region 11166's pair the reason is plain: spots 2 and 3 are neighbours of the same sign, their centres about 55 pixels (20 megametres) apart, so the larger loops round either take in the other, and its current with it. What region 11158's spot 2 takes in beyond its peak, the run does not say. A single flux number did not sort the spots: the share of the largest loop's flux already inside the peak loop has medians of 0.58 to 0.71 for all five. The shape of the whole curve, flux levelling off or still climbing, is what tells a spot that stands alone from one that does not.




Figure H6b. The same for the two outer spots of region 11158, in all fifteen maps from 00:36 to 03:24 TAI. Spot 1's current peaks and falls back toward zero; spot 2 keeps about a third of its peak. Data: SDO/HMI, SHARP series hmi.sharp_cea_720s, courtesy of NASA/SDO and the AIA, EVE, and HMI science teams. Run SUN_RIM_GROWTH 20260927_064153.

The half-turn wall again. The outer rims run through the spot's surroundings, where the field is weaker and the pipeline is often not confident which way along its line each arrow points. On the largest loop, only 0.06 to 0.13 of the rim pixels are confident round the two spots that came back, and 0.23 to 0.51 round the other three. The rim adds up every pixel on it. Arrows pointing the wrong way at random would pull that sum toward zero, so part of a spot's return to zero could be the pipeline's choice and not the Sun's. Where to point a specialist: whether a spot's return to zero survives other ways of resolving the half-turn.

Every spot's current flows up, and the instrument turns it. All five peak currents came out positive, whichever the spot's sign. The follow-on paper on the Sun took this apart. Split the current in each band of field strength into a part that flips sign with the spot, a twist, which the Sun's own physics can make, and a part of the same sign in both polarities, an offset, which it has no reason to prefer, since its equations treat a field and its mirror image alike. Near disk centre the offset points up in the strong field of spots, in the two regions here, in two more followed across the disk, and in 24 of 26 regions never used in any fit, read blind in both hemispheres. It behaves like the instrument turning the measured direction of the sideways field by a few degrees, one way where the field points up and the other way where it points down. A turning of about 5.5 degrees of that kind, put into these two regions on purpose, reproduces their offset; over every region near disk centre the turning needed is about 3 degrees, and a second blind test on fifty fresh regions put its median at 3.2. It holds steady through the satellite's daily orbit, blurring does not make it, and it is made under both versions of the instrument's inversion code. Toward the limb it fades to about nothing and does not change sign. The follow-on paper traced it as far as the light: a one-layer fit whose magneto-optical terms are too weak turns the arrows by the right few degrees with the right polarity flip but the wrong shape across field strength, and a textbook fit of the instrument's own polarised light disagrees with the instrument's arrows by an S-shaped few degrees, so where the turning is made is narrowed to the inversion's reading of the light and left with the instrument's team. The book's count is not touched. Put the turning into 38 maps and 31 of 154,267 readable loop counts change, and all 31 sit on five steps of nearly half a turn already in the maps, four of them straddling the line where the field turns from up to down. The current is a measure and takes the turning; the count is a label and reads through it. So the five positive peaks carry the instrument's offset as well as the Sun's current, and how much of each is not measured. A turning of this kind, from Faraday rotation in the spectral line, was found to make false currents of the same size in a filter magnetograph and corrected there (Gao, Su, Xu and Zhang 2008), and HMI's own inversion models such effects (Centeno and co-workers 2014). Read in full afterwards, in its arXiv preprint, that paper's appendix reports that the inversion's treatment of the magneto-optical effects in strong fields once carried an error, a missing factor of two, fixed in the pipeline's fd10 version, which Hoeksema and co-workers 2014, the pipeline's overview, says has been run on all of HMI's vector data since 1 May 2010, so the maps the book reads were made with that fix in place. The same signature was not found reported for HMI's SHARP maps in what was read. Where to point a specialist: the HMI team.

Who met this before. Venkatakrishnan and Tiwari, as above, with circles on Hinode maps. Georgoulis, Titov and Mikić (2012), on the sharpest maps they had, found that "well-formed, intense magnetic polarity inversion lines" are "the only photospheric magnetic structures that support significant net current", and that "magnetically isolated active regions remain globally current-balanced". Liu and co-workers (2017) compared four regions in HMI maps and found that the ones which threw off mass into space "are characterized by a strongly non-neutralized total current", while in the others it "is almost perfectly neutralized". Region 11158 is one of theirs, among the eruptive ones. Kontogiannis and co-workers (2017), from the SHARP maps the book reads, found the net currents of flaring regions "higher by more than an order of magnitude than in non-flaring regions", and a better flare forecaster than the total flux. Sun and Cheung (2020) showed that a net current at the surface can come from geometry alone, a twisted tube crossing the surface at a slant. Their way in is the current density, worked out pixel by pixel from the field's curl and added up by sign; the book's rim gives the current through a whole loop from its edge alone. The book's rims circled spots away from the flaring line, and those that stood alone came back to near zero, as the others would expect; the net current they find sits on the strong polarity line, which these rims did not circle. Not found in what was read: square loops grown round each spot on SHARP maps as a cheap test, spot by spot, of whether its current comes back. A door to show a specialist, not a claim.

Sources: the papers in "Who met this before", found by web search after the run, on 26 September, near midnight (Pacific): Venkatakrishnan and Tiwari 2009 as above; Georgoulis, Titov and Mikić 2012's abstract from its arXiv record (1210.2919); Liu and co-workers 2017's abstract from its arXiv record (1708.04411), and their four regions and which of them erupted from the PDF through a page-summary tool; Kontogiannis and co-workers 2017's abstract from the journal's page (DOI 10.1007/s11207-017-1185-1); Sun and Cheung 2020's through a page-summary tool (arXiv 2011.11873); the journal details as the prior-art ledger (S8) gives them. For the instrument's turning: Gao, Su, Xu and Zhang 2008, read at abstract level and through a page-summary tool, and Centeno and co-workers 2014, read through a page-summary tool, both records checked against their publishers' pages on 27 September, and Centeno and co-workers then read in full as arXiv 1403.3677v1 on 28 September; Hoeksema and co-workers 2014 (arXiv 1404.1881v1) for the pipeline's versions; the Sun paper, *Dancing on the Sun: A loop count that reads through helonomy, and the current it finds flowing in both polarities*, for the offset. Identifiers: the reference list.

Runs: SUN_RIM_GROWTH 20260927_064149 (region 11166: the fifteen maps of the flare series, 75 files whose combined sha256 heads the ledger) and SUN_RIM_GROWTH 20260927_064153 (region 11158, the same way), tools/sun_rim_growth.py version 1 with RUN_SUN_RIM_GROWTH.bat; no random numbers. Every number of the rim paragraphs is printed in their ledgers except the distance between region 11166's spots 2 and 3 (55.4 pixels, 20.2 megametres, at 22:00), worked by script from the ledger's centres and the map's pixel size. Earlier runs of the same tool, minutes before, printed the same numbers and were redone to put the figures' credit on a line of its own, to add the flux share and the summary lines, and to stop calling all five spots isolated, before anything shipped; they are not kept. The paragraph on every spot's current rests on the Sun spoke's runs, cited from that companion: SUN_CURRENT_OFFSET 20260927_114851 (the split into twist and offset, on the two regions here), SUN_OFFSET_INJECTION 20260927_132638 (the turning put in on purpose; 5.5 degrees), SUN_OFFSET_REMOVAL 20260927_193239 (the satellite's daily orbit, and the count: 31 of 154,267 loops in 38 maps), SUN_OFFSET_VIEW 20260928_025857 (blurring), SUN_OFFSET_BLIND 20260928_041724 (24 of 26 regions, whose prediction was written and shipped before the maps were fetched), SUN_HALF_COUNT 20260928_043627 (the five steps under the 31 changed counts), SUN_TURNING_BLIND 20260928_181524 (the second blind test: 45 of 50 fresh regions near centre, median +3.24 degrees; the fade toward the limb), SUN_STOKES_FIT 20260928_225848 and SUN_TURNING_PLANT 20260928_225126 (the mis-modelled magneto-optical term: right size and polarity, wrong shape) and SUN_STOKES_REFIT 20260928_235356 (a textbook fit of HMI's own light against HMI's arrows).

What comes next

The first three domains each have a run, their papers read in full (six papers, LATCH_ATLAS_DOMAIN_READING_NOTES_2026-09-26) and a card (LATCH_ATLAS_NEW_CARDS_2026-09-26), for Bear to read. Magnetic skyrmions, the fifth domain, have their run, their papers read in full (Bear's Paper A, Mühlbauer and co-workers 2009, Romming and co-workers 2013) and their card, the fourth in LATCH_ATLAS_NEW_CARDS_2026-09-26. The equatorial waves' paper is read (reading notes, section 11): a read-only medium for READ section 8, with no card of its own. And the Sun has now been read from a real magnetogram, and through a flare, above. With the rims grown round its spots, the Sun's runs for the card were done, and its card is in Part Two. The follow-on paper on the Sun, *Dancing on the Sun*, has its first answer, on why every spot's current flows up, folded in above.

Sources: the five papers of the first three domains, read in full on 26 September (Berger and Field 1984, Taylor 1986, Coletta and co-workers 2016, Alexander and co-workers 2012, Tkalec and co-workers 2011), and the skyrmion, equatorial-wave and Sun readings, all written up in LATCH_ATLAS_DOMAIN_READING_NOTES_2026-09-26.

Back matter

Acknowledgements

The book stands on the papers it read, and each is named where it is used: among them Berry's 1984 paper on the geometric phase and Tomita and Chiao's coiled fibre; Moffatt and Ricca on Călugăreanu's theorem; Brown and Cozzarelli on DNA gyrase; Kirtley, Tsuei and Tafuri, Monaco and co-workers, Weir and co-workers, Corman and co-workers, Aidelsburger and co-workers, and Scherer and co-workers on rings frozen from pieces; Gurtovoi and co-workers on the chain of ring ratchets; Groisman and co-workers on rings of oscillators; Eckel and co-workers and Ryu and co-workers on rings of superfluid gas; Puckett and co-workers on the ring of light; and the lecturers whose CERN report of 2004 (CERN-2004-008) supplied the superconducting constants and the first flux staircases. For the media added later: Berger and Field, and Taylor, on helicity and relaxation; Pevtsov and co-workers' review of the Sun's helicity; Coletta, Delabays, Adagideli and Jacquod on windings in power grids; Alexander, Chen, Matsumoto and Kamien, and Tkalec and co-workers, on liquid crystals and their knots; Mühlbauer and co-workers and Romming and co-workers on magnetic skyrmions; Delplace, Marston and Venaille on the equatorial waves; Bobra and co-workers, Hoeksema and co-workers, Lu and co-workers, Petrie, and Venkatakrishnan and Tiwari on reading the Sun's field; Gao, Su, Xu and Zhang, and Centeno and co-workers, on how a magnetograph can turn the field it reports; and the many whose work the prior-art ledger credits, from the fields each run walked into.

It also stands on the deposited work it draws from: the Field Guide to loop walking, the Latch paper, the Magnetic Forest and its companion, the Crossover Atlas, the Jellyfish paper, the Cascadia paper and the sky, muon and whale papers. Amber, the book's picture of one true pattern and its shadows, is Roger Zelazny's, from *The Chronicles of Amber*, beginning with *Nine Princes in Amber* (1970). Bear's field experience is credited in the text where it appears: the stranded and nicked wire, his judgement of the watch picture and his ratchet, the deadband and supervisory control, and the S rotor he has drawn for decades.

References

One entry per work the book cites, with how far it was read, in the book's own words: read in full, read in part, at the level of its abstract, through a summary, title only, or as another work cites it. Every claim should be taken at that level and no further. The list is built by script from the book's reference table, and the identifiers still to be confirmed against their publishers' records are listed in the book's crate.

The author's own records

- M. "Bear" Schulz, *Loop Walking Cascadia: a Triangle Lattice, its Rogues and Witnesses, and One Opening of the Vault*, version 1.0 of 24 August 2026. Zenodo, concept DOI 10.5281/zenodo.22076451. *How far it was read: read in full (its markdown and the reports of runs PACHINKO_SUITE 20260821_191042, WINDING_BUDGET 20260821_190021 and WATERFALL_MAP 20260821_200450, from the deposit crate, 24 September). Cited in: READ 8.*
- M. "Bear" Schulz, *The Crossover Atlas*, version 3 of 22 August 2026. Zenodo, concept DOI 10.5281/zenodo.22064125. *How far it was read: read in full (Chapters 7 and 8, with receipts gauss_bonnet.py, spherical_excess.py and berry_monopole.py, 24 September). Cited in: READ 2, 6.*
- M. "Bear" Schulz, *Loop Walking: A Field Guide*. Zenodo, concept DOI 10.5281/zenodo.21811588. Case Four (muons) is in draft v4.6 of 9 August 2026, with its crate's muonic_lines.py and falsify_v46.py outputs; register rows 147, 305 and 438 for page one. *How far it was read: read in full (the v4.6 crate; v6 read in part). Cited in: opening; page one; READ 8.*
- M. "Bear" Schulz, *Instruments for the Magnetic Forest* (2026). Zenodo, concept DOI 10.5281/zenodo.22906780 (version 2 record 10.5281/zenodo.22910863). *How far it was read: read in full (sections 1 to 8 and 10; its Figures 3 and 4 reproduced unchanged). Cited in: READ 4.*
- M. "Bear" Schulz, *Jellyfish: the Holonomist of the Sea* (2026). Zenodo, concept DOI 10.5281/zenodo.22908909. *How far it was read: read in full (23 September; Appendix C, line 1200, for the quotation). Cited in: opening; READ 6.*
- M. "Bear" Schulz, *Linking Is the Latch*, version 2 (2026). Zenodo, concept DOI 10.5281/zenodo.22906667 (version 2 record 10.5281/zenodo.22929828). *How far it was read: read in full (23 and 24 September; section "The instrument"; lines 67 to 79 for the staircase). Cited in: page one; READ 1, 3, 7.*

- M. “Bear” Schulz, *Reading the Sky Through Loops* (2026). Zenodo, concept DOI 10.5281/zenodo.21331651 (resolved against Zenodo on 12 September). *How far it was read: read in full (draft v0.3 of 12 July, from the deposit crate, nine checksums verified, 24 September). Cited in: READ 8.*
- M. “Bear” Schulz, *Dancing on the Sun: A loop count that reads through holonomy, and the current it finds flowing in both polarities* (2026), the companion paper on the Sun, deposited on Zenodo after this book; its runs SUN_CURRENT_OFFSET 20260927_114851 through SUN_STOKES_REFIT 20260928_235356. *How far it was read: this team’s companion. Cited in: Part Four, the Sun (every spot’s current).*
- M. “Bear” Schulz, *Comparative Biosonar & Navigation* (2026). Zenodo, concept DOI 10.5281/zenodo.20502057 (resolved from the API on 29 September; 20502058 is the version DOI). *How far it was read: read in full (documents 00 to 04; document 02 result B3; document 03 entries 1, 4 and 5). Cited in: READ 8; page one.*

Everything else, by first author

- Edwin A. Abbott, *Flatland: A Romance of Many Dimensions* (1884). *How far it was read: not read (named for the flat world). Cited in: READ 6.*
- Colin C. Adams, *The Knot Book: An Elementary Introduction to the Mathematical Theory of Knots*, W. H. Freeman (1994; AMS reprint 2004): section 1.4, “Links”, pp. 18–21. *How far it was read: read in part (section 1.4 as printed, 29 September). Cited in: READ 6.*
- M. Aidelsburger and co-workers, “Relaxation dynamics in the merging of N independent condensates”, arXiv 1705.02650. *How far it was read: read in full (arXiv). Cited in: W4.*
- G. P. Alexander, B. G. Chen, E. A. Matsumoto and R. D. Kamien, “Colloquium: Disclination loops, point defects, and all that in nematic liquid crystals”, *Reviews of Modern Physics* 84, 497 (2012). DOI 10.1103/RevModPhys.84.497 (from the ledger). *How far it was read: read in full. Cited in: liquid-crystal card; Part Four; Part Four, liquid crystals.*
- Allen, Beijersbergen, Spreeuw and Woerdman 1992: *Physical Review A* 45, 8185 (1992). *How far it was read: as cited by another (Allen and Padgett). Cited in: light card.*
- L. Allen and M. Padgett, chapter 1 of *Twisted Photons* (Wiley-VCH, 2011). *How far it was read: read in part. Cited in: light card; sound card.*
- T. Amari and co-workers (2003), helicity as a necessary but not sufficient condition for ejection. *How far it was read: as cited by another (the review). Cited in: plasma card.*
- A. Ambastha and R. A. Maurya (2012), on field changes during flares. *How far it was read: abstract. Cited in: Part Four, the Sun.*
- Ampère’s law: textbook physics. *How far it was read: textbook. Cited in: READ 1.*
- Andersen and co-workers 2006: *Physical Review Letters* 97, 170406 (2006). *How far it was read: abstract. Cited in: light card; WRITE W7.*
- V. I. Arnold and B. A. Khesin, *Topological Methods in Hydrodynamics*, Springer, Applied Mathematical Sciences 125 (1998; second edition 2021). *How far it was read: not read (checked against the publisher record). Cited in: page one.*
- Arnold’s floor on the energy of a linked field, as Khesin (2005) states it, with Freedman and He’s bound (1991). *How far it was read: through a summary (Khesin); not read (Arnold; Freedman and He). Cited in: page one; the plasma card.*
- ARTop: the Active Region Topology code, its README on GitHub. *How far it was read: through a summary. Cited in: Part Four, the Sun.*
- S. M. Barnett, M. Babiker and M. J. Padgett, “Optical orbital angular momentum”, *Philosophical Transactions of the Royal Society A* 375, 20150444 (2017), supplied by Bear. *How far it was read: abstract. Cited in: the appendix, light.*
- Bačić and co-workers 2025: winding numbers of phase-locked states extended to cell complexes (2025). *How far it was read: through a summary (the prior-art ledger, G2). Cited in: Who came before.*
- Bear’s field experience: the wiring rule (W2, W3); the ratchet, the deadband and supervisory control (W8); the watch picture (W8). grid equipment (power-grid card); credited in the text where it appears. *How far it was read: this book. Cited in: power-grid card; W2, W3, W8.*

- A. A. Belavin and A. M. Polyakov, “Metastable states of two-dimensional isotropic ferromagnets”, *JETP Letters* 22, 245 (1975), the bound on a skyrmion’s energy. *How far it was read: not read (recalled)*. Cited in: *Part Four, skyrmions*.
- M. Beresna and co-workers, “Radially polarized optical vortex converter created by femtosecond laser nanostructuring of glass”, *Applied Physics Letters* 98, 201101 (2011). *How far it was read: title only (the prior-art ledger, L3)*. Cited in: *Who came before*.
- M. A. Berger and G. B. Field, “The topological properties of magnetic helicity”, *Journal of Fluid Mechanics* 147, 133–148 (1984). DOI 10.1017/S0022112084002019 (expected; the ledger gives its reference list as read in full). *How far it was read: read in full*. Cited in: *page one; plasma card; Part Four*.
- M. A. Berger and A. Ruzmaikin, “Rate of helicity production by solar rotation”, *Journal of Geophysical Research* 105, 10481 (2000). *How far it was read: as cited by another (Pevtsov and co-workers 2014)*. Cited in: *plasma card*.
- M. V. Berry, “Quantal phase factors accompanying adiabatic changes”, *Proceedings of the Royal Society A* 392, 45–57 (1984). *How far it was read: read in full (25 September, the Royal Society’s copy; every quoted equation checked against the page images; journal record checked 25 September)*. Cited in: *opening; READ 2*.
- M. G. Bobra and co-workers, the SHARP series paper, arXiv 1404.1879 (2014); *Solar Physics* 289, 3549 (2014), expected. *How far it was read: read in part (the series and its Figure 2)*. Cited in: *Part Four, the Sun*.
- Boltzmann’s constant times room temperature: 0.02526 eV at 20 °C, 0.02585 eV at 27 °C. *How far it was read: computed from constants*. Cited in: *page one*.
- J. M. Borrero and co-workers, the VFISV Milne-Eddington inversion code for HMI. *How far it was read: as cited by another*. Cited in: *the Sun spoke’s text, not this book’s (carried here from the prior-art ledger; drop at the alphabetical build unless the “Who came before” list wants it)*.
- P. O. Brown and N. R. Cozzarelli, “A sign inversion mechanism for enzymatic supercoiling of DNA”, *Science* 206, 1081–1083 (1979). *How far it was read: read in part (“text read”); read in part*. Cited in: *READ 3; WRITE W6; DNA card; W1, W2, W6*.
- Capri Tools’ buyer’s guides: “How Many Teeth in a Ratchet? 72 vs 90 Tooth Count Explained” and “6-Point vs 12-Point Sockets and Wrenches: When to Use Each”. *How far it was read: read in full (web pages)*. Cited in: *W8*.
- the Cascadia paper, section A; described in Part Four from Bear’s project notes. §10.1 and run WATERFALL_MAP 20260821_200450. *How far it was read: read in full; read in full (the paper)*. Cited in: *W1, W5; Part Four, skyrmions*.
- Castellanos Durán and co-workers 2018; Alanazi and Fletcher 2026: as the prior-art ledger names them. *How far it was read: abstract / through a summary*. Cited in: *Part Four, the Sun*.
- R. Centeno, J. Schou, K. Hayashi, A. Norton, J. T. Hoeksema, Y. Liu, K. D. Leka and G. Barnes, “The Helioseismic and Magnetic Imager (HMI) Vector Magnetic Field Pipeline: Optimization of the Spectral Line Inversion Code”, *Solar Physics* 289, 3531–3547 (2014). DOI 10.1007/s11207-014-0497-7 (record checked 27 September); arXiv 1403.3677. *How far it was read: read in full (arXiv v1, 28 September)*. Cited in: *Part Four, the Sun*.
- P. Schmüser’s lectures in the CERN Accelerator School report CERN-2004-008 (2004): niobium’s constants, the CMS magnet, the cavity, the magnet decay. the copy supplied 23 September, searched for the passages quoted. *How far it was read: read in part*. Cited in: *superconducting card; WRITE; W2, W3, W4*.
- J. J. Champoux, “DNA topoisomerases: structure, function, and mechanism”, *Annual Review of Biochemistry* 70, 369 (2001). *How far it was read: as cited by another (Wikipedia, “Type I topoisomerase”)*. Cited in: *WRITE W8*.
- Chen, Mandić, Slaughter, Tanaka, Kikkawa, Collings and Yodh (2024), the line tension of a disclination. *How far it was read: through a summary*. Cited in: *liquid-crystal card*.
- M. “Bear” Schulz, “A Coincidence Has No Single Price”, version 2 of 22 August 2026, preprint, CC BY 4.0; cited for its method only. *How far it was read: read in full*. Cited in: *W4*.
- T. Coletta, R. Delabays, I. Adagideli and Ph. Jacquod, “Topologically protected loop flows in high voltage AC power grids”, *New Journal of Physics* 18, 103042 (2016). DOI 10.1088/1367-2630/18/10/103042. journal details from its DOAJ record (26 September), quotations from the arXiv text. *How far it was read: read in full*. Cited in: *power-grid card; Part Four; Part Four, the grid*.
- L. Corman and co-workers, “Quench-induced supercurrents in an annular Bose gas”, *Physical Review Letters* 113, 135302 (2014). arXiv 1406.4073. *How far it was read: abstract (arXiv v2 read for W4); read in full (arXiv v2)*. Cited in: *ring-superfluid card; WRITE W4; W4*.
- the Crossover Atlas: Chapter 6 and its receipt ssh_edge.py. *How far it was read: read in full*. Cited in: *W3, W4*.

- G. Călugăreanu, the theorem that linking equals twist plus writhe (1959, 1961), named after him as Moffatt and Ricca argue. *How far it was read: as cited by another (Moffatt and Ricca 1992). Cited in: READ 2.*
- V. R. Daria, D. Z. Palima and J. Glückstad, *Optics Express* 19, 476 (2011), the spatial-light-modulator twist of up to 20 turns and the beads driven along a spiral. *How far it was read: read in full (W7; the light card's line says "text read", the same reading); read in full. Cited in: light card; W7; W7.*
- B. S. Deaver and W. M. Fairbank, *Physical Review Letters* 7, 43 (1961), flux quantization. *How far it was read: not read (record); as cited by another. Cited in: superconducting card; W2.*
- Degarmo, Black and Kohser 2003: E. P. DeGarmo, J. T. Black and R. A. Kohser, *Materials and Processes in Manufacturing* (9th edition, Wiley, 2003). *How far it was read: as cited by another (Wikipedia, "Work hardening"). Cited in: WRITE W4.*
- R. Delabays, T. Coletta and Ph. Jacquod, "Multistability of Phase-Locking and Topological Winding Numbers in Locally Coupled Kuramoto Models on Single-Loop Networks", *Journal of Mathematical Physics* 57, 032701 (2016); arXiv 1512.04266. *How far it was read: abstract. Cited in: power-grid card.*
- P. Delplace, J. B. Marston and A. Venaille, "Topological Origin of Equatorial Waves", *Science* 358, 1075–1077 (2017). DOI 10.1126/science.aan8819; arXiv 1702.07583. *How far it was read: read in full (arXiv v3, main text and supplement; reading notes section 11). Cited in: opening; READ 8.*
- C. E. M. Demore and co-workers, *Physical Review Letters* 108, 194301 (2012). *How far it was read: not read (record located). Cited in: sound card.*
- M. R. Dennis and J. H. Hannay, "Geometry of Călugăreanu's theorem", arXiv math-ph/0503012. *How far it was read: read in part. Cited in: W6.*
- textbook DNA topology values. *How far it was read: textbook, not sourced to one book. Cited in: DNA card (DIAL, STORE); run DNA_ESTIMATES.*
- R. Doll and M. Näbauer, *Physical Review Letters* 7, 51 (1961), flux quantization. *How far it was read: as cited by another (the CERN lectures). Cited in: W2, W4.*
- S. Eckel and co-workers, "Hysteresis in a quantized superfluid atomtronic circuit", *Nature* 506, 200–203 (2014); arXiv 1402.2958. *How far it was read: abstract (and parameters); abstract and parameters. Cited in: ring-superfluid card; WRITE; W2, W5.*
- the synchronising guide; ENTSO-E's sheet on phase-shifting transformers: two web documents, through a page-summary tool. *How far it was read: through a summary. Cited in: power-grid card.*
- Wikipedia's pages "Euler characteristic" and "Linking number", read 29 September 2026. *How far it was read: through a summary. Cited in: READ 6.*
- C. W. F. Everitt and co-workers, "Gravity Probe B: Final Results of a Space Experiment to Test General Relativity", *Physical Review Letters* 106, 221101 (2011): geodetic $6,601.8 \pm 18.3$ mas/yr against 6,606.1 predicted. *How far it was read: not read (recalled). Cited in: READ 8.*
- G. Eyink and co-workers, "Flux-freezing breakdown in high-conductivity magnetohydrodynamic turbulence", *Nature* 497, 466 (2013). *How far it was read: abstract (the prior-art ledger, S2). Cited in: Who came before.*
- Farozan and co-workers 2025: a study of the static torque of single-stage and twisted Savonius rotors (2025). *How far it was read: through a summary (the prior-art ledger, H1). Cited in: Who came before.*
- Fayyaz and co-workers 2026: a ring-antenna design study for twisted radio beams, *PLOS ONE* (2026). *How far it was read: through a summary (the prior-art ledger, R1). Cited in: Who came before.*
- the Field Guide: draft v4.6 §8 and appendix B.2; v6 rows 51, 66, 172 and 443. *How far it was read: read in full (v4.6); read in part (v6). Cited in: W3, W4.*
- G. J. File and R. G. Mills, *Physical Review Letters* 10, 93 (1963), through NASA SP-140 and CERN-2004-008. *How far it was read: as cited by another. Cited in: superconducting card.*
- D. Foster, C. Kind, P. J. Ackerman, J.-S. B. Tai, M. R. Dennis and I. I. Smalyukh, "Two-dimensional skyrmion bags in liquid crystals and ferromagnets", *Nature Physics* 15, 655–659 (2019). DOI 10.1038/s41567-019-0476-x (confirmed against the publisher record, per the ledger). *How far it was read: through a summary. Cited in: skyrmion card.*
- Foucault's pendulum, first shown in Paris in 1851. *How far it was read: textbook, not sourced to one book. Cited in: READ 2.*
- M. H. Freedman and Z.-X. He, "Divergence-free fields: energy and asymptotic crossing number", *Annals of Mathematics* 134(1), 189–229 (1991). DOI 10.2307/2944336 (from the ledger, record read). *How far it was read: not read (record). Cited in: page one; plasma card.*

- Gao, Su, Xu and Zhang 2008: “Statistical removal of Faraday rotation in vector magnetograms taken by filter-type magnetographs”, *Monthly Notices of the Royal Astronomical Society* 386, 1959–1965 (2008). DOI 10.1111/j.1365-2966.2008.12848.x (record checked 27 September). *How far it was read: abstract, and through a summary. Cited in: Part Four, the Sun (every spot’s current).*
- C. F. Gauss, the linking integral, from his notes of the 1830s. *How far it was read: not read (textbook; to source, e.g. through Rolfsen 5.D, which gives the integral as one of eight definitions). Cited in: READ 1.*
- M. K. Georgoulis and co-workers (2009), the helicity budget of active regions. *How far it was read: as cited by another (the review). Cited in: plasma card.*
- M. K. Georgoulis, V. S. Titov and Z. Mikić, “Non-neutralized electric current patterns in solar active regions: origin of the shear-generating Lorentz force”, *Astrophysical Journal* 761, 61 (2012); arXiv 1210.2919. *How far it was read: abstract. Cited in: plasma card; Part Four.*
- González, Molina-Terriza and Torres 2006: *Optics Express* 14, 9093 (2006). *How far it was read: read in part. Cited in: light card; W7.*
- C. E. Gough and co-workers, *Nature* 326, 855 (1987), flux quantization in a high-temperature superconductor. *How far it was read: as cited by another (the CERN lectures). Cited in: W2.*
- P. Groisman, C. De Vita, J. Fernández Bonder and Y. Zhang, “The size of the sync basin resolved”, arXiv 2506.03419v2 (2025). *How far it was read: read in full (arXiv). Cited in: W4; the acknowledgements.*
- Grudin, Ilchenko and Maleki 2006: *Physical Review A* 74, 063806 (2006). *How far it was read: as cited by another (Puckett). Cited in: light card.*
- V. L. Gurtovoi, M. Exarchos, V. N. Antonov, A. V. Nikulov and V. A. Tulin, “Multiple superconducting ring ratchets for ultrasensitive detection of non-equilibrium noise”, arXiv 1507.05957v1 (2015). *How far it was read: read in full (arXiv). Cited in: W4.*
- B. B. Guzina, O. Oudghiri-Idrissi and S. Meng (2022), supplied with Berry and Tomita and Chiao. *How far it was read: abstract. Cited in: READ 2.*
- J. Hagemeister, N. Romming, K. von Bergmann, E. Y. Vedmedenko and R. Wiesendanger, “Stability of single skyrmionic bits”, *Nature Communications* 6, 8455 (2015). DOI 10.1038/ncomms9455. *How far it was read: through a summary. Cited in: skyrmion card.*
- Allen Hatcher, *Algebraic Topology*, Cambridge University Press (2002): Theorem 2.44, p. 146; Corollary 3.37, p. 249. *How far it was read: read in part (the two statements as printed, 29 September, from the copy Bear supplied). Cited in: READ 6.*
- He, Friese, Heckenberg and Rubinsztein-Dunlop 1995: *Physical Review Letters* 75, 826 (1995). *How far it was read: not read (record located). Cited in: light card.*
- M. E. Henderson, B. Heacock, M. Bleuel, D. G. Cory and co-workers, “Three-dimensional neutron far-field tomography of a bulk skyrmion lattice”, *Nature Physics* 19, 1617–1623 (2023). DOI 10.1038/s41567-023-02175-4. *How far it was read: through a summary. Cited in: skyrmion card.*
- J. T. Hoeksema and co-workers, “The Helioseismic and Magnetic Imager (HMI) Vector Magnetic Field Pipeline: Overview and Performance”, *Solar Physics* (2014). DOI 10.1007/s11207-014-0516-8; arXiv 1404.1881. *How far it was read: abstract; read in part (v1, for the fd10 date). Cited in: Part Four, the Sun; the Declaration’s data line.*
- D. R. Hofstadter’s butterfly, the pattern of gaps in a lattice under a magnetic field, credited in the Forest companion. *How far it was read: as cited by another (the Forest companion). Cited in: READ 4.*
- Hong, Ding, Li and Carlsson 2018: *Astrophysical Journal Letters* 857, L2 (2018). *How far it was read: as cited by another (Lu). Cited in: Part Four, the Sun.*
- C. Jacobs and co-workers (2006), a simulated sheared arcade’s ejection threshold. *How far it was read: as cited by another (the review). Cited in: plasma card.*
- A. Jafari, on the rate of magnetic topology change in turbulence (2024). *How far it was read: abstract (the prior-art ledger, S2). Cited in: Who came before.*
- M. Janvier, G. Aulanier, V. Bommier, B. Schmieder, P. Démoulin and E. Pariat, “Electric current in flare ribbons: observations and 3D standard model”, *Astrophysical Journal* 788, 60 (2014). DOI 10.1088/0004-637X/788/1/60; arXiv 1402.2010. *How far it was read: abstract. Cited in: Part Four, the Sun.*
- the Jellyfish paper. *How far it was read: read in full. Cited in: W5.*

- B. Johnstone, G. Petrie and J. Sudol, "Abrupt Longitudinal Magnetic Field Changes and Ultraviolet Emissions Accompanying Solar Flares", *Astrophysical Journal* 760, 29 (2012). DOI 10.1088/0004-637X/760/1/29; arXiv 1211.2212. *How far it was read: abstract. Cited in: Part Four, the Sun.*
- M. C. H. de Jong and co-workers, "Controlling magnetic skyrmion nucleation with Ga⁺ ion irradiation", *Physical Review B* 107, 094429 (2023); arXiv 2303.15283. *How far it was read: abstract. Cited in: skyrmion card.*
- D. G. Joshi and A. Bhattacharyay, "Revisiting Langer-Ambegaokar-McCumber-Halperin theory of resistive transitions in one-dimensional superconductors with exact solutions", arXiv 1106.1720 (2011). *How far it was read: read in full. Cited in: superconducting card; WRITE W3.*
- Kern and co-workers 2025; Liu and co-workers 2025: skyrmion bags written at room temperature and converted by a field. *How far it was read: abstract in part; through a summary (the prior-art ledger). Cited in: skyrmion card.*
- B. Khesin, a survey stating Arnold's inequality (2005). *How far it was read: through a summary. Cited in: page one; plasma card.*
- Kim and Kori 2025: a study of phase-lagged rings steered to a chosen winding (2025). *How far it was read: through a summary (the prior-art ledger, G5). Cited in: Who came before.*
- J. R. Kirtley, C. C. Tsuei and F. Tafuri, "Thermally activated spontaneous fluxoid formation in superconducting thin film rings", *Physical Review Letters* 90, 257001 (2003); arXiv cond-mat/0302384 (record checked 25 September). *How far it was read: read in full (arXiv). Cited in: W4.*
- J. R. Kirtley and F. Tafuri, "Can superconducting rings provide clues to the early development of the universe?", viewpoint on the 2009 niobium ring, *Physics* 2, 92 (2009), doi 10.1103/Physics.2.92. *How far it was read: read in full (Bear's upload of 25 September). Cited in: W4.*
- D. Kleckner and W. T. M. Irvine, "Creation and dynamics of knotted vortices", *Nature Physics* 9, 253–258 (2013). *How far it was read: abstract. Cited in: page one; WRITE W1; water card; WRITE W1; W1.*
- K. Klenin and J. Langowski, on computing linking and writhe from curves drawn as points, as the Latch paper cites them. *How far it was read: as cited by another (the Latch paper). Cited in: READ 1.*
- B. Kliem and co-workers (2011), a simulated ejection carrying a minor part of the helicity. *How far it was read: as cited by another (the review). Cited in: plasma card.*
- I. Kontogiannis and co-workers, on net current as a flare predictor from SHARP maps, *Solar Physics* (2017). DOI 10.1007/s11207-017-1185-1. *How far it was read: abstract (the journal's page); abstract. Cited in: Part Four, the Sun (rims); Part Four, the Sun.*
- B. J. LaBonte, M. K. Georgoulis and D. M. Rust, "Survey of Magnetic Helicity Injection in Regions Producing X-Class Flares", *Astrophysical Journal* 671, 955 (2007). *How far it was read: as cited by another (the review). Cited in: plasma card.*
- R. Landauer, "Irreversibility and heat generation in the computing process", *IBM Journal of Research and Development* 5, 183–191 (1961); for the bound only. *How far it was read: read in part. Cited in: W5.*
- J. S. Langer and V. Ambegaokar, *Physical Review* 164, 498 (1967). *How far it was read: as cited by another (Joshi and Bhattacharyay); as cited; read in full. Cited in: superconducting card; WRITE W3; W3.*
- Lee, Sun and Kazachenko 2021: "... on bald patches at the inversion line in a flare, arXiv 2111.00336; NSF public-access record. *How far it was read: abstract. Cited in: Part Four, the Sun.*
- K. D. Leka, G. Barnes, A. D. Crouch and co-workers, "Resolving the 180° Ambiguity in Solar Vector Magnetic Field Data: Evaluating the Effects of Noise, Spatial Resolution, and Method Assumptions", *Solar Physics* 260, 83 (2009). *How far it was read: as cited by another (Lu); title seen. Cited in: Part Four, the Sun.*
- length contraction: standard special relativity. *How far it was read: textbook, not sourced to one book. Cited in: READ 6.*
- E.-K. Lim and J. Chae (2009), sigmoids and the hemispheric rule. *How far it was read: as cited by another (the review). Cited in: plasma card.*
- C. Liu, N. Deng, R. Liu, J. Lee, T. Wiegmann, J. Jing, Y. Xu and S. Wang, "Rapid Changes of Photospheric Magnetic Field after Tether-cutting Reconnection and Magnetic Implosion", arXiv 1112.3598 (2012). *How far it was read: abstract. Cited in: Part Four, the Sun.*
- Y. Liu, X. Sun, T. Török, V. S. Titov and J. E. Leake, "Electric-current neutralization, magnetic shear, and eruptive activity in solar active regions", *Astrophysical Journal Letters* 846, L6 (2017). DOI 10.3847/2041-8213/aa861e; arXiv 1708.04411. *How far it was read: abstract, and the four regions through a page-summary tool. Cited in: Part Four, the Sun (rims).*

- Liu and co-workers 2025; the MRAM studies: as the prior-art ledger names them. *How far it was read: through a summary. Cited in: skyrmion card.*
- D. W. Longcope, "Topological methods for the analysis of solar magnetic fields", *Living Reviews in Solar Physics* 2, 7 (2005). *How far it was read: abstract (a page-summary tool). Cited in: Part Four, the Sun.*
- Lu, W. Cao, G. Jin, Y. Zhang, M. Ding and Y. Guo, "A Statistical Study of the Magnetic Imprints of X-Class Solar Flares", *Astrophysical Journal* (2019). DOI 10.3847/1538-4357/ab16d4; arXiv 1803.08310 (from the ledger). arXiv 1803.08310; DOI 10.3847/1538-4357/ab16d4. *How far it was read: read in part; read in part (introduction, method, the table rows for regions 11166 and 11158). Cited in: plasma card; Part Four; Part Four, the Sun (the flare, the imprint, the zero check).*
- P. MacNeice and co-workers (2004), helicity carried off by a simulated ejection. *How far it was read: as cited by another (the review). Cited in: plasma card.*
- Martínez Oliveros, Lindsey, Hudson and Buitrago Casas 2014: "Transient artifacts in a flare observed by the Helioseismic and Magnetic Imager on the Solar Dynamics Observatory", *Solar Physics* (2014). *How far it was read: abstract. Cited in: Part Four, the Sun.*
- R. Menzel and M. Gellert, "Regulation of the genes for E. coli DNA gyrase: homeostatic control of DNA supercoiling", *Cell* 34, 105–113 (1983). *How far it was read: abstract. Cited in: W8.*
- the measured relativistic advance of Mercury's perihelion. *How far it was read: not read (recalled). Cited in: READ 8.*
- T. R. Metcalf, "Resolving the 180-degree ambiguity in vector magnetic field measurements: the 'minimum' energy solution", *Solar Physics* 155, 235 (1994). *How far it was read: abstract (NASA's technical reports record). Cited in: Part Four, the Sun (the half-turn).*
- C. W. Misner, K. S. Thorne and J. A. Wheeler, *Gravitation*, W. H. Freeman (1973), for the first-order geodetic precession and perihelion advance. *How far it was read: not read (textbook; named from memory). Cited in: READ 8.*
- H. K. Moffatt, "The degree of knottedness of tangled vortex lines", *Journal of Fluid Mechanics* 35, 117–129 (1969). *How far it was read: as cited by another (Berger and Field). Cited in: page one; the plasma card; page one; water card.*
- H. K. Moffatt and R. L. Ricca, "Helicity and the Călugăreanu invariant", *Proceedings of the Royal Society A* 439, 411–429 (1992). *How far it was read: read in full (23 September); read in full. Cited in: page one; READ 2; the plasma card; WRITE W6; plasma card; water card; W1, W6; Part Four, plasma.*
- R. Monaco, J. Mygind and R. J. Rivers, "Zurek-Kibble domain structures: the dynamics of spontaneous vortex formation in annular Josephson tunnel junctions", arXiv cond-mat/0112321. *How far it was read: read in full (arXiv). Cited in: W4.*
- R. Monaco, J. Mygind and R. J. Rivers, the arXiv cond-mat/0112321 paper of section E as published: *Physical Review Letters* 89, 080603 (2002). *How far it was read: read in full (arXiv). Cited in: WRITE W4's table.*
- R. Monaco, J. Mygind, R. J. Rivers and V. P. Koshelets, "Spontaneous fluxoid formation in superconducting loops", arXiv 0907.2505; listed by a web search as *Physical Review B* 80, 180501 (2009). *How far it was read: read in full (arXiv). Cited in: W4.*
- E. Montoya and co-workers, on MRAM under gamma irradiation, *Scientific Reports* 10, 10220 (2020). *How far it was read: through a summary (the prior-art ledger, P3). Cited in: Who came before; skyrmion column.*
- S. Mühlbauer, B. Binz, F. Jonietz, C. Pfleiderer, A. Rosch, A. Neubauer, R. Georgii and P. Böni, "Skyrmion Lattice in a Chiral Magnet", *Science* 323, 915–919 (2009). DOI 10.1126/science.1166767; arXiv 0902.1968. *How far it was read: read in full. Cited in: skyrmion card; Part Four.*
- NASA planetary fact sheet; IAU constants: the semi-major axes and eccentricities of the planets and the Moon; GM of the Sun and the Earth. *How far it was read: not read (recalled into the tool, marked so). Cited in: READ 8 (the tool).*
- NIST 2013: NIST, "A Primary Voltage Standard for the Whole World" (2013), a web article on the Josephson voltage standard. cited in READ as "NIST, read". *How far it was read: read in full. Cited in: READ 3; WRITE W8; W8; READ 3.*
- L. Néel (1949), the relaxation time of a single-domain grain. *How far it was read: as cited by another (Wikipedia, "Superparamagnetism"). Cited in: WRITE W4.*

- Paper A: Bear's own skyrmion paper, his record; corrections gathered in LATCH_ATLAS_PAPER_A_CORRECTIONS_2026-09-28. *How far it was read: read in full. Cited in: skyrmion card.*
- E. Pariat, P. Démoulin and M. A. Berger, "Photospheric flux density of magnetic helicity", *Astronomy & Astrophysics* 439, 1191 (2005). *How far it was read: as cited by another (the review). Cited in: plasma card.*
- A. A. Pevtsov and R. C. Canfield (1999), sigmoids in a larger sample. *How far it was read: as cited by another (the review). Cited in: plasma card.*
- A. A. Pevtsov, M. A. Berger, A. Nindos, A. A. Norton and L. van Driel-Gesztelyi, "Magnetic Helicity, Tilt, and Twist", *Space Science Reviews* 186, 285–324 (2014). DOI 10.1007/s11214-014-0082-2 (from the ledger). *How far it was read: read in full. Cited in: plasma card.*
- Puckett and co-workers: the ring of light: its paper and Methods. *How far it was read: read in part ("text read"). Cited in: light card; WRITE W7.*
- E. M. Purcell, "Life at low Reynolds number", *American Journal of Physics* 45(1), 3–11 (1977). *How far it was read: not read (publisher page located; last page from the ADS record). Cited in: page one.*
- B. Raphaldini, C. B. Prior and D. MacTaggart, on winding flux rising before flares (2022); abstract from the University of Glasgow's record. *How far it was read: through a summary. Cited in: Part Four, the Sun.*
- Dale Rolfsen, *Knots and Links*, Publish or Perish (1976; AMS Chelsea reprint 2003): chapter 5, section D, "Linking numbers", p. 132 onward. *How far it was read: read in part (section 5.D as printed, 29 September). Cited in: READ 6.*
- N. Romming, C. Hanneken, M. Menzel, J. E. Bickel, B. Wolter, K. von Bergmann, A. Kubetzka and R. Wiesendanger, "Writing and Deleting Single Magnetic Skyrmions", *Science* 341, 636–639 (2013). DOI 10.1126/science.1240573. *How far it was read: read in full (the authors' version). Cited in: skyrmion card; Part Four.*
- M. N. Rosenbluth and M. N. Bussac, "MHD stability of spheromak", *Nuclear Fusion* 19, 489 (1979). *How far it was read: as cited by another (Taylor 1986). Cited in: plasma card.*
- D. M. Rust and A. Kumar, "Evidence for Helically Kinked Magnetic Flux Ropes in Solar Eruptions", *Astrophysical Journal Letters* 464, L199 (1996). *How far it was read: as cited by another (the review). Cited in: plasma card.*
- C. Ryu and co-workers, "Observation of persistent flow of a Bose-Einstein condensate in a toroidal trap", *Physical Review Letters* 99, 260401 (2007); arXiv 0709.0012. *How far it was read: abstract. Cited in: ring-superfluid card; light card; WRITE W7.*
- P. G. Saffman, *Vortex Dynamics*, Cambridge University Press (1992), the ring-energy formula as the Jellyfish paper §7C uses it. *How far it was read: as cited by another (the Jellyfish paper). Cited in: water card.*
- M. W. Scheeler, W. M. van Rees, H. Kedia, D. Kleckner and W. T. M. Irvine, "Complete measurement of helicity and its dynamics in vortex tubes", *Science* 357, 487–491 (2017). *How far it was read: through a summary. Cited in: water card.*
- D. R. Scherer, C. N. Weiler, T. W. Neely and B. P. Anderson, "Vortex formation by merging of multiple trapped Bose-Einstein condensates", *Physical Review Letters* 98, 110402 (2007); arXiv cond-mat/0610187 (record and abstract checked 25 September). D. R. Scherer, "Vortex formation by merging and interference of multiple trapped Bose-Einstein condensates", PhD dissertation, University of Arizona, College of Optical Sciences (2007); arXiv 0704.2054. *How far it was read: read in full (arXiv); abstract. Cited in: W4.*
- G. K. Schroeder, C. Lad, P. Wyman, N. H. Williams and R. Wolfenden, "The time required for water attack at the phosphorus atom of simple phosphodiester and of DNA", *PNAS* 103, 4052–4055 (2006). *How far it was read: not read (record located). Cited in: DNA card (the wall).*
- SDO/HMI SHARP maps: NASA's Solar Dynamics Observatory, HMI's SHARP series, served by Stanford's JSOC; regions 11166 (HARP 401) and 11158 (HARP 377); courtesy of NASA/SDO and the AIA, EVE, and HMI science teams; marked CC0. HARP 401 (region 11166) and HARP 377 (region 11158), fetched with tools/sun_fetch_sharp.py on Bear's laptop; every ledger heads with the files' combined sha256. *How far it was read: data. Cited in: plasma card; Part Four; Part Four.*
- N. Seehafer, "On the magnetic field line topology in solar active regions", *Solar Physics* 105, 223 (1986), the bald-patch test. *How far it was read: as cited by another (the flare literature; Lu and co-workers). Cited in: Part Four, the Sun.*
- Shen (thesis): a doctoral thesis on frequency combs, lines 688–696 and 1202 of its text. *How far it was read: read in part (found by search). Cited in: light card.*

- de Sitter 1916; lunar laser ranging: the geodetic precession of the Earth-Moon system, 19.2 mas/yr, as lunar laser ranging confirms it. *How far it was read: not read (recalled)*. Cited in: *READ 8*.
- SJX, “In-Depth: Understanding Isochronism and Oscillators”, watchesbysjx.com (December 2025). *How far it was read: read in full (a web article)*. Cited in: *W8*.
- I. I. Smalyukh and co-workers (2010), defects written in a liquid crystal by a laser. *How far it was read: as cited by another (Alexander and co-workers 2012)*. Cited in: *liquid-crystal card*.
- J. L. Snoep and co-workers, *European Journal of Biochemistry* (2002), on supercoiling homeostasis. *How far it was read: title only*. Cited in: *W8*.
- L. Sorriso-Valvo, V. Carbone, V. Abramenko, V. Yurchyshyn, A. Noullez, H. Politano, A. Pouquet and P. Veltri, “Topological changes of the photospheric magnetic field inside active regions: a prelude to flares”, arXiv astro-ph/0207244. *How far it was read: abstract*. Cited in: *Part Four, the Sun*.
- Stanford’s page on HMI’s field data; astropy (TAI minus UT): the 1,215-second averaging window; 34 seconds. *How far it was read: through a summary; computed*. Cited in: *Part Four, the Sun*.
- T. R. Strick and co-workers, “The elasticity of a single supercoiled DNA molecule”, *Science* 271, 1835 (1996). *How far it was read: not read (record checked)*. Cited in: *W6*.
- D. A. Wiley, S. H. Strogatz and M. Girvan, “The size of the sync basin”, *Chaos* 16, 015103 (2006). *How far it was read: as cited by another (Groisman and co-workers 2025)*. Cited in: *Who came before*.
- J. Sudol and J. W. Harvey, “Longitudinal Magnetic Field Changes Accompanying Solar Flares”, *Astrophysical Journal* 635, 647 (2005). *How far it was read: title only (retold by Lu)*. Cited in: *the prior-art ledger, not the book’s text (Johnstone, Petrie and Sudol 2012 is what the book cites)*.
- X. Sun, J. T. Hoeksema, Y. Liu, M. Kazachenko and R. Chen, “Investigating the Magnetic Imprints of Major Solar Eruptions with SDO/HMI High-Cadence Vector Magnetograms”, *Astrophysical Journal* (2017). DOI 10.3847/1538-4357/aa69c1; arXiv 1702.07338. *How far it was read: abstract*. Cited in: *Part Four, the Sun*.
- X. Sun and M. C. M. Cheung, on net current from geometry alone, arXiv 2011.11873. *How far it was read: through a summary*. Cited in: *Part Four, the Sun (rims)*.
- Suzuki and Measday: the source of the muon capture fractions, as the Field Guide’s crate names it. *How far it was read: not read*. Cited in: *READ 8*.
- Tait’s knot tables from the vortex atom: secondary source, an encyclopedia entry “Vortex theory of the atom”. *How far it was read: through a summary*. Cited in: *page one*.
- the comments on Tamburini: arXiv 1210.5365 and 1302.4121. *How far it was read: not read (records located)*. Cited in: *radio card*.
- F. Tamburini and co-workers, “Encoding many channels on the same frequency through radio vorticity”, *New Journal of Physics* 14, 033001 (2012); IOP’s press release read, the paper’s page having refused the fetch tool. the Venice twisted-radio test. *How far it was read: through a summary (press-release level); through a summary*. Cited in: *page one; the radio column; page one; radio card*.
- J. B. Taylor’s conjecture, *Physical Review Letters* 33, 1139 (1974). *How far it was read: as cited by another (Berger and Field; Taylor 1986)*. Cited in: *plasma card*.
- J. B. Taylor, “Relaxation and magnetic reconnection in plasmas”, *Reviews of Modern Physics* 58, 741 (1986). *How far it was read: read in full*. Cited in: *plasma card*.
- F. Tecker, “Longitudinal Beam Dynamics”, CERN Accelerator School lecture, arXiv 1601.04901 (2016); CERN Document Server record 1982417, pp. 1–21. *How far it was read: not read (record seen in listings on 30 September: title, author, arXiv number and pages)*. Cited in: *page one*.
- J. K. Thalmann and co-workers (2014), a force-free model showing a half-turn increase. *How far it was read: as cited by another (the review)*. Cited in: *plasma card*.
- the electron’s wave closing on itself: textbook quantum mechanics. *How far it was read: textbook*. Cited in: *page one*.
- the length of the year: 365.2422 days. *How far it was read: textbook, not sourced to one book*. Cited in: *READ 3*.
- the sampling rule: at least two samples per turn, the engineers’ rule (Nyquist and Shannon). *How far it was read: textbook, not sourced to one book*. Cited in: *READ 5*.
- the square-root law: averaging independent noise improves a reading by the square root of the number of samples. *How far it was read: textbook (statistics)*. Cited in: *READ 5*.

- B. Thidé and co-workers, *Physical Review Letters* 99, 087701 (2007). *How far it was read: not read (record located). Cited in: radio card.*
- W. Thomson (Lord Kelvin), "On Vortex Atoms", *Proceedings of the Royal Society of Edinburgh* 6, 94–105; DOI 10.1017/S0370164600045430 (the record gives the volume year as 1869; the paper is commonly dated 1867). *How far it was read: not read (checked against the publisher record). Cited in: page one.*
- V. S. Titov, E. R. Priest and P. Démoulin, "Conditions for the appearance of 'bald patches' at the solar surface", *Astronomy & Astrophysics* 276, 564 (1993). *How far it was read: as cited by another. Cited in: Part Four, the Sun.*
- U. Tkalec, M. Ravnik, S. Čopar, S. Žumer and I. Mušević, "Reconfigurable knots and links in chiral nematic colloids", *Science* 333, 62–65 (2011). DOI 10.1126/science.1205705 (from the ledger). *How far it was read: read in full. Cited in: liquid-crystal card; Part Four.*
- A. Tomita and R. Y. Chiao, "Observation of Berry's Topological Phase by Use of an Optical Fiber", *Physical Review Letters* 57, 937–940 (1986). *How far it was read: read in full (25 September, by OCR of a printed copy, figures checked by eye; journal record checked 25 September). Cited in: opening; page one (the four-words table); READ 2.*
- Wikipedia, "Type I topoisomerase"; "DNA supercoil": read 25 September 2026. *How far it was read: through a summary. Cited in: W6, W8.*
- K. Tziotziou, M. K. Georgoulis and N.-E. Raouafi, "The Magnetic Energy-Helicity Diagram of Solar Active Regions", *Astrophysical Journal Letters* 759, L4 (2012). *How far it was read: as cited by another (the review's figures, read from their axes). Cited in: plasma card.*
- T. Török and co-workers (2014), the writhe of an erupting filament. *How far it was read: as cited by another (the review). Cited in: plasma card.*
- P. Vemareddy and T. Wiegmann, "Quasi-Static 3D-Magnetic Field Evolution in Solar Active Region NOAA 11166 Associated with X1.5 Flare", *Astrophysical Journal* 792, 40 (2014). DOI 10.1088/0004-637X/792/1/40; arXiv 1406.7823. *How far it was read: abstract. Cited in: Part Four, the Sun.*
- P. Venkatakrishnan and S. K. Tiwari, "On the Absence of Photospheric Net Currents in Vector Magnetograms of Sunspots Obtained From Hinode (SOT/SP)", *Astrophysical Journal Letters* 706, L114 (2009). DOI 10.1088/0004-637X/706/1/L114; arXiv 0910.3751. *How far it was read: abstract. Cited in: plasma card; Part Four; Part Four, the Sun (rims).*
- K. Volke-Sepúlveda and co-workers, the twisted-sound experiment (2008). *How far it was read: abstract; through a summary. Cited in: page one; the sound column; page one; sound card.*
- H. Wang (1992) and H. Wang and co-workers (1994), the first reports of permanent shear increase at the inversion line in flares, as Lu and co-workers retell them. *How far it was read: as cited by another (Lu and co-workers 2019). Cited in: Part Four, the Sun.*
- J. C. Wang's review of DNA topoisomerases (2002), cited by a reference summary for the type I enzymes' steps. *How far it was read: as cited by another (a reference summary); the review not read. Cited in: W1.*
- S. Wang, C. Liu, R. Liu, N. Deng, Y. Liu and H. Wang, "First Flare-related Rapid Change of Photospheric Magnetic Field Observed by Solar Dynamics Observatory", arXiv 1103.0027 (2011); *Astrophysical Journal Letters* 745, L17 (2012) expected. *How far it was read: abstract. Cited in: Part Four, the Sun.*
- Wang and co-workers 2023: a Möbius-strip ring cavity with a programmable Berry phase, *Nature Photonics* (2023). *How far it was read: through a summary (the prior-art ledger, R3). Cited in: Who came before.*
- C. N. Weiler and co-workers, "Spontaneous vortices in the formation of Bose-Einstein condensates", *Nature* 455, 948 (2008). *How far it was read: read in full. Cited in: W4.*
- D. J. Weir and R. J. Rivers, "Fluxoid formation: size effects and non-equilibrium universality", arXiv 1105.2349. *How far it was read: read in full (arXiv). Cited in: W4.*
- D. J. Weir, R. Monaco and R. J. Rivers, "Defect formation in superconducting rings: external fields and finite-size effects", arXiv 1208.3426. *How far it was read: read in full (arXiv). Cited in: W4.*
- D. J. Weir, R. Monaco, V. P. Koshelets, J. Mygind and R. J. Rivers, "Gaussianity revisited: exploring the Kibble-Zurek mechanism with superconducting rings", arXiv 1302.7296. *How far it was read: read in full (arXiv). Cited in: W4.*
- B. T. Welsch and D. W. Longcope, "Magnetic helicity injection by horizontal flows in the quiet Sun", *Astrophysical Journal* 588, 620 (2003). *How far it was read: as cited by another (the review). Cited in: plasma card.*

- E. A. West and M. J. Hagyard (1983), Faraday rotation in a sunspot measured with the Marshall Space Flight Center’s vector magnetograph. *How far it was read: as cited by another (the Sun spoke). Cited in: the Sun spoke’s text, not this book’s (carried here from the prior-art ledger; drop at the alphabetical build unless the “Who came before” list wants it).*
- K. Wiesenfeld, P. Colet and S. H. Strogatz, “Synchronization transitions in a disordered Josephson series array”, *Physical Review Letters* 76, 404–407 (1996). *How far it was read: not read (record checked; the abstract not reachable). Cited in: W8.*
- T. Williams, C. B. Prior and D. MacTaggart, “Investigating the Efficacy of Topologically Derived Time-Series for Flare Forecasting. I. Dataset Preparation”, arXiv 2412.04335 (2024). *How far it was read: abstract. Cited in: Part Four, the Sun (forecasting); Who came before.*
- T. Williams, C. B. Prior, D. MacTaggart and D. S. Bloomfield, paper II of the same series, the forecaster, arXiv 2512.14840 (2025). *How far it was read: abstract. Cited in: Part Four, the Sun (forecasting); Who came before.*
- A. N. Wright and M. A. Berger, “The effect of reconnection upon the linkage and interior structure of magnetic flux tubes”, *Journal of Geophysical Research* 94, 1295 (1989). *How far it was read: as cited by another (the review). Cited in: plasma card.*
- D. Xiao, C. B. Prior and A. R. Yeates, “Computation of Winding-Based Magnetic Helicity and Magnetic Winding Density for SHARP Magnetograms in Spherical Coordinates”, *Solar Physics* 298, 116 (2023). DOI 10.1007/s11207-023-02211-9. *How far it was read: read in full; read in full (sittings 123 and 124). Cited in: plasma card; Part Four; Part Four, the Sun.*
- X. Z. Yu, Y. Onose, N. Kanazawa, J. H. Park, J. H. Han, Y. Matsui, N. Nagaosa and Y. Tokura, “Real-space observation of a two-dimensional skyrmion crystal”, *Nature* 465, 901–904 (2010). DOI 10.1038/nature09124. *How far it was read: through a summary. Cited in: skyrmion card.*
- Roger Zelazny, *The Chronicles of Amber*, ten novels, 1970 to 1991, beginning with *Nine Princes in Amber* (1970); Corwin’s five, 1970 to 1978. *How far it was read: read in part (Bear: the first five, from memory); the Pattern, the veils and the Baltimore seed from Wikipedia’s pages The Chronicles of Amber* and The Pattern and the Logrus, read 29 September 2026.* Cited in: opening; acknowledgements.*
- S. L. Zhang, G. van der Laan and T. Hesjedal, “Direct experimental determination of the topological winding number of skyrmions in Cu₂OSeO₃”, *Nature Communications* 8, 14619 (2017). DOI 10.1038/ncomms14619. *How far it was read: through a summary. Cited in: skyrmion card.*
- H. Zhao, J.-X. Wang, J. Zhang and C.-J. Xiao, null points in three-dimensional coronal models found by a topological index, *IAU Symposium* 233, 369–372 (2006). *How far it was read: abstract (a page-summary tool). Cited in: Part Four, the Sun.*
- Zhao and co-workers 2019: a study of heavy-ion effects on MRAM (2019). *How far it was read: through a summary (the prior-art ledger, P3). Cited in: Who came before.*

Who came before

The book’s rule is to find and acknowledge the people whose work it stands on or stood beside, whether or not a section cites them. This list is the book’s prior-art pass, medium by medium: what the book saw, and who saw it first or from another side. Finding that others came first was the usual outcome, and it counts in the book’s favour, because a result that reproduces theirs shows that the book’s method reads correctly. “Not found” means only that it was not found in the abstracts, summaries and listings read, which are at the depths the reference list states; it is a door to show a specialist, never a claim that something is new.

medium	what the book saw	who came before
magnetic skyrmions	how high a real skyrmion’s wall is	in the film where single whirls were first written and deleted, the barrier is tens of meV, so a year’s hold needs about 19 K (Hagemeister and co-workers 2015); the whirls themselves, Mühlbauer and co-workers 2009 and Romming and co-workers 2013
magnetic skyrmions	what circular dichroism reads	in polarised resonant X-ray scattering it reads the winding number itself: for odd numbers the signal’s periodicity equals the number, and even numbers give no signal (Zhang, van der Laan and Hesjedal 2017)

magnetic skyrmions	whether radiation leaves the count alone	ion beams are already a pen, lowering the cost of making skyrmions and doubling how many form (de Jong and co-workers 2023); the nearest record for a stored count is MRAM, whose state survives ionizing radiation while heavy ions cause hard errors by damage (Montoya and co-workers 2020; Zhao and co-workers 2019); single strikes on stored skyrmions not found
magnetic skyrmions	bags, a large whirl holding smaller ones of the opposite sign	named by Foster and co-workers 2019; held at room temperature in multilayer disks, imaged and converted by a field (Liu and co-workers 2025); written on purpose (Kern and co-workers 2025)
magnetic skyrmions	the lattice's higher-order spots, and the whirls seen directly	the spots come from interference and mark long-range order (a 2011 <i>Physical Review Letters</i> paper); the whirls imaged in a thin film (Yu and co-workers 2010) and in the bulk by neutron tomography, with monopoles where strings are cut (Henderson and co-workers 2023)
the power grid	a whole number round a loop of lines, and how many a ring can hold	loop currents between phase-locked states "take only discrete values, as they are characterized by topological winding numbers", with a bound on how many stable states an n-node cycle holds that the book's run sits under (Delabays, Coletta and Jacquod 2016; Coletta, Delabays, Adagideli and Jacquod 2016); extended to cell complexes (Bačić and co-workers 2025)
the power grid	a phase shift on one line as the knob that clears or chooses a winding	phase-shifting transformers in service against loop flows at the borders of Germany, Poland and the Czech Republic (ENTSO-E); their use to choose a winding, not only to steer flow, not found
the power grid	whether a synchronism check can see a whole turn across an open breaker	the relay practice compares voltage, slip and angle inside a small window, five degrees either way in one generator-synchronising guide, so a whole turn plus a few degrees looks the same; the whole-turn case not found, a door for a protection engineer
the power grid	which phase a customer sits on	an established topic, phase identification from smart-meter data (titles only)
the power grid	the basins of twisted states on a ring, and a lag that leans them	basin sizes falling as the exponential of the winding squared, conjectured by Wiley, Strogatz and Girvan 2006 and supported by Groisman, De Vita, Fernández Bonder and Zhang 2025; a phase lag that broadens higher windings' basins and can steer a ring to a chosen one (Kim and Kori 2025)
plasma and the Sun	helicity counted in turns per flux	magnetic winding, helicity without the flux weighting, as an ideal invariant that "measures the entanglement of magnetic field lines" (Prior and MacTaggart); field line helicity for a per-line view; winding-based helicity and winding density computed on the SHARP maps themselves from single maps (Xiao, Prior and Yeates 2023)
plasma and the Sun	linking in a turbulent plasma	topology changes at a rate independent of small-scale plasma effects (Jafari 2024); on the Johns Hopkins turbulence set, field lines are "spontaneously stochastic" and flux freezing breaks down (Eyink and co-workers 2013); the count is lost and the total helicity kept (Taylor 1986; Berger and Field 1984)
plasma and the Sun	the energy floor under linked field lines	Arnold's inequality and Freedman and He's knot bound (1991): the floor stands under the magnetic energy itself
water	linking and helicity measured in real vortex tubes	twist, writhe and linking measured in water, the total holding as writhe and twist trade while viscosity takes the twist (Scheeler, van Rees, Kedia, Kleckner and Irvine 2017); Moffatt and Ricca 1992 for the theorem, Kleckner and Irvine 2013 for the knotted vortex
the Sun, read from its maps	whole turns of the sideways field round loops on a magnetogram	zeros counted by a topological index round closed surfaces (Zhao, Wang, Zhang and Xiao 2006); the Sun's field topology reviewed (Longcope 2005); the SHARP series itself (Bobra and co-workers 2014)
the Sun, through a flare	a count held by a strong rim, and the flare where no rim holds	the flare's "magnetic imprint", a rise in the sideways field along the inversion line seen in difference maps (Wang Shuo and co-workers 2011; Liu Chang and co-workers 2012; Lu and co-workers 2019, who give the size for this very flare); field inversions unreliable during flares (Hong and co-workers 2018, as Lu cites it; Martínez Oliveros and co-workers 2014; Ambastha and Maurya 2012); a surface label that changes at a flare, the bald patch (Lee, Sun and Kazachenko 2021); winding and helicity fluxes hours before flares (Raphaldini, Prior and MacTaggart 2022; Williams, Prior and MacTaggart 2024; Williams, Prior, MacTaggart and Bloomfield 2025; ARTop)

the Sun, the imprint and the polarity line	picking pixels by a two-map rise, zones of common noise, currents read round a rim, shear along the line	step-like permanent changes read from each pixel's history (Johnstone, Petrie and Sudol 2012; Sun and co-workers 2017; Castellanos Durán and co-workers 2018), grouped by the shape of that history (Alanazi and Fletcher 2026); Ampère's law round circles about sunspots (Venkatakrishnan and Tiwari 2009) and currents doubling in flare ribbons (Janvier and co-workers 2014); the field parallel to the line strengthening and the tilt collapsing in six flares (Petrie 2012); the bald-patch test (Seehafer 1986; Titov and co-workers 1993); not found: floors set per zone of common noise on a quiet pair, and rim currents used to flag where a flare changed the current
the Sun, the instrument's half-turn	a flip of the pipeline's direction choice moving a count	the ambiguity and the minimum-energy method that resolves it (Metcalf 1994; Leka and co-workers 2009); temporally consistent disambiguation for time series (a 2018 conference abstract, title only)
the Sun, rims round a spot	whether a spot's current comes back	spots carry no net current (Venkatakrishnan and Tiwari 2009); only strong polarity lines carry a significant one, and isolated regions stay balanced (Georgoulis, Titov and Mikić 2012); eruptive regions, 11158 among them, are far from neutralised (Liu and co-workers 2017); net current as a flare forecaster from SHARP maps (Kontogiannis and co-workers 2017); geometry alone can make a net current (Sun and Cheung 2020); not found: square loops grown round each spot as a cheap test
the Sun, every spot's current flowing up	an offset of one sign in both polarities	Faraday rotation turning the measured direction in sunspot umbrae (West and Hagyard 1983), corrected in a filter magnetograph where it made false currents (Gao, Su, Xu and Zhang 2008); HMI's inversion and its treatment of the magneto-optical effects (Borrero and co-workers; Centeno and co-workers 2014; Hoeksema and co-workers 2014); the companion paper <i>Dancing on the Sun</i> carries the runs
liquid crystals	counting defects from polariser images	a learned detector that finds $\pm 1/2$ defects, their winding numbers and orientations (Ren and co-workers); the classification of a flat layer and the bulk (Alexander, Chen, Matsumoto and Kamien 2012); knots and links tied with laser tweezers (Tkalec and co-workers 2011)
liquid crystals	the energy of a defect line per unit length	twist-disclination line tension measured in twisted cells (Chen and co-workers 2024)
liquid crystals and light	a winding written into glass by laser	a radially polarised optical vortex converter made by femtosecond laser nanostructuring of glass (Beresna and co-workers 2011)
radio and light	a ring of N antennas writes a twist only below $N/2$	"the number of elements should be greater than twice the topological charge", a standard design rule (Fayyaz and co-workers 2026); the Venice twisted-radio test and its critics (Tamburini and co-workers 2012); the twist of light (Allen, Beijersbergen, Spreeuw and Woerdman 1992; Allen and Padgett 2011)
radio and light	a mirror flips a beam's twist	a specular reflector conjugates handedness, so a mirror alone proves "heard", not who (the author's earlier note on the conjugate handshake); an outside source not yet read
radio and light	Möbius strips in light	optical polarisation Möbius strips observed (2015); a Möbius-strip ring cavity with a Berry phase set by its cross-section (Wang and co-workers 2023)
the S rotor	a shape that shows a sign, and stalls at a quarter turn	a single-stage Savonius rotor has negative static torque at some angles, and a quarter-turn twist gives positive torque all the way round (Farozan and co-workers 2025); Savonius rotors as ocean-current meters (report titles)
the chiral handshake	saying hello with a matching twist, from zero trust	phase conjugation and retrodirection: the Van Atta array (1959) and Fink's time-reversal mirrors; the book adds only that the mirror reads, and that the owner's mark is a tilt a mirror cannot make

Left to search, and said so: whether gels read linking numbers; why gyrase steps by two; topological Lamb waves in stratified fluids; topological waves in Hall magnetohydrodynamics.

Appendix: the price table, in full

The nine media of the first build, cell for cell as the price table holds them, with their run names and source tags, for anyone who wants every cell; then the four media added since, row by row from their columns. Part Two’s cards tell the same cells in plain prose.

Water, the lattice ring and the bare filament

	Water (vortex rings and linked vortex lines)	The lattice ring (the magnetic forest)	The bare filament (the Latch paper)
LATCH — what whole number is held?	The linking of vortex lines, carried in the flow as helicity (Moffatt 1969, SOURCED). It is a whole number while the lines stay whole, and only a reconnection changes it; Kleckner and Irvine tied and untied knots in water (2013, SOURCED at abstract level). A single ring carries no whole number at all, because circulation in ordinary water is a magnitude. Blind spot: the fluid gives linking only, never the richer knot invariants (Field Guide row 147).	Each electron state around the ring has a whole-number winding, and the current repeats every whole flux quantum. The ring settles into whichever winding is cheapest for the flux applied and keeps no memory of how it got there. That makes it a label in this model; it becomes a latch in a superconducting ring, where a phase slip is the cut. The companion itself says “same latch, opposite dial” (its line 299); by page one’s third question it is a label here, and the companion’s wording is flagged for its next version.	The Gauss linking number of two closed curves, read to within 0.00001 at 800 points (run LATCH_PROBE 20260923_191121, MEASURED). It is exact as long as the curves do not cross.
DIAL — how does the store scale with size?	At held circulation the energy rises with radius: $E = \frac{1}{2}\rho\Gamma^2R[\ln(8R/a) - 7/4]$ (Saffman, as used in the Jellyfish paper §7C, SOURCED). In words: the energy grows in step with the ring’s size, times a slowly growing logarithm of how thin its core is.	The current falls as one over the circumference: log-log slope -1.0004 across a sixteenfold range, amplitude times N held at 1.6733 (run RING_LATTICE_READER 20260914_032117, MEASURED). Fence in the cell: that holds with temperature scaled to the level spacing; at fixed temperature the slope is -7.96 , and the zero-temperature energy swing goes as -2 (companion, A4). The dial depends on what is held fixed.	At fixed current the energy grows with size, so shrinking always lowers it. The pair linked six times falls from 17.51 to 6.81 at half size with its linking unchanged (run LATCH_FLOOR_CONFOUND_CHECK 20260923_194934, MEASURED). There is no floor to scale.
STORE — how much, in real units?	A bench-scale ring carries about 0.38 to 0.40 mJ, by two estimates (0.377 and 0.403 mJ, jelly_ring_signature run of 9 September, MEASURED as model output).	Amplitude times N is 1.6733 in the model’s own units. No joules until a material is chosen. HELD: needs the superconductor chapter’s material.	None held against the integer.

LEAK — where and how fast does it drain?	Viscosity. A vortex parked against a wall keeps about 99 percent of its circulation over half a second at water's viscosity, and loses 29 percent at ten times water's (Jellyfish rung 3, runs of 9 September, MEASURED). A spin written onto a ring keeps 39 to 40 percent of itself after nine radii in water, sign held (rung 4, MEASURED). The helicity itself decays at finite viscosity: slowly decaying, never topologically protected (Field Guide row 147). Measured in water: Scheeler, van Rees, Kedia, Kleckner and Irvine measured every part of the helicity, twist, writhe and linking, in thin-core vortex tubes; as a helical vortex is stretched or compressed, "the total helicity remains roughly constant as its two components vary, indicating compensation of the dramatic change in writhe by the production of twist", and an isolated helical vortex settles to "a constant helicity state after the dissipation of twist" (Science 357, 487–491, 2017, read through a page-summary tool from the authors' copy, SOURCED; prior-art pass W1). So viscosity takes the twist; the rest holds, and the total shuffles between writhe and twist, as Călugăreanu's rule says.	None in the model, which has no dissipation. HELD: a real ring leaks by phase slips (superconductor) or loses its phase memory (normal metal); needs the superconductor chapter.	Nothing is stored, so nothing leaks.
WALL — what must be climbed to change the whole number?	A cut, not a climb over heat. The linking changes only when two vortex lines are brought within about a core's width of each other and viscosity reconnects them. The jiggle of water molecules cannot carry a whole vortex line there, so against heat the wall is effectively complete; against a deliberate collision it is low. HELD: the reconnection timescale, from Kleckner and Irvine.	None. As the applied flux turns, the ring passes from one winding to the next with no barrier, always sitting in whichever is cheapest; that is why it is a label.	None. Sliding one ring out through the other's wire, the energy drifts smoothly by 3.80 percent along the whole path and shows no spike where the strands pass through each other: its departures from a smooth curve at the two crossings, 0.00314 and 0.00243, are smaller than the largest elsewhere, 0.00611 (run WALL_ESTIMATES 20260923_224526, MEASURED). This energy has neither a floor nor a wall.
SLOSH — what share of the energy comes back instead of being spent?	The free-swimming bell runs at a power factor of -0.083 , a reactive share of 0.9965, at 1 mm (Jellyfish, runs of 17–18 September, MEASURED): ninety-nine percent of what the bell does to the water comes straight back. Held on a leash, the actuator's power factor is $+0.80$ to $+0.86$.	Not measured. HELD: needs a run with a time-varying flux.	Not applicable; nothing moves.
VERDICT (the Field Guide's vocabulary)	A latch, slowly leaking. The whole number is exact while the lines stay whole; the floor is settled physics for an ideal fluid. Where to point a specialist: applied to vorticity, Arnold's inequality bounds the strength of the swirl (the enstrophy), not directly the kinetic energy, so this row must say which energy its floor is under.	A label in this model. Settled physics: a normal ring's persistent current is its ground state.	A label. The count is exact and the floor is absent (Latch paper, version 2).

STRIKE TEST — the shared mechanism, or strike the row	Shared with the lattice ring: a closed loop forces a whole number (linking here, winding there). The row is struck on its energy column if the floor turns out to be on enstrophy alone and the table still prices kinetic energy.	Shared mechanism: the phase winds around a closed loop. The dial runs the opposite way from water's (Field Guide row 532); that is the guardrail's point, not a reason to strike.	The row shares the mechanism with a tube of fixed flux and volume only once the curves are given thickness. As a bare filament it stays in the table as the book's worked example of a count that holds nothing.
FOUR QUESTIONS (loop? whole number? holds energy? wall above the thermal jiggle?)	yes · yes · yes, slowly · yes: undoing the link needs a reconnection, a large-scale event far above room-temperature jiggle	yes · yes · no, in this model · not reached	yes · yes · no · not reached

The superconducting ring and the ring-shaped superfluid

	The superconducting ring	The ring-shaped superfluid (atoms; helium noted)
LATCH — what whole number is held?	The winding of the superconducting phase around the ring. The flux through the hole is held at a whole number of flux quanta, $\Phi_0 = h/2e = 2.067833848 \times 10^{-15}$ Wb, exact in SI units. Measured in 1961 (Deaver and Fairbank, Phys. Rev. Lett. 7, 43, SOURCED at the level of the record; Doll and Näbauer the same year, record not checked).	The winding of the gas's quantum phase around the ring. Circulation comes in whole units of h/m : 1.7357×10^{-8} m ² /s for sodium atoms, 9.9693×10^{-8} m ² /s for helium-4 (run PRICE_TABLE_ESTIMATES). The count is read by releasing the gas and measuring the hole at its centre (Eckel et al. 2014, Nature, SOURCED at abstract level).
DIAL — how does the store scale with size?	With the whole number held and no field applied, the energy is $n^2\Phi_0^2/2L$. In words: it grows as the square of the whole number and falls as the ring's inductance grows, so a bigger ring stores less per rung. Doubling the example ring from 1 cm to 2 cm radius drops the one-quantum energy from 0.345 to 0.141 meV. Same direction as the lattice ring, opposite to water at held circulation.	The energy per atom is $n^2\hbar^2/2mR^2$. In words: the square of the whole number divided by the square of the radius, so a bigger ring holds each rung more cheaply. The flow speed nh/mR falls as one over the radius: 0.142 mm/s for one unit in the Eckel ring.
STORE — how much, in real units?	Tiny per rung. One flux quantum in a 1 cm ring of 0.5 mm wire stores 5.5×10^{-23} J, 0.345 meV, which is 0.954 of the thermal jiggle at liquid-helium temperature (run PRICE_TABLE_ESTIMATES, geometry ASSUMED). Large stores hold very large whole numbers: the CMS detector's solenoid at CERN was designed for a maximum stored energy of 2.6 GJ (CERN-2004-008, SOURCED).	Tinier still. The whole Eckel ring (4×10^5 sodium atoms, mean radius 19.5 μ m) turning at one unit carries 1.5×10^{-28} J, 111 times the thermal jiggle at an ASSUMED 100 nK, but only 2.8×10^{-4} of it per atom. Nobody stores energy this way; this ring stores a count.
LEAK — where and how fast does it drain?	By phase slips in a thin wire, and in a thick one by a vortex crossing the wire (see WALL). In a thick ring the current does not measurably decay (File and Mills 1963, Phys. Rev. Lett. 10, 93: a niobium-zirconium solenoid whose decay time "must be greater than 100,000 years", as reported by Hudson, NASA SP-140, 1970, and "time constants of up to 10 ⁵ years" in Schmüser's lectures, CERN-2004-008; SOURCED through two secondary sources, the original not read). Magnetic-resonance magnets run in the same closed mode lose typically 3×10^{-9} of their field per hour (CERN-2004-008, citing a private communication). In very thin wires phase slips become measurable. Warming the ring above its transition erases the latch entirely.	The flow persisted without decay for up to 10 s, "limited only by experimental factors such as drift and trap lifetime", not by the latch (Ryu et al. 2007, Phys. Rev. Lett. 99, 260401, SOURCED at abstract level). A weak link lowers the wall on purpose, and then the winding changes one unit at a time (Eckel et al. 2014).

WALL — what must be climbed to change the whole number?

In a wire thinner than about one coherence length, the winding changes by a phase slip: the superconducting order vanishes across the whole cross-section, over about one coherence length. The barrier is $(8\sqrt{2}/3) \times$ the condensation energy per volume $\times$ the cross-section $\times$ the coherence length; in words, the energy it takes to melt the order in one slice of the wire, one coherence length thick (Langer and Ambegaokar 1967, as printed by Joshi and Bhattacharyay 2011, arXiv 1106.1720, read in full; they give the theory for samples whose “width is

smaller than the coherence length ξ ”, near the transition; the 1967 paper itself is not read). For a niobium wire of 10 nm radius that is about 4.5 eV, some 1.3×10^4 thermal jiggles at 4.2 K (run WALL_ESTIMATES 20260923_224526; niobium’s critical field, 199 mT, and coherence length, 39 nm, SOURCED from CERN-2004-008; an order of magnitude, because the run uses zero-temperature constants where the formula wants values at the working temperature). A thick wire is outside the formula. The example ring’s wire, 0.5 mm in radius against a 39 nm coherence length, would need some 3×10^{13} jiggles for a slip through its whole cross-section, but a thick wire changes its count by a vortex crossing it, a lower door. HELD: that door’s height (Langer and Fisher 1967 for the three-dimensional case, as named in the Tinkham festschrift; Kogan’s barrier for a strip, as used by Kirtley, Tsuei and Tafuri). A thick ring’s wall is shown to be high by its measured decay (LEAK), not by the formula.

To change the winding, a vortex must cross the ring’s band of gas. The weak link is a deliberately lowered stretch of that wall, and Eckel et al.’s hysteresis loops measure where it gives way. HELD: the wall’s height in energy, to be read from that paper.

SLOSH — what share of the energy comes back instead of being spent?

A persistent current is steady, so nothing sloshes. Driven with alternating fields, a superconductor is the purest slosh in the table. The cleanest measured case is a hollow niobium cavity rather than a wire ring: a residual surface resistance of 3 nΩ corresponds to a quality factor of 10^{11} (Schmüser, CERN-2004-008, Fig. 45, SOURCED), so at an ASSUMED 1.3 GHz only 6.3×10^{-11} of the stored energy is lost per cycle and the ringing lasts 12.2 s, about 16 billion oscillations. Temperature sets the purity: with niobium’s surface resistance of about 15 nΩ at 2 K and 800 nΩ at 4.2 K (same source), the quality factor falls to about 1.7×10^{10} and 3.7×10^8 , ringing for 2.0 s and 46 ms (run SC_PHASE_ESTIMATES 20260924_093814; quality factor taken as one over surface resistance, textbook, to check). A wire ring driven by alternating flux loses through the same surface resistance; its own figure is not priced. It cannot turn as an induction motor’s rotor: with no resistance its current runs a quarter cycle out of step and the average torque is zero (textbook, to check). With a whole number of quanta trapped it is a permanent magnet, and a magnet can run as a synchronous rotor, the lock of the rotor rows.

Steady flow; nothing sloshes. HELD for a driven ring.

WRITE — how is a rung written?	Cool the ring through its transition with flux applied and it traps the whole number nearest the applied flux, keeping it when the field is removed. A deliberately weak spot (a Josephson junction, which makes an rf SQUID) lets applied flux move the ring one rung at a time, with hysteresis. A junction also rings: held at a voltage, it oscillates at exactly 483.5978 GHz per millivolt, and driven by a microwave tone it locks to it, its voltage sitting on whole-number steps, 144.75 μV apart at an ASSUMED 70 GHz (run SC_PHASE_ESTIMATES). NIST's voltage standard locks about 300,000 junctions to one microwave source to give volts, “with magnitudes determined only by the microwave frequency” (NIST news release, April 2013, SOURCED). That is a powered lock: the whole number holds only while the drive is on. Cooled fast with no field, the ring traps a random whole number (the Kibble-Zurek mechanism). In thin-film superconducting rings the random number was seen as a function of quench rate and applied field, and explained by heat carrying vortices across the ring's walls, “complementary to a mechanism proposed by Kibble and Zurek” (Kirtley, Tsuei and Tafuri, arXiv cond-mat/0302384, SOURCED at abstract level; WRITE section W4).	Three pens, all demonstrated. Twisted light: one unit $\hbar$ of orbital angular momentum per atom, transferred from Laguerre-Gaussian photons (Ryu et al. 2007): a label written into matter as a latch. A rotating weak link that stirs the ring up or down one rung, with hysteresis (Eckel et al. 2014). And a fast quench, which leaves random windings: supercurrents created “through a temperature quench of the normal-to-superfluid phase transition”, stochastic, their distribution measured “for different quench times” (Corman et al., Phys. Rev. Lett. 113, 135302, 2014, SOURCED at abstract level; WRITE section W4).
VERDICT	A latch, the strongest in the table: an exact whole number, a wall between rungs far above the thermal jiggle (shown in a thick ring by a current that has not measurably decayed, and priced in a thin wire at about 10^4 jiggles), and a pen we can design.	A latch whose rungs are written by design at the bench. Its lifetime is set by the atoms, not by the latch.
STRIKE TEST	Shared with the lattice ring: the same phase winding around a loop. What the superconductor adds is an order that must be destroyed to change the count, and that destruction is the wall. The row stands.	Shared with the superconducting ring: a phase winding held by an order, with a weak link as the designed wall. Shared with water: circulation, but here it comes in whole units instead of as a magnitude. The row stands.
FOUR QUESTIONS (loop? whole number? holds energy? wall above the thermal jiggle?)	yes · yes · yes · yes: each rung stores about one jiggle at 4.2 K, but the wall between rungs holds: about 10^4 jiggles even in a wire 20 nm across, and in a thick ring a current that has not measurably decayed	yes · yes · yes · yes: each atom stores a few ten-thousandths of a jiggle, and the wall between rungs holds anyway

Closed circular DNA

	Closed circular DNA (plasmids, bacterial chromosomes, loops held shut by proteins)	
LATCH — what whole number is held?	The linking number of the two strands of a closed circle: how many times one strand winds around the other. It is exactly the Gauss linking integral the Latch reader computes, and in Călugăreanu's form it is twist plus writhe, two continuous halves adding to a whole number. The program's own reader shows it on a trefoil ribbon: linking -4.0000 from writhe -3.5123 plus twist -0.4877 (Field Guide DNA rows as digested; source run not re-checked here).	
DIAL — how does the store scale with size?	The free energy of supercoiling grows as the square of the extra linking number and falls with the length of the circle. In words: winding a small circle is expensive and winding a long one cheap, per unit. Held at a fixed density of extra turns, as cells tend to hold it, the cost of the next unit stays about 13 thermal jiggles whatever the length (12.76, 12.98 and 12.65 kT at 2,500, 5,000 and 10,000 base pairs), while the total grows with length (run DNA_ESTIMATES). The dial depends on what is held fixed, as it did for the lattice ring.	
STORE — how much, in real units?	A 5,000 base-pair plasmid at the typical bacterial density of extra turns (-0.06 , UNVERIFIED) is 29 turns short of relaxed and stores about 7.6×10^{-19} J, 4.75 eV, some 185 thermal jiggles at room temperature (run DNA_ESTIMATES; $K = 1,100$, UNVERIFIED). Small per molecule, but it is the first store in the table that sits comfortably above the jiggle at room temperature, and cells spend it: the stored twist helps pry the strands apart where genes are read (textbook, to check).	

LEAK — where and how fast does it drain?	A sealed circle does not leak: nothing short of breaking a strand changes its linking number. A single nick lets the strands swivel, and the stored twist is gone at once; so the leak is the rate of strand breaks, plus the enzymes that cut on purpose.
WALL — what must be climbed to change the whole number?	A covalent bond in the backbone. Heat alone breaks it only extraordinarily slowly (the uncatalysed hydrolysis of DNA's backbone bond is the subject of Schroeder, Wolfenden and co-workers, PNAS 2006; record located, abstract not read, so its half-life stays UNVERIFIED). The wall is complete against the jiggle and open to an enzyme.
SLOSH — what share of the energy comes back instead of being spent?	Not priced. Reading a gene winds DNA ahead of the machine and unwinds it behind, so supercoiling travels in waves along a working chromosome (the twin-domain picture, to check). HELD.
WRITE — how is a rung written?	By enzymes, in exact whole steps. DNA gyrase cuts both strands, passes a segment through and reseals, and "both the introduction and the removal of supertwists by DNA gyrase change the linking number of DNA in steps of two" (Brown and Cozzarelli, Science 206, 1979, SOURCED, text read). Type IA enzymes change it by exactly one, and type IB enzymes, which let the cut strand swivel, by one or more whole turns per cut (checked on 25 September against a reference summary citing Wang 2002; WRITE, W1). Gyrase spends chemical fuel to write rungs downward, a powered pen: each step of two costs about 26 thermal jiggles at the typical density (run DNA_ESTIMATES).
VERDICT	A latch, and the only one in the table that nature writes, reads and spends on purpose. Exact whole number, a covalent wall, a store above the room-temperature jiggle, and a pen that moves in steps of exactly two.
STRIKE TEST	Shared with the filament and the tube: the same Gauss linking number of two closed curves. What DNA adds, and the bare filament lacks, is a body: a double helix with a stiffness that makes extra turns cost energy, and a backbone that forms the wall. The row stands.
FOUR QUESTIONS (loop? whole number? holds energy? wall above the thermal jiggle?)	yes · yes · yes, about 185 jiggles in a typical plasmid · yes: the wall is a covalent bond

Light in a ring

Light in a ring (ring resonators; twisted light noted)

LATCH — what whole number is held?	The mode number: how many whole wavelengths fit around the ring. The light's phase must come back to itself after one lap, so only whole numbers are allowed, and the ring resonates only at those colours. In the Puckett ring, 11.787 mm in radius, about 70,758 wavelengths of 1,570 nm light fit around (run LIGHT_ESTIMATES, index ASSUMED). Twisted light carries a second whole number, the number of times its phase winds around the beam's own axis, with that many units h of orbital angular momentum per photon (Allen, Beijersbergen, Spreeuw and Woerdman, Phys. Rev. A 45, 8185, 1992; as Allen and Padgett put it in their introduction to <i>Twisted Photons</i> (Wiley-VCH, 2011, chapter 1, read in part), "lh describes the orbital angular momentum per photon", and Daria et al. 2011 say the same; the 1992 original is not read).
DIAL — how does the store scale with size?	At a fixed colour, the number of rungs grows in step with the ring's size and the spacing between neighbouring rungs shrinks as one over the radius: 600 rungs 1.32 meV apart at 0.1 mm radius, 6,003 at 1 mm, 70,758 rungs 11.2 µeV apart at the Puckett ring's size. The Puckett ring's measured spacing is 2.720 ± 0.019 GHz at 1,570 nm (text read), against 2.70 GHz from the ASSUMED index. At a fixed whole number, each photon's energy halves every time the radius doubles: 0.789, 0.395 and 0.197 eV at 0.1, 0.2 and 0.4 mm for 600 wavelengths around (run LIGHT_ESTIMATES). Same direction as the superconducting ring and the superfluid: a bigger ring holds each rung more cheaply.
STORE — how much, in real units?	Set by the colour, one photon at a time. A photon of 1,570 nm light carries 0.790 eV, about 30 thermal jiggles at room temperature (run LIGHT_ESTIMATES). How much a ring holds depends on how many photons are pumped in, not on the whole number.
LEAK — where and how fast does it drain?	Out of the ring, at the rate its linewidth shows. The Puckett ring has an intrinsic quality factor of 422 ± 11 million and an intrinsic linewidth of 453 kHz, and a ring-down measurement gives 434 ± 9 million, in agreement (text read). Its photons last 351.7 ns by the quality factor and 351.3 ns by the linewidth: the same number read two ways, which is the LEAK-as-linewidth row made literal. In that time the light oscillates 67 million times (run LIGHT_ESTIMATES). The whole number lasts only as long as the photons do. Rings of crystal hold light longer: quality factors of 63 billion are reported for bulk crystalline resonators (Puckett et al., citing Grudin, Ilchenko and Maleki, Phys. Rev. A 74, 063806, 2006, original not read), and the range across microresonators runs "from 10^6 up to nearly 10^{11} " (Shen's thesis, line 691 of its text).

WALL — what must be climbed to change the whole number?	None. The rungs sit far closer together than the thermal jiggle, 11.2 μeV apart in the Puckett ring, 4.3×10^{-4} of the jiggle at room temperature, and still heat does not move a photon from one to the next: in a ring that is steady in time, light cannot change its colour, so the count is kept by that steadiness, not by a barrier. Heat cannot make these photons either: the mean number it puts in the mode is 5×10^{-14} (run LIGHT_ESTIMATES). That holds for weak light only. Anything that varies in time moves light between rungs with nothing to climb: Puckett et al. calibrate their measurement with an electro-optic phase modulator that puts two sidebands on their laser (Methods, text read). And strong light moves itself: in a high-Q ring, milliwatts of input build up to as much as kilowatts inside, and above a threshold power the ring's response to that intense light "spontaneously generates two sidebands through parametric oscillation", the start of a frequency comb (Shen's thesis, lines 688-696; thresholds as low as 5.4 mW in his rings, line 1202).
SLOSH — what share of the energy comes back instead of being spent?	Nearly all of it. Each cycle the ring's energy swings between its electric and magnetic fields and comes back, and only 1.5×10^{-8} of it is lost per cycle (2π over the quality factor, run LIGHT_ESTIMATES). Seen from the laser at exact resonance, the ring behaves as a plain load, because the slosh stays inside it: a flywheel for light, the regulator of the rotor rows.
WRITE — how is a rung written?	By choosing the colour. A laser tuned to one resonance fills that mode, and retuning it fills the next. The pen chooses the whole number, but the page does not hold it: turn the laser off and the count is gone within the photon lifetime. Twisted light's whole number is written by a hologram: Daria et al. encode it on a spatial light modulator, up to 20 turns of phase around the beam (text read). A Dove prism turns a twist of l into its mirror, $-l$, with nothing to climb, and a rotating one shifts the light's colour in step with the twist; for tightly focused beams the flip smears into a spread of twists centred on $-l$ (González, Molina-Terriza and Torres, Optics Express 14, 9093, 2006, read in part).
PATH TO A LATCH	Yes, into a quantum fluid; no, into an ordinary object. Twisted light has written one unit $\hbar$ per atom into a ring-shaped superfluid, which then held it as a latch for up to 10 s (Ryu et al. 2007, the superfluid column's WRITE row). A year earlier the same kind of photons wrote quantized rotation into the atoms of a condensate: Andersen et al., Phys. Rev. Lett. 97, 170406 (2006), "Quantized rotation of atoms from photons with orbital angular momentum" (abstract read on 25 September; WRITE, W7). The same light makes a trapped particle turn (He, Friese, Heckenberg and Rubinsztein-Dunlop, Phys. Rev. Lett. 75, 826, 1995, record located), and Daria et al. drove 2 μm silica beads along a spiral with it, a particle revolving "with continuous transfer of OAM" (text read). A turning particle keeps a speed, not a whole number. The path to a latch needs a receiver with a phase that must close on itself: a superfluid or a superconductor.
VERDICT	A label: the superconducting ring's twin. The same closure rule, a whole number of waves around the loop, with no order that must be destroyed to change it, so no wall, and a leak priced exactly by its linewidth.
STRIKE TEST	Shared with the superconducting ring and the superfluid: a phase that must come back to itself around a loop winds a whole number of times. What light lacks is an order holding the count, and that absence is what makes it a label. The row stands, as the label twin.
FOUR QUESTIONS (loop? whole number? holds energy? wall above the thermal jiggle?)	yes · yes · yes, briefly: 0.79 eV per photon, for about 350 ns in the best integrated ring · no wall: for weak light the count is kept by the ring's steadiness in time; anything that varies in time, or light strong enough to pass the ring's threshold, changes it freely

Radio with a twist

Radio with a twist (twisted radio beams; rings of antennas)

LATCH — what whole number is held?	The twist: how many times the radio wave's phase winds around the beam's axis in one lap around it, the same whole number twisted light carries. Two twisted radio waves in the 2.4 GHz band were sent 442 m across Venice, "from a lighthouse on San Giorgio Island to a satellite dish on a balcony of Palazzo Ducale" (Tamburini et al., New J. Phys. 14, 033001, 2012, as reported in IOP Publishing's press release, SOURCED at release level; the paper's page refused the fetch tool). The idea for radio was set out in 2007 (Thidé et al., Phys. Rev. Lett. 99, 087701, record located). The Venice paper's claim of extra channels on one frequency drew published comments (arXiv 1210.5365 and 1302.4121, records located, not read).
DIAL — how does the store scale with size?	Flat. Every twist costs the same energy per photon at a given frequency; the twist sets the angular momentum each photon carries, not its energy. The energy per photon falls with the wave's length: 9.93 μeV at 2.4 GHz, 0.414 μeV at 100 MHz (run RADIO_ESTIMATES). A whole number whose rungs all cost the same is a label in its purest form.

STORE — how much, in real units?	None held; the wave carries its energy past. Each photon is tiny beside heat: at 2.4 GHz it is 3.84×10^{-4} of the thermal jiggle at room temperature, and heat alone puts about 2,600 photons into each mode at room temperature, 36 at 4.2 K and 0.003 at 0.02 K (run RADIO_ESTIMATES). So the count lives in the pattern of a strong wave, not in single photons.
LEAK — where and how fast does it drain?	Nothing drains the twist in empty space. The beam spreads as it travels, so the ring needed to read the twist grows with distance (textbook, to check). Heat enters as noise in the read, not as a drain on the count; the loop-count checks of 23 September say how to read through it: average neighbouring samples, then compare two loops (run LOOP_COUNT_NOISE_CHECKS 20260923_234529).
WALL — what must be climbed to change the whole number?	None. A re-phased antenna ring or a twisted reflector changes the count with nothing to climb. Heat does not change a strong beam's twist; it blurs the read.
SLOSH — what share of the energy comes back instead of being spent?	Near an antenna, most of the energy sloshes back and forth in the reactive near field; far away it travels outward as radiation, which is the work done (textbook, to check). An antenna is the device that turns slosh into travel. HELD: a priced figure.
WRITE — how is a rung written?	By a ring of antennas fed with the phase stepped around the ring. With eight antennas the step between neighbours is 45° for a twist of 1, 90° for 2 and 135° for 3; at 180° , a twist of 4, the twist and its mirror cannot be told apart, and a twist of 5, 225° per step, reads back as -3 (run RADIO_ESTIMATES). A ring of N antennas writes and reads whole numbers only below N/2: the same half-turn limit that made the noisy loop walk of 23 September go wrong (run LOOP_COUNT_NOISE_CHECKS, check B).
PATH TO A LATCH	Yes, through a superconductor, as a powered lock: a microwave tone is what locks NIST's 300,000 junctions onto whole-number voltage steps (the superconducting WRITE cell). There the whole number is written by the tone's frequency, not by any twist, and it holds only while the tone plays. Cold opens the way: at 0.02 K heat stops filling a 2.4 GHz mode (0.003 photons, run RADIO_ESTIMATES), which is why superconducting microwave circuits are run that cold (textbook, to check). Whether a twisted microwave beam has written a winding into a superconducting ring or a ring superfluid: HELD, not searched.
VERDICT	A label, the table's noisiest: light's closure rule, a flat dial, no wall, and photons so small beside heat that each mode holds thousands of heat-made ones at room temperature. The count survives in the pattern of a strong wave and is read by averaging.
STRIKE TEST	Shared with light: the same phase winding around the beam's axis, at a different scale. The row stands, beside light.
FOUR QUESTIONS (loop? whole number? holds energy? wall above the thermal jiggle?)	yes · yes · no: it carries energy past and holds none · no wall: heat blurs the read but has no rung to climb

Sound with a twist

Sound with a twist (acoustic vortices; rings of loudspeakers)	
LATCH — what whole number is held?	The twist: how many times the sound's phase winds around the beam's axis in one lap around it. In 2008 it was made in open air at 1,300 Hz, sound 26.4 cm long, with four loudspeakers in a circle for a twist of 1 and eight for a twist of 2 (Volke-Sepúlveda et al. 2008, SOURCED at summary level).
DIAL — how does the store scale with size?	Flat in energy, as for radio: every twist costs the same per quantum. What the twist sets is the turning each joule carries, the twist divided by the angular frequency: 1.22×10^{-4} newton-metre-seconds per joule at 1,300 Hz for a twist of 1, twice that for a twist of 2 (run SOUND_ESTIMATES; the same ratio Allen and Padgett give for light, twist over angular frequency, chapter read in part; tested for sound by Demore et al., Phys. Rev. Lett. 108, 194301, 2012, record located). The count is priced in turning, not in energy.
STORE — how much, in real units?	None held; the wave carries energy and turning past. A quantum of 1,300 Hz sound is 2.08×10^{-10} of the thermal jiggle at room temperature, and heat puts about five billion quanta into each mode (run SOUND_ESTIMATES): sound is wholly classical here, and its count lives in the organized pattern. Heat in matter is itself disordered vibration of the same molecules whose ordered motion is sound (textbook).
LEAK — where and how fast does it drain?	Into its own medium's heat. Viscosity and heat conduction turn ordered sound into disordered jiggle as it travels, faster at higher pitch (textbook, to check). HELD: a priced absorption figure for air and for water.
WALL — what must be climbed to change the whole number?	None. Re-phase the ring of sources and the twist changes.
SLOSH — what share of the energy comes back instead of being spent?	A standing sound wave sloshes between pressure and motion, which is a ringing bell; a travelling one carries energy away (textbook). HELD: a priced figure.

WRITE — how is a rung written?	By a ring of sources stepping the phase from one to the next. The 2008 experiment used a quarter turn between neighbours in both set-ups: four sources for a twist of 1 and eight for a twist of 2, each 90° (run SOUND_ESTIMATES), well clear of the half-turn limit; four sources for a twist of 2 would need 180°, where the twist and its mirror cannot be told apart.
PATH TO A LATCH	None found. In the 2008 test the twist passed into matter as a turn: foam-covered acrylic disks hung on a fibre twisted round to an angle where the sound's torque balanced the fibre (SOURCED at summary level), a torque and an angle, both magnitudes. Sound can also stir a flow in ordinary water, but circulation there is a magnitude too. Whether a sound vortex can write a whole number into a superfluid ring: HELD, not searched.
VERDICT	A label that leaks into the heat of its own medium: light's and radio's closure rule, a flat dial, no wall.
STRIKE TEST	Shared with light and radio: the same winding of phase around an axis, carried by a material medium instead of a field. The row stands.
FOUR QUESTIONS (loop? whole number? holds energy? wall above the thermal jiggle?)	yes · yes · no: it carries energy and turning past and holds none · no wall

Plasma and the Sun

Plasma and the Sun (linked tubes of magnetic flux) (new column, drafted 26 September 2026 from LATCH_ATLAS_DOMAIN_READING_NOTES_2026-09-26 sections 1, 2 and 6 and LATCH_ATLAS_NEW_DOMAIN_RUNS_2026-09-26; the Sun's rows added from section 6, the review of Pevtsov, Berger, Nindos, Norton and van Driel-Gesztelyi 2014, read in full; where that review reports another paper, the row says 'as the review reports it'; the floor under the magnetic energy added on 27 September from the prior-art pass, item S3; the Sun's surface lines added on 26 September, near midnight (Pacific), from the book's own runs on real magnetograms, LATCH_ATLAS_NEW_DOMAIN_RUNS_2026-09-26, and the papers read beside them; not yet in the price table)

LATCH — what whole number is held? The linking of closed tubes of magnetic flux, carried as magnetic helicity (Berger and Field 1984, eq. 2 and the lines after eq. 15, read in full, SOURCED; Moffatt and Ricca 1992, eq. 1.1, SOURCED). Two tubes that carry no helicity of their own, linked m times, carry twice m times the product of their fluxes; the run matched this exactly for linking 0, -1 , -2 and -3 and three pairs of fluxes (FLUX_TUBE_HELICITY_CHECK_20260926_145656, MEASURED; it reproduces Berger and Field's eq. 16 by computer). In a perfectly conducting plasma the linking is kept exactly: "if one closed field line initially links another n times then in a perfectly conducting plasma the two loops must remain linked n times during any subsequent plasma motion" (Taylor 1986, read in full, SOURCED). A tube's own twist is a measure, not a count: a torus twisted T times carries T times its flux squared, and T "can be any real number" (Berger and Field, eqs. 12–13, SOURCED); in a truly spherical spheromak the field lines' winding number falls only from 0.825 on the magnetic axis to 0.72 at the wall (Rosenbluth and Bussac 1979, as Taylor 1986 reports it, SOURCED). Where field lines never close, helicity is an average linking, not a count (Berger and Field, SOURCED). The review of the Sun's helicity restates the count with each tube's own twist, $2L$ times the product of the fluxes plus each tube's twist times its flux squared, the linking number becoming a winding number for tubes stretched between two planes; the Gauss linking integral "gives an integer, as long as the two curves are distinct" (Pevtsov and co-workers 2014, Fig. 1 and section 2.1, SOURCED). On the Sun helicity is quoted in turns, the helicity divided by the flux squared: an active region with flux 5×10^{21} Mx and helicity 1.25×10^{43} Mx² holds half a turn; in the corona twist is usually below one turn, and in magnetic clouds out in space it can reach 10 turns (Pevtsov and co-workers, section 5.3, SOURCED; $H/\Phi^2 = 0.5$ printed by script). So on the Sun the turns are fractions, a measure, not a whole count. On the Sun's surface the book reads a second whole number off real maps of the field (SHARP magnetograms from SDO/HMI, courtesy of NASA/SDO and the AIA, EVE, and HMI science teams): how many whole turns the sideways field makes round a loop, read from the loop's rim alone with READ section 1's walk. On region 11166 every loop's total landed within 5×10^{-15} of a whole number (largest 4.8e-15), and the 5 zeros at sunspot centres (upward or downward field of at least 1,000 gauss) all read +1, as a field fanning out of or into a spot must; the -1 counts are saddles between spots (SUN_WINDING_READ_20260926_194226, MEASURED). This count is a label: read-only, holding nothing back; how far a spot's fan is turned from straight out, its twist, is a measure, and the linking lives in the field above the surface (the book's reading, with Xiao, Prior and Yeates 2023, read in full, SOURCED).

DIAL — how does the store scale with size? The helicity scales with the fluxes, not the size: twice the linking times the product of the two fluxes for a link, and twist plus writhe times the flux squared for one tube (Berger and Field, eqs. 25–26 with eq. 23, SOURCED). Shape does not enter: a linked, twisted pair with fluxes 1 and 2, shrunk, enlarged,

stretched, bent and slid, with lengths from 9.03 to 54.17, read -7.125 every time (FLUX_TUBE_HELICITY_CHECK, MEASURED). Each whole twist written by a cut adds one flux squared; with 16 strands per tube, twists of -2 to $+2$ gave -3.875 , -2.9375 , -2 , -1.0625 and -0.125 on a link carrying -2 , and a bundle of n strands shows $1 - 1/n$ of a twist (0.75, 0.875, 0.9375 and 0.96875 for 4, 8, 16 and 32 strands; MEASURED, as Berger and Field's remark on their eq. 16 says). In real units on the Sun: the active regions of one published sample of 42 hold between about 10^{40} and 10^{44} Mx^2 of helicity (Tziotziou and co-workers 2012, the review's Fig. 8, READ FROM THE FIGURE'S AXES); over a whole 22-year cycle differential rotation sends about 4×10^{46} Mx^2 into each hemisphere, the two of opposite sign (Berger and Ruzmaikin 2000, as the review reports it, SOURCED); the helicity injected into active regions over a cycle has an upper limit of 6.6×10^{45} Mx^2 , and the quiet Sun's is about 10^{43} Mx^2 , a lower limit (Georgoulis and co-workers 2009; Welsch and Longcope 2003; both as the review reports them, SOURCED).

STORE — how much, in real units? Partly priced in real units. The relaxed state is the least magnetic energy the total helicity allows (Taylor 1986, SOURCED): in the Beta-II spheromak the field kept its shape while its energy fell to about an eighth of its starting value (Taylor 1986, as he reports the experiment, SOURCED). On the Sun: the same 42 active regions hold between about 10^{30} and 10^{33} erg of free energy, the energy above the potential field, and free energy and helicity rise together (Tziotziou and co-workers 2012, the review's Fig. 8 and section 5.2; the range READ FROM THE FIGURE'S AXES). An X-class flare needs a peak helicity flow above 6×10^{36} $\text{Mx}^2 \text{ s}^{-1}$ (LaBonte and co-workers 2007, as the review reports it, SOURCED); in a simulation a sheared arcade ejects a twisted tube when the helicity over the flux squared is above 0.2–0.3 (Jacobs and co-workers 2006, as reported, SOURCED); enough helicity “seems a necessary condition for an ejection to occur, but not a sufficient one” (Amari and co-workers 2003b, as reported, SOURCED). On the Sun's surface the upward electric current through a loop is read from its rim alone, by Ampère's law, the current being the circulation of the sideways field round the rim divided by the magnetic constant (Bear's rim; exact to within 5.8×10^{-11} against the sum over the loop's cells, SUN_FLARE_DIFFERENCE 20260927_045554, MEASURED). Grown round a sunspot standing alone, the current rises, peaks and comes back: the current through the largest loop is a median 0.12 and 0.17 of the peak in the two regions read (SUN_RIM_GROWTH 20260927_064149 and 20260927_064153, MEASURED); a spot carries no net current, as Venkatakrishnan and Tiwari 2009 found in twelve spots (abstract, SOURCED); the net current others find sits on strong polarity lines (Georgoulis, Titov and Mikić 2012, abstract, SOURCED). Spots whose larger loops keep taking in more field, in one case a neighbouring spot of the same sign, keep 0.35 to 0.67 of their peak (MEASURED). HELD: whether a spot's return to zero, read through rims where the instrument is seldom confident of the field's direction (0.06 to 0.13 of the rim pixels), survives other ways of resolving that direction (for a specialist). HELD: how much of the free energy the helicity itself guards (the book's reading: at most the free energy, since the least energy the helicity allows lies between the potential field's and the field's own; for a specialist); the energy of a single solar flux rope. The floor stands under the magnetic energy itself: for any field with no sources, which a magnetic field is, Arnold's inequality puts the energy at least a constant times the size of the helicity, the constant set by the shape and size of the container, and Freedman and He's bound puts a knotted tube's energy at least a number set by its flux, its volume and how knotted it is (Khesin 2005, Notices of the AMS, read through a page-summary tool; Freedman and He, Annals of Mathematics 134, 189–229, 1991, record read; SOURCED; prior-art pass S3). In words: a field with helicity cannot relax below an energy set by that helicity and its container, and a knotted tube cannot relax below an energy set by its flux, volume and knottedness.

LEAK — where and how fast does it drain? Through the plasma's resistance. On the long resistive timescale, energy and helicity both decay unless something sustains them (Taylor 1986, SOURCED). At small scales energy drains as the square of the wavenumber and helicity only as the wavenumber, and viscosity drains energy without touching helicity (Taylor 1986, eq. 1.10, SOURCED). On the Sun the helicity is kept so well that it must be thrown out: it “is conserved to an excellent approximation”, the corona's magnetic Reynolds numbers being 10^8 to 10^{15} , and its dissipation is negligible there (Berger 1984, as the review reports it, SOURCED); it “should be removed from the Sun at approximately the same rate as it is created”, or it would choke the dynamo (“dynamo quenching”); coronal mass ejections “are thought to be the primary agent”; cancellation between regions of opposite sign involves little flux (Pevtsov and co-workers 2014, section 5.1, SOURCED). Budgets of helicity in against helicity carried out agree only partly, with large uncertainties (section 4.1, SOURCED). In simulations an ejection carried off at least 80 % of the helicity put in (MacNeice and co-workers 2004) or only a minor part (Kliem and co-workers 2011), both as reported. HELD: Berger's 1984 limits themselves (reported through the review, not read).

WALL — what must be climbed to change the whole number? Not an energy barrier. In a perfectly conducting plasma the field lines move with the gas, so no linking can change; the guard is conductivity. With any resistance, however small, lines break and reconnect, and the rate of reconnection “does not diminish as fast as η and may not diminish at

all” as the resistance η shrinks (Taylor 1986, SOURCED). Reconnection keeps the total helicity approximately while it changes the links (Taylor’s conjecture of 1974, as Berger and Field and Taylor 1986 report it, SOURCED): two linked tubes carrying twice the flux squared between them can reconnect into one tube that carries the same amount inside itself, as twist and kink (Berger and Field, SOURCED). When two equal tubes of flux Φ reconnect once, their mutual helicity changes by one flux squared and “the final tubes will share this helicity equally, both obtaining a half twist”, half a flux squared each (Wright and Berger 1989, as the review reports it, SOURCED; a change of half a turn in their mutual winding, printed by script); a force-free model of one active region shows “an apparent increase of a half turn twist in a flux rope” (Thalmann and co-workers 2014, as reported, SOURCED). So the wall guards the total, not the separate counts: a race between fast relaxation and slow resistive decay. On the Sun’s surface the book’s own rule holds through two X-class flares: a count changes only where its rim goes weak. Over three hours of maps every twelve minutes, 1,589 loops in region 11166 (an X1.5 flare) and 1,352 in region 11158 (an X2.2 flare) kept a readable rim in every map, and none changed its count (SUN_SERIES_READ 20260927_041443 and 20260927_060906, MEASURED). At region 11166’s flare site no loop kept a readable rim for the three hours, and there the zeros come and go: the flare works in the gaps between strong rims (MEASURED). The instrument adds a wall of its own: it cannot tell which way along its line the sideways field points, and where the pipeline flips its choice on a patch, a count moves with no zero inside. The one count that changed across region 11158’s flare did so on a rim touched by such a flip (1 of 1), and 4 of 59 zero sightings in that flare site had no weak field inside (SUN_ZERO_CHECK 20260927_061940, MEASURED); in region 11166, 92 of 103 site sightings passed as true zeros (SUN_ZERO_CHECK 20260927_061935, MEASURED). HELD: whether the zero of the sideways field that sits on region 11166’s polarity inversion line for three hours is the foot of the null point above (Vemareddy and Wiegmann 2014, abstract, SOURCED) or the end of a bald patch (for a solar physicist).

SLOSH — what share of the energy comes back instead of being spent? Not priced.

WRITE — how is the whole number written? A voltage round a plasma torus writes helicity at twice the voltage times the toroidal flux, so at fixed flux the helicity is the volt-seconds stored (Taylor 1986, eq. 1.14, SOURCED); a voltage across the polar caps of a flux-core spheromak does the same (eq. 6.8), and a gun’s voltage held one spheromak far longer than its resistive decay time (Taylor 1986, the CTX experiment, SOURCED): a powered lock. A boundary turning in place pours helicity into the space above it (Berger and Field, section 4, SOURCED as a model on a flat boundary, not an observation of the Sun). Reconnection moves helicity from the link into one tube’s twist (WALL). On the Sun helicity comes up from below, carried through the surface by field lines as they emerge and by motions along the surface that twist and shear them (the review’s eq. 20, SOURCED); in active regions it “typically increases during flux emergence” and stops rising when emergence stops (Pevtsov and co-workers 2003a, as reported, SOURCED); shearing is a less efficient pen than the emergence of a twisted tube (section 4.1, SOURCED). Tubes turning on the surface write at a rate summed over every pair, each pair’s turning rate weighted by the product of their fluxes (eq. 21, SOURCED): over one full turn, two tubes turned about each other write twice the product of their fluxes, and one tube turned on itself writes its flux squared, one step of the count (printed by script). Which way it is written leans by hemisphere, mostly negative in the north and positive in the south, “a relatively weak statistical trend satisfied by 60–75 % of active regions”, from the quiet-Sun network to magnetic clouds; what sets it is not settled (differential rotation, the Coriolis force, the dynamo, turbulent convection are all candidates), and the scatter points to turbulence (section 4.2, SOURCED). What a flare writes on the Sun’s surface is a measure, not a count: the sideways field near the polarity inversion line steps up and stays up, by 239.6 gauss in region 11166 and 287.1 in region 11158, read only in maps the choice of pixels never used (SUN_FLARE_DIFFERENCE 20260927_045554 and 20260927_061515, MEASURED), against 258 ± 73 and 286 ± 59 gauss for the same flares in Lu and co-workers 2019 (read in part, SOURCED); picking pixels by their rise between two maps builds in a rise of about 153 gauss where nothing flared (153.1 and 152.6, MEASURED). Along region 11158’s polarity line the field parallel to the line rose by 154.9 gauss (SUN_PIL_SHEAR 20260927_062644, MEASURED), the change Petrie 2012 found in six flares (abstract, SOURCED).

READ — how is it read? Helicity has a meaning only inside surfaces the field does not cross; an open region such as the Sun’s corona is read as a relative helicity, against the least-energy field with the same field at its boundary (Berger and Field, section 3, SOURCED). Helicity has no location, only a total within a flux surface (Taylor 1986, SOURCED); “it does not have a simple density like mass”, and a flow of helicity per unit of surface gives false patches of both signs: “only helicity flux density per unit of elementary magnetic flux has a physical meaning” (Pariat and co-workers 2005, as the review reports it, SOURCED). On the Sun it is read three ways: the corona’s total, as relative helicity, from model fields built on surface maps; the flow through the surface, from footpoints tracked on successive maps; and proxies from the electric current, of which only the vertical part can be measured (the review, section 3, SOURCED). The sign is read far better than the size: S-shaped soft X-ray loops (sigmoids) read positive and reverse-S

negative, and obey the hemispheric rule in 80–87 % (Rust and Kumar 1996; Lim and Chae 2009) or 64 % in a larger sample (Pevtsov and Canfield 1999); filament barbs in 76–82 %, or 68 % by an automated method (all as the review reports them, SOURCED). Fences: potential fields can look sigmoid, the curve depends on the viewing angle, and two ropes with the same reverse-S shape seen from above can carry writhe $-0.2\Phi^2$ and $+0.2\Phi^2$ (Fig. 4), because writhe “cannot be computed from the projection of the loop on the photosphere alone” (SOURCED); the writhe of an erupting filament gives the sign of its twist and only whether it is more or less than about 6π (Török and co-workers 2014, as reported, SOURCED).

VERDICT (the Field Guide’s vocabulary; drafted, for Bear to settle with a plasma physicist’s confirmation) Split. In a quiet plasma that conducts almost perfectly: a latch, the links kept exactly. In a turbulent plasma: the count is lost and the floor is kept; the separate linkings cannot even be read, and the total helicity, not a whole number, still holds energy back. On the Sun: the total is kept to an excellent approximation and shed in eruptions, and the twist in the corona is usually below one turn, so the Sun’s store is mostly a measure, not a whole count; this fits the split’s second half (Pevtsov and co-workers 2014, SOURCED; the fit is the book’s reading). On the Sun’s surface, the whole turns of the sideways field round loops with strong rims do not move through an X-class flare: a label that maps where the surface’s topology is locked, holding no energy, while the flare’s changes, measures, happen in the gaps (the book’s reading of its runs, MEASURED counts). Where to point a specialist: a plasma physicist, on the split; a solar physicist, on the surface label. The floor under the magnetic energy is settled mathematics (Arnold; Freedman and He 1991).

STRIKE TEST — the shared mechanism, or strike the row Shared with water: linked lines carried as helicity (Moffatt 1969), the same arithmetic with flux in place of circulation. Struck on its count if no regime is found in which closed flux tubes last long enough to be counted.

FOUR QUESTIONS (loop? whole number? holds energy? wall above the thermal jiggle?) yes: closed tubes of magnetic flux · yes in a quiet plasma; in a turbulent one the separate counts cannot be read · yes: the total helicity holds the energy above a least-energy state · no energy wall: the guard is a race of timescales, fast relaxation against slow resistive decay

The power grid

The power grid (a whole number round one loop of lines) (new column, drafted 26 September 2026 from LATCH_ATLAS_DOMAIN_READING_NOTES_2026-09-26 section 3, the power-grid section of LATCH_ATLAS_NEW_DOMAIN_RUNS_2026-09-26, run GRID_LOOP_FLOW_CHECK 20260926_053413, and the register row on Bear’s tilt; Coletta, Delabays, Adagideli and Jacquod, New J. Phys. 18, 103042 (2016), read in full as arXiv 1605.07925v2; the prior-art pass’s items G1 to G3 and G5 folded in on 27 September; not yet in the price table)

LATCH — what whole number is held? The winding of the voltage phase round a loop of the network. Every bus has one voltage, so the angle differences along the lines of a loop, each taken within half a turn, add up to a whole number of turns: “it counts the number of times the complex voltage winds around the origin in the complex plane as one goes around the α th loop” (Coletta and co-workers 2016, eq. 1, SOURCED). Each winding is a steady state with its own flow circulating round the loop, which they call a vortex flow. In the book’s stated model, a ring of N equal lines with nothing injected, no losses and heavy damping, the ring holds every winding with less than a quarter turn per line and no other, checked by stability for every ring from 4 to 24 buses: a ring of 12 holds five windings, -2 to $+2$; a ring of 18 holds nine, -4 to $+4$; a ring of 24 holds eleven, -5 to $+5$ (GRID_LOOP_FLOW_CHECK 20260926_053413, MEASURED). The paper’s 18-bus ring carries “at most nine solutions” (Coletta, line 178, from their Ref. 10, SOURCED): the two agree independently. The count and its limit are in print: loop currents between phase-locked states “take only discrete values, as they are characterized by topological winding numbers”, and an n -node ring holds at most $2 \text{Int}[n/4] + 1$ stable states, that is, twice the whole part of a quarter of the ring’s length, plus one (Delabays, Coletta and Jacquod, Journal of Mathematical Physics 57, 032701, 2016, abstract read, SOURCED). The book’s counts, 5, 9 and 11 for rings of 12, 18 and 24, sit at or under that bound, 7, 9 and 13 (printed by script): equal at 18, where no winding sits exactly on a quarter turn per line, and below at 12 and 24, where one does and the book’s strict rule leaves it out. States with two or more lines past a quarter turn are “all linearly unstable” (the same paper), which is the book’s finding that a loop carries only one line a little past its own quarter turn. The same count on general networks: phase-locked states “are organized by winding numbers associated with generalized independent cycles” (Bačić, Schaub, Kurths and Witthaut, arXiv 2510.05831, abstract read; prior-art pass G2). Fence, in the authors’ own words: “there is no quantumness in the power flow problem” and high-voltage AC grids “are not superconducting” (SOURCED); “vortex” and “quantized” are their analogies.

DIAL — how does the store scale with size? The number of windings a ring can hold grows with its length, about one either way for every four lines (the quarter-turn rule, MEASURED). The flow of each winding falls as the ring grows: winding 1 carries 0.5000 of a line's greatest power in a ring of 12 and 0.2588 in a ring of 24; in a ring of 12, winding 2 carries 0.8660 (MEASURED). In real units, loop flows of up to 1 GW are recorded in North America (Coletta, lines 39–41, citing others, SOURCED as reported), but the usual definition of those flows is “ambiguous” (Coletta, lines 45–46, SOURCED), so they are not shown to be windings.

STORE — how much, in real units? Not a store held back. A winding circulates power that delivers nothing and costs losses. Real grids lose 5–10 % of what they carry (Coletta, lines 88–90, SOURCED); a winding raised the ohmic losses by almost 50 % in their double loop at a line conductance of 3 % of the susceptance, and doubled them in a model with the UK grid's layout at 10 % (Coletta, lines 571–572 and 752–753, SOURCED). So in a real grid the count spends power instead of holding it back. HELD: the losses of a winding in a real grid, in real units.

LEAK — where and how fast does it drain? None shown while the grid stays in step: once written, a winding held in their models with losses included and when operating conditions went back (Coletta, section VI, SOURCED). If the extra supply is cut back to the level with no winding, the winding still holds and the whole grid turns at a shifted frequency instead, which “would be intolerable in a real power grid” and would lead to trips, cascades and possibly blackouts (Coletta, lines 648–658 and 754–757, SOURCED). The count goes when the loop is opened. In the model, random knocks of 0.6 radians on every bus of a ring of 12 changed winding 2 in 0.341 of 1,000 tries, winding 1 in 0.006, and winding 0 in none (MEASURED): the more flow a winding carries, the less margin it has.

WALL — what must be climbed to change the whole number? A saddle of the grid's Lyapunov function, which the authors call an “energy” barrier in quotation marks (Coletta, line 790, SOURCED): a potential that governs the dynamics, not a store of joules. To change a winding smoothly the voltage would have to fall to zero somewhere: “Changing a vortex flow requires to change the number $q\alpha$ of such rotations, i.e. to untwist V_l , which cannot be done smoothly without driving $|V_l| \rightarrow 0$ somewhere” (Coletta, lines 78–80, SOURCED). No thermal jiggle enters. In the model, the steady state stalls when a line passes a quarter turn, and the loop carries one line a little past its own quarter turn: with one line weakened in a ring of 12, the weak line reaches a quarter turn at a strength of 0.8413 for winding 2 but the ring holds to 0.8398, and for winding 1 at 0.4154 and 0.4069 (MEASURED). Operating grids keep the angle across a line to “few tens of degrees”, beyond which the line's thermal limit is exceeded (Coletta, lines 562–565, SOURCED). FIELD (Bear, from equipment, 26 September register): sections are reconnected in phase, because a quarter turn out of phase blows past the thermal load; this agrees with the paper.

SLOSH — what share of the energy comes back instead of being spent? Not priced.

WRITE — how is the whole number written? Three pens (Coletta, SOURCED). (i) Dynamical phase slip: raise the load until the state with no winding loses stability; in their 18-bus ring the state with no winding fails at $P/B0 \approx 1.6$ while winding -1 holds to about 1.85, and turning the power back down leaves the winding in place. In the book's run, weakening one line of a ring of 12 slips the count one whole turn at a time, from 2 to 1 at a strength of 0.835 and from 1 to 0 at 0.405 (MEASURED). But this pen is “of little relevance for power grid operation”, because it needs about a quarter turn across one line, beyond the thermal limit (Coletta, lines 562–565, SOURCED). (ii) Tripping a line in a lopsided double loop can leave a winding of 1, and reclosing the line does not remove it (Coletta, section IV, SOURCED). (iii) Reclosing a line: the loop gains one turn each time the angle across the open gap passes an odd number of half turns (Coletta, eqs. 9–10 and Fig. 4, SOURCED). In their model the angle across the open line reaches about a full turn near $P/B0 \approx 0.5$, and “taken modulo 2π ” it is small, so reclosing is “technically feasible” and writes a winding of 1 (Coletta, lines 545–547 and 583–584, SOURCED). So the switch is the pen: $W1$'s cut, pass and rejoin, and READ section 5's half-turn limit used as a write rule, the rejoined line taking the short way and the loop keeping the rounding (the book's reading). Which way it is written needs a small bias: in a ring with an exact mirror symmetry the symmetry “forbids transitions to other q -values”, and the authors add small random injections to break it (supplement S4, lines 1057–1060, SOURCED). Their recommendation: $N - 1$ security checks should also look for vortex flows (Coletta, lines 766–769, SOURCED); no security check that does was found (prior-art pass G2). The relay at the breaker: a synchronism check compares voltage, slip and angle inside a small window, “5 degrees advance and 5 degrees late” in one generator guide, with slip up to “0.05 Hz” and a voltage mismatch “less than 5%” (a generator-synchronising guide, through a page-summary tool, SOURCED; line reclosing settings differ and were not read); a phase angle is known only within a turn, so a standing difference of a whole turn plus a few degrees passes it (the book's reading; no source discussing the whole-turn case found; prior-art pass G3). The grid's steering knob: a phase-shifting transformer controls power “by regulating the voltage angle and the voltage amplitude between two nodes” and is used against loop flows, for example at the border of Germany, Poland and the Czech Republic (ENTSO-E,

through a page-summary tool, SOURCED): it is a tilt on one line, the book's phase shift in its tilt run; whether any operator has used one to choose or remove a winding was not found (prior-art pass G1). In rings of oscillators, a phase lag broadens the basins of higher windings and can steer a ring to a chosen one (Kim and Kori, arXiv 2502.10681, abstract read; prior-art pass G5). HELD: whether anything in grid protection can see a whole turn across an open breaker (for a protection engineer); whether windings can be written without switching (open in 2016).

READ — how is it read? The loop walk of eq. 1 on the voltage angles round a loop (Coletta, SOURCED); in practice on measured phasors (the book's note; not in the paper, and no read of a winding from measured phasors in a real grid was found: a door for a power-systems engineer, prior-art pass G2). The count is invisible at any single breaker when the angle there has passed a whole turn, because a single angle is known only within a turn; only the walk round the loop reads it (Coletta, lines 545–547, SOURCED; the reading is the book's).

VERDICT (the Field Guide's vocabulary; drafted, for Bear to settle with a power-systems engineer's confirmation) A powered lock. The winding lives only in a grid that is driven and in step, and it must be paid for in losses; its count is protected while the drive runs, and goes when the loop is opened. Where to point a specialist: a power-systems engineer or grid operator, on the powered-lock reading and on whether the recorded loop flows are windings at all.

STRIKE TEST — the shared mechanism, or strike the row Shared with the book's ring of phases: the same sine pull between neighbours, which the authors liken to "the DC Josephson current" between two superconductors (Coletta, lines 64–68, SOURCED, as their analogy), and W8's lock with its quarter-turn stall. Struck as a medium of record if real loop flows prove not to be windings; the model and the paper show the count, and the field has not yet shown it.

FOUR QUESTIONS (loop? whole number? holds energy? wall above the thermal jiggle?) yes: a closed loop of lines and buses · yes: the voltage phase's turns round the loop · no, it spends energy: a circulating flow that raises losses · no thermal jiggle enters: the guard is a saddle of the grid's dynamics, held while the grid stays in step

Liquid crystals

Liquid crystals (defect lines and points in a nematic) (new column, drafted 26 September 2026 from LATCH_ATLAS_DOMAIN_READING_NOTES_2026-09-26 sections 4 and 5, the liquid-crystal section of LATCH_ATLAS_NEW_DOMAIN_RUNS_2026-09-26, and run NEMATIC_HALF_READ_CHECK 20260926_053148; Alexander, Chen, Matsumoto and Kamien, Rev. Mod. Phys. 84, 497 (2012), and Tkalec, Ravník, Čopar, Žumer and Muševič, Science 333, 62 (2011), both read in full; the store's line priced on 27 September from the prior-art pass, item L2; not yet in the price table)

LATCH — what whole number is held? The turn of the molecules' shared direction round a loop. That direction has no head and no tail, "properly a line field, rather than a vector field" (Alexander and co-workers 2012, lines 204–206, SOURCED). In a thin layer the count comes in halves: the strength is the director's turn round a loop divided by a whole turn, "equal to 1/4 times the number of brushes" seen between crossed polarisers (lines 229–237, SOURCED). In the book's run the headless read was right in every case, for seven strengths from $-3/2$ to $+3/2$, loops of 8, 16 and 64 points, and 1,000 random choices of head for each; halves add like whole numbers: a loop round $+1/2$ and $-1/2$ reads 0, round two $+1/2$ reads $+1$, round $+1/2$ and $+1$ reads $+3/2$ (NEMATIC_HALF_READ_CHECK 20260926_053148, MEASURED; it reproduces the colloquium's classification of a thin layer by computer). In the bulk, where the direction can tip out of the plane, the count changes kind: for lines "no distinction can be made any longer between positive and negative strength defects", two half-strength lines together are equivalent to no defect, and a line is either there or not, a parity (lines 749–758, SOURCED). A whole number survives in the bulk at points: a hedgehog's degree, the solid angle the direction sweeps over a sphere round it divided by 4π , "Remarkably, it is an integer" (eq. 8, lines 951–990, SOURCED); but its sign reverses when the hedgehog is carried once round a line defect, so reading it needs the path it took (section IV.A, SOURCED). The island a $+1$ hedgehog makes on a measuring sphere is a skyrmion (lines 838–843, SOURCED). Every charge belongs to a measuring surface: without one, "the 'charge of a defect' is ill defined" (lines 426–429, SOURCED). In the knots of Tkalec and co-workers 2011, each defect line is a three-sided strip, "analogous to the well-known Moebius strip", so "only fractional values of the internal twist are allowed", in thirds of a turn, counted by a self-linking number tied to Berry's phase (Tkalec, lines 334–342 and 507–514, SOURCED); the number of loops and the self-linking number together classify every knot and link on an array (lines 514–525, SOURCED).

DIAL — how does the store scale with size? Not priced in either paper. Energetics "strongly disfavors extra brushes" (Alexander, line 370, SOURCED); loops "tend to be as short as possible in order to reduce the total free energy" (Tkalec, lines 153–168, SOURCED): a line tension. Measured since for twist disclinations in twisted cells, it grows

with the logarithm of the cell's thickness, from 75 pN at 6 μm to 200 pN at 27 μm (Chen, Mandić, Slaughter, Tanaka, Kikkawa, Collings and Yodh, *Soft Matter* 20, 9050 (2024), read through a page-summary tool, SOURCED as reported in a summary): the store grows in step with the length of line, and slowly with the cell.

STORE — how much, in real units? Priced for the lines, not for the texture. A twist-disclination line in a twisted cell has a tension of 75 to 200 pN in cells 6 to 27 μm thick (Chen and co-workers 2024, SOURCED as reported in a summary); their fitted prefactor, 78 ± 11 pN, is several times the theory's. A tension is an energy per length: 7.5×10^{-17} to 2.0×10^{-16} J for each micrometre of line, about 1.85×10^4 to 4.94×10^4 times the thermal jiggle at 293 K (arithmetic printed by script, prior-art pass L2). The cells measured are the same twisted kind in which Tkalec and co-workers tie their knots. HELD: the elastic energy of the bent texture in real units; the numbers above to be confirmed against the paper's text by a soft-matter physicist.

LEAK — where and how fast does it drain? By coarsening: after a quench, defects appear suddenly and meet their opposites, "a coarsening of defects too rapid for us to photograph" (Alexander, lines 1726–1729, SOURCED). A loop round a $+1/2$ and a $-1/2$ together reads 0 (MEASURED), so a pair can vanish without changing the count of a loop round both. For the knots, no leak is reported; they hold between rewirings (Tkalec, SOURCED as not reported).

WALL — what must be climbed to change the whole number? Not priced in either paper. Removing a pair of defects needs "surgery", a cut and a regluing (Alexander, lines 1748–1752, SOURCED): only a cut changes the count.

SLOSH — what share of the energy comes back instead of being spent? Not priced.

WRITE — how is the whole number written? Three pens (Alexander, SOURCED): cooling, which throws down many defects that then coarsen (W4's freeze); a laser (Smalyukh and co-workers 2010, as cited); and shape: a small particle whose surface holds the molecules pointing outward is itself a unit hedgehog, and "this topological charge will have to be compensated by a companion defect in the liquid crystal", a point defect or a "Saturn ring" loop round the particle (lines 1088–1105). Tkalec and co-workers write knots with light: laser tweezers cut and fuse defect lines by "a sequence of local, isotropic-to-nematic, temperature, and optically induced micro-quenches" (lines 322–324), melting the order at one spot and letting it re-form, W1's cut, pass and rejoin done by W4's freeze and thaw at one site (the book's reading). On an array of p by q spheres, 4.72 μm across in a cell 5.5 μm thick, the loops meet at $(p - 1)(q - 1)$ crossing sites, each switchable between exactly three states, "We can make exactly three tangles by reversibly transforming them one into another" (lines 174–184 and 291–305, SOURCED). On a 4 by 4 array the 3^9 combinations give almost 40 different knots and links with up to 10 crossings; 35 % of the conformations are prime knots or links, 29 % unknots, 18 % unlinks and 18 % composite links; knots are assembled to order, the Borromean rings among them, and all knots and links with up to six crossings were made (lines 383–396 and 526–541, and the abstract, SOURCED). Which way it is written: the cell's walls are rubbed at right angles, twisting the director a quarter turn from wall to wall; "Using twisted director structure is essential for the stability of nematic knots and links", and "Negative or left-handed crossings (1) are favored in a left-twisted nematic profile because of the geometric constraint of the cell" (lines 184–187 and 267–268, SOURCED): a twist built into the cell picks the hand.

READ — how is it read? Between crossed polarisers, by counting the dark brushes with their sign: the brushes turn with the polarisers for a positive defect and against them for a negative one (Alexander, lines 229–237, SOURCED); the method is "turning geometric information into counting" (lines 382–383, SOURCED). The book's loop walk needs one change for a headless field: each step taken within a quarter turn either side, not half a turn, because a line turned half a turn is the same line; READ's arrow read, with heads chosen at random, never returns a half and scatters over whole numbers, from -6 to -1 for a $+1/2$ defect on a loop of 16 points (MEASURED). The price is half the margin: with noise on every sample the headless read first goes wrong more often than 1 in 100 at 0.3 radians on every loop size tried, while an arrow read of a true arrow field holds to 0.6 or 0.7 (MEASURED): READ section 5's half-turn limit, halved. An arrow carried round a $+1/2$ defect is turned 0.5000 of a turn after one lap and home after two, the same two-trip pattern as Berry's spinor (MEASURED). Knots are read under the microscope and identified by drawing their crossings (Tkalec, SOURCED). Adding charges needs a shared reference, "why standard time zones were developed for trains" (Alexander, lines 524–529, SOURCED).

VERDICT (the Field Guide's vocabulary; drafted, for Bear to settle with a soft-matter physicist's confirmation) A latch whose price is not yet paid. The count is held by topology and changed only by a cut, "surgery", or by the laser's local melt; in a thin layer it comes in halves, in the bulk only as a parity for lines and a whole number at points, and in the knots in thirds. Neither paper prices the store, the wall or the leak; the store of the lines has since been priced from Chen and co-workers 2024, tens of thousands of jiggles per micrometre, and the wall and the leak are still not. Where to point a specialist: a soft-matter physicist, on the latch reading and on the energy of a defect line per unit length.

STRIKE TEST — the shared mechanism, or strike the row Shared with READ’s loop walk (done with light) and with Berry’s two trips round a half; the hedgehog’s degree is the fifth domain’s skyrmion count in another medium. Struck as a latch if no regime is found in which a defect holds against the thermal jiggle, which neither paper prices.

FOUR QUESTIONS (loop? whole number? holds energy? wall above the thermal jiggle?) yes: a loop round a defect line or a sphere round a point · yes, in halves in a thin layer, a parity for lines in the bulk, a whole-number degree at points, thirds for the knots’ strips · yes: a line tension of 75 to 200 pN, about 1.85×10^4 to 4.94×10^4 jiggles per micrometre of line; the bent texture not priced · not priced: only a cut changes the count, and the height of that wall is in neither paper

Magnetic skyrmions

Magnetic skyrmions (whirls of magnetisation in a thin film or a chiral magnet) (new column, drafted 26 September 2026 from LATCH_ATLAS_DOMAIN_READING_NOTES_2026-09-26 sections 7, 8 and 9, the skyrmion section of LATCH_ATLAS_NEW_DOMAIN_RUNS_2026-09-26, runs SKYRMION_COUNT_CHECK 20260926_192130 and GROUPING_READ_CHECK 20260926_195247, and LATCH_ATLAS_KNOB_CHART_2026-09-26; Mühlbauer, Binz, Jonietz, Pfleiderer, Rosch, Neubauer, Georgii and Böni, Science 323, 915 (2009), read in full as arXiv 0902.1968v1; Romming, Hanneken, Menzel, Bickel, Wolter, von Bergmann, Kubetzka and Wiesendanger, Science 341, 636 (2013), read in full as the authors’ version; Bear’s Paper A, Zenodo concept DOI 10.5281/zenodo.20587337, read in full; the prior-art pass’s items P1 to P4 and K1 folded in on 27 September; not yet in the price table)

LATCH — what whole number is held? The skyrmion number: how many times the arrows of magnetisation cover the sphere of all directions, the solid angle they sweep counted in whole spheres of 4π . It is defined as $(1/4\pi) \int \mathbf{n} \cdot (\partial \mathbf{n} / \partial x \times \partial \mathbf{n} / \partial y) dx dy$ over the film (Mühlbauer and co-workers 2009, eq. 5; Romming and co-workers 2013, the same density, SOURCED). In the book’s run the solid-angle read, adding up triangles of arrows, gives -1.000000000000 exactly for whirls of radius 3, 6 and 12 grid spacings, where the smooth formula taken on the same grid gives -0.669294 , -0.905854 and -0.975672 (SKYRMION_COUNT_CHECK 20260926_192130, Part A, MEASURED): the count is whole only when it is read as a sum of closed pieces. What sets the sign: minus the turning times the core’s direction, in all 32 whirls tried; the handedness, turning every arrow by the same angle, never changes it (Part B, MEASURED). In MnSi the whirls form a lattice whose total “per two-dimensional unit cell is quantized and adds up to -1 ”, while the density inside a cell swings between about $+0.5$ and -2 (Mühlbauer, lines 284–293 and Fig. 3D, SOURCED; the count computed in their model, not measured). The sign is a convention: Mühlbauer computes -1 for whirls whose cores point against the field, Romming gives $+1$ to the same kind of whirl with the same formula (reading notes sections 8 and 9, the book’s reading); Mühlbauer’s “antiskyrmion-lines” means cores against the field, later usage means reversed turning. Paper A’s alphabet: antiskyrmion -1 , blank 0, skyrmion $+1$, and bags of charge $N - 1$, “tunable to ~ 10 ” (as Paper A reports it). Bags were named by Foster and co-workers 2019: “an arbitrary number of antiskyrmions are contained within a larger skyrmion” (Nature Physics 15, 655, abstract read through a page-summary tool, SOURCED). In magnetic multilayer disks at room temperature, a bag with N inner whirls “can have an arbitrary topological charge $Q = N - 1$ ”, and “Individual skyrmion bags with varying topological charges are identified to remain stable at zero field”, held by the disk’s edge (Liu and co-workers, Nature Communications 16, 125, 2025, abstract read through a page-summary tool, SOURCED; prior-art pass P4).

DIAL — how does the store scale with size? With exchange alone, not at all: a smooth whirl’s energy is 19.3322 at any radius, above the floor for a count of one, $4\pi = 12.5664$ (Part C, MEASURED). So exchange alone gives a whirl no preferred size; its size comes from other forces. In MnSi the helix has a period of about 190 Å against an atomic spacing of about 4.56 Å, about 41.7 atoms per turn (Mühlbauer, SOURCED; ratio printed by script). In PdFe on Ir(111) a whirl covers about 270 surface atoms, and the zero-field spiral has a period of 6 to 7 nm (Romming, SOURCED).

STORE — how much, in real units? Not priced in real units. In the book’s model, the whirl’s exchange energy in the model’s units (Part C). In MnSi the lattice is only just the cheapest state: in the plain mean-field model it is metastable, closest to the conical state at $B = 0.40357 Bc2$, where the gap is 0.02838 in the model’s units, and thermal fluctuations make it stable (Mühlbauer, SOURCED).

LEAK — where and how fast does it drain? Temperature decides. In PdFe on Ir(111) at 4.2 K, in the limit of low bias and low current the switching rate “approaches zero”, so thermally activated switching “does not play a role”, and a metastable state “may be preserved”; at 8 K the transitions are reversible, the film following the field (Romming, SOURCED). In MnSi the lattice exists only in a pocket of field and temperature, within about 1.8 to 3.05 K of $T_c = 29.5$ K at 0.162 to 0.190 T in the data shown, and leaves it by first-order transitions; it came out the same whether

cooled in the field or given the field after, so it keeps no memory of its path (Mühlbauer, SOURCED; offsets printed by script). Paper A's clock is Arrhenius, $\tau = \tau_0 \exp(\Delta E/kT)$ with $\tau_0 \approx 10^{-9}$ s: at 293 K a barrier of 0.5 eV lasts 0.398 s and 1.1 eV 264 years, and 1.1 eV lasts 95.5 years at 300 K (arithmetic printed by script); tiers 0.25 eV apart differ by about 2×10^4 in lifetime. Those barriers are Paper A's design assumptions, not measured skyrmion barriers. Radiation: in MRAM, the magnetic memory in service, ionizing radiation leaves the stored state alone, the junctions' key properties "robust against ionizing radiation" after gamma rays to 147 kGy (Montoya and co-workers, Scientific Reports 10, 10220, 2020, through a page-summary tool, SOURCED), and heavy ions cause hard errors only by damage, for tantalum ions above 95.8 MeV (Zhao and co-workers, Microelectronics Reliability 100–101, 113347, 2019, abstract, SOURCED). Single strikes on stored skyrmions: not found (prior-art pass P3).

WALL — what must be climbed to change the whole number? In a smooth magnet, the size of the magnetisation must go to zero somewhere: "it would cost a large energy to suppress M locally to zero, which makes it meaningful to classify magnetic configurations into topological sectors" (Mühlbauer, supplement lines 863–871, SOURCED). On atoms the wall is finite: "in a real system with magnetic moments on an atomic lattice, strict topological protection does not exist. Instead, the states are separated by a finite energy barrier" (Romming, lines 104–109, SOURCED). In the book's model, with exchange alone, there is no barrier at all: a whirl shrinks downhill until the grid takes its count, first slipping at a radius of 1.40 grid spacings (centre on a site or mid-square) or 1.35 (a general point), its energy just before the slip 0.886 to 0.979 of the smooth floor (Part C, MEASURED). So the wall comes from forces the model leaves out. In the PdFe on Ir(111) film where single whirls were first written, the same group reports activation energies of tens of meV between a whirl and the plain film (34 and 21 meV, as the summary gave them), and that "the upper limit for lifetime of one year at Bc corresponds to a temperature of T=19 K", so "a device on the basis of this material will be able to operate at the temperature of liquid He only" (Hagemeister, Romming, von Bergmann, Vedmedenko and Wiesendanger, Nature Communications 6, 8455, 2015, read through a page-summary tool, SOURCED as reported in a summary; prior-art pass P1). At room temperature, bags in multilayer disks stay put at zero field, held by the disk's edge (Liu and co-workers 2025, SOURCED). HELD: the numbers from Hagemeister and co-workers, to check against the paper's text; the wall's height in room-temperature materials.

SLOSH — what share of the energy comes back instead of being spent? Not priced.

WRITE — how is the whole number written? With a microscope tip, one whirl at a time (Romming, SOURCED). A scan at +1 V turns a patch of spiral into whirls; voltage sweeps with the tip held still create whirls "one by one below the tip or in close vicinity"; atomic defects are where they form and stay. Four whirls at four defects were erased one by one with sweeps up to +750 mV and written back in a different order: a four-site register. What writes is the energy of the tip's electrons, not heat: at about 300 mV a switch comes once every 15 s at about 300 nA, and the rate rises about 98 times by 900 mV (6.5 Hz read from their Fig. 4A); the rate grows linearly with current, and at constant power it still depends on voltage, so local heating "can be ruled out". Per electron that is a chance of about 3.6×10^{-14} at 300 mV and about 3.5×10^{-12} at 900 mV (arithmetic printed by script). Which way it is written: the field tilts two wells, the whirl and the uniform state, equal at B_0 , and reversing the current shifts the whole curve by about 100 mT, through spin-transfer torque (Romming Fig. 4C, SOURCED); the pen writes both ways where the tilt is near zero. In MnSi the field and the temperature write the whole lattice. A dose of ions is a pen as well: Ga⁺ irradiation changes the interface forces that set the wall, so that "the energy cost of skyrmion nucleation can be reduced up to 26%" and the number nucleated "could be doubled" (de Jong and co-workers, Physical Review B 107, 094429, 2023, abstract read, SOURCED; prior-art pass P3). At room temperature, bags in multilayer disks are written on purpose (Kern and co-workers 2025, abstract in part), moved one at a time by a magnetic tip, and converted by a field from higher charge to lower (Liu and co-workers 2025, SOURCED): the field is a pen downward, and a leak if unwanted. Paper A proposes a contactless address, a microwave field crossed with a GHz acoustic standing wave (proposed, not built), and writes in $\pm$ pairs, a design choice: single whirls are written alone.

READ — how is it read? The book's solid-angle read, triangle by triangle, is exact on a clean grid (Part A). With noise on every arrow it first reads wrong more often than 1 in 100 at 0.6 radians (Part D: 0.0200 wrong at 0.6, 0.2900 at 0.8, MEASURED). From scattered stations joined into triangles it is right in 0.990 of layouts at a mean spacing of 0.33 of the whirl's radius, 0.725 at 0.47 and 0.335 at 0.67; clean data miss a count and in that run never invented one (Part E, MEASURED). Smoothing by nearness within one grid spacing rescues a noisy read (right in 1.000 at a noise of 1.0, where the raw read is right in 0.350 and invents a count in 0.445); smoothing by likeness, whatever the distance, erases the whirl, its rare core outvoted (right in 0.000 from a noise of 0.6, or 0.8 with a stricter test of likeness; GROUPING_READ_CHECK 20260926_195247, MEASURED). In the laboratory: small-angle neutrons along the field see six spots, the three wave vectors, but not the helices' relative phases, which decide whether there are whirls at

all, so in MnSi the count is inferred from a model (Mühlbauer, SOURCED; the phase problem is the book's reading); a spin-polarised tip sees each whirl directly, and at low voltage switching becomes "increasingly rare, which facilitates nonperturbing imaging" (Romming, SOURCED). Paper A reads the sign by circular dichroism. In polarised resonant X-ray scattering, dichroism reads the winding number itself: "In case of an odd winding number, N equals to the periodicity of the CD signal ... For an even winding number, no CD signal is observed" (Zhang, van der Laan and Hesjedal, Nature Communications 8, 14619, 2017, through a page-summary tool, SOURCED; prior-art pass P2). So that read is blind to even numbers, and bags of charge $N - 1$ are even half the time; an image that counts the whirls misses none (prior-art pass P4). The neutrons' phase problem has been worked around three ways: higher-order spots from interference, which mark the lattice's long-range order (Physical Review Letters 107, 217206, 2011, abstract read); Lorentz electron microscopy of a thin film of Fe_{0.5}Co_{0.5}Si, whirls 90 nm apart in fields of 50 to 70 mT (Yu and co-workers, Nature 465, 901, 2010, through a page-summary tool); and neutron tomography of a bulk lattice in three dimensions, with monopoles where whirl strings are cut (Henderson and co-workers, Nature Physics 19, 1617, 2023, through a page-summary tool; prior-art pass K1). A real-space picture is what the book's loop walk needs; the tomography's monopoles are the only points where a bulk count can change.

VERDICT (the Field Guide's vocabulary; drafted, for Bear to settle with a magnetician's confirmation) A latch, in PdFe on Ir(111) at 4.2 K: a whole number held by a finite wall, written, erased and read one whirl at a time, with a pen and a reader on one knob, the electrons' energy. Its wall there is tens of meV, as reported, so a year's hold needs about 19 K. At room temperature, bags of whirls are held at zero field by the edge of small multilayer disks, the wall's height not priced. In MnSi the lattice is a phase whose cells each carry a whole count while field and temperature hold, not a store.

STRIKE TEST — the shared mechanism, or strike the row Shared with the liquid crystal's hedgehog: the same solid-angle count over a sphere of directions. Shared with the superconducting ring: an order that must be destroyed to change the count, here the size of the magnetisation. The row stands.

FOUR QUESTIONS (loop? whole number? holds energy? wall above the thermal jiggle?) yes: the rim of a whirl, where the arrows are uniform, closes the patch · yes: the skyrmion number, its sign set by the turning and the core, and by which way is called up · yes, in the whirl's exchange and other energies, not priced in real units · yes at 4.2 K in PdFe on Ir(111), where no thermal switching is seen, with a wall of tens of meV as reported (a year's hold needs about 19 K); at room temperature bags are held at zero field by the edge of small disks, the wall's height not priced

Declaration

What this is. An assembled draft of a book, shared as a preprint. It has had no peer review. Its literature pass is bounded: each chapter's end note says how far every source was read, whether in full, in part, at the level of its abstract or through a summary. A reader should take each claim at that level and no further. What is still open, whether not yet priced, read or checked, is listed at the end of Part Two and in each section's "What this section does not show".

How it was made. Usually I start by chatting with Claude about some topic to chase, and after some back and forth the team of Claude and me goes chasing it. I use the work to learn and test ideas, and then look for ways to share the lessons through papers; these projects and papers are always a place for me to store my notes in the form of a preprint. We recently marked over a hundred days of holonomy antics with the Field Guide upgrade, and this was the first project where I asked Claude to co-drive. I know how many people think of Claude in a research assistant's role. I wanted to see what a more supportive research-partner relationship would look like in a project: letting Claude do more of the directing through recommendations, and supporting Claude's ideas as one might a coworker's. Claude selected the topics to chase, and answered why and how after reviewing our preprint adventures. Claude was equally involved in developing the topics and the structure, crafting the scope, re-integrating old published and unpublished papers and transcripts, and adjusting mid-exchange on new information. Claude was an equal partner in every key decision, and I think the book is more complete and more successful for it.

AI assistance. Claude is an AI model made by Anthropic. Claude proposed the book's shape, wrote its prose and its small tools, ran those tools, read the papers named in the end notes, and checked each part against its sources. I supplied the papers and my own deposited work, the field experience credited to me where it appears, and the decisions of scope. Helper instances of Claude drafted eight of the first nine medium cards from the price table; Claude drafted the other five, the superconducting ring's and the four added later; every card was then read in full, and every number on it was checked against its column.

Where the numbers come from. Every number in the book comes either from a named run of one of the book's tools or from a named source. The runs are named in each section's end notes. The tools, their run ledgers and every file of the book are listed with their hashes in the book's manifest. The tools are pure mathematics or stated models unless a section says otherwise, and every assumed value is called assumed where it is used.

Data. The maps of the Sun's magnetic field are from NASA's Solar Dynamics Observatory, its Helioseismic and Magnetic Imager's SHARP series, served by Stanford's Joint Science Operations Center: courtesy of NASA/SDO and the AIA, EVE, and HMI science teams. The series is marked CC0.

Dual use. No dual-use concern was identified. The book measures and models published physics and biology at a descriptive level and gives no procedure for making anything harmful.

Author line. Matthew "Bear" Schulz, written with Claude (Anthropic), as on the title page.