

Dancing on the Sun

A loop count that reads through holonomy, and the current it finds flowing in both polarities

Matthew “Bear” Schulz

Independent researcher · ORCID 0009-0000-0956-5743 · mattschulz97@protonmail.com

Preprint. This paper has not been peer reviewed. Its search of the literature is limited, and where a paper was read only in part, the text says so. It comes from a part-time research team that works in the open and gives its tools away. Every number in it is printed by a tool that ships with it.

What this paper found

The Sun’s surface field has a sideways part, and at every point it points somewhere. Walk a small square loop round a patch of the surface, and add up how far that direction turns as you go. The total is always a whole number of turns: the loop’s **count**. This paper reads counts on the maps that NASA’s Solar Dynamics Observatory makes of the Sun’s field. The aim is to learn what such a count can see, and what the instrument does to it.

1. **A count on a strong rim holds, even through a great flare.** Loops whose whole rim sits in clearly measured field kept their counts through two X-class flares: 1,589 loops round one active region and 1,352 round another, with not one change. Of the 106 among them that counted a turn, 103 never had two neighbouring arrows on their rims more than a quarter turn apart, in any map. (Figure 3.)
2. **Every spot’s current seems to flow up.** Near the centre of the Sun’s disk, the electric current in the strong field of sunspots has a part that points up whether the spot’s field points up or down. The Sun’s own physics treats a field and its mirror image alike, so it has no reason to prefer one sign in both polarities. It was read blind twice, each time with the prediction written before the maps were fetched. Near disk centre, 24 of 26 regions never used in any fit showed it, and then 45 of 50 more. (Figures 4, 5 and 14.)
3. **It behaves like the instrument.** A small turning of the measured sideways field, one way where the field points up and the other where it points down, reproduces it when put into the maps on purpose. It holds steady through the satellite’s daily orbit, and blurring does not make it. Tested blind, it appears in maps made by both versions of the instrument’s inversion code: 17 of 18 regions under the older code and 28 of 32 under the newer. Fitted on the two regions where it was first seen, it is about 5.5 degrees; over all regions near centre it is about 3, and the second blind test put its median at 3.2. It fades toward the edge of the disk without changing sign. (Figures 6 and 14.)
4. **It is made in the instrument’s frame, and the light carries a mirror image of it.** Read two ways at once, about the Sun’s vertical and about our line of sight, the turning at mid-disk is about 4 degrees of the line-of-sight kind and about none of the vertical kind, which is why it looks as if it fades toward the edge when read the Sun’s way. In the instrument’s own polarised light, the direction of the polarisation differs from the reported arrow by an amount that flips with the field’s polarity and changes sign from the bottom of the spectral line to its wings, and a textbook one-layer model of the line makes that pattern. 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 its turning changes sign near 1,000 gauss and the maps’ does not, at any plausible width of the line. And fitting the instrument’s own light with a textbook one-layer model, our arrows differ from the instrument’s by an S-shaped few degrees. Where the turning is made is narrowed to the inversion’s reading of the light, and the rest is for the instrument’s team. (Figures 7 to 10.)
5. **The count reads through that turning.** Put the turning into 38 maps, and 31 of 154,267 readable loop counts change. All 31 sit on five places already in the maps, where two neighbouring arrows point nearly opposite ways. No turning smaller than an eighth of a turn can move a loop whose rim steps are all under a quarter turn. (Figure 11.)
6. **Two counts, with opposite blind spots.** The instrument measures the line the field lies along, and a program chooses which way along the line each arrow points. A second count that follows only the line cannot be moved by that choice. A smooth turning moves it three to four times as often as the arrow count, because its weak point sits at a quarter turn instead of a half. Where it comes out odd, a “half count”, it marks a place radar engineers have called a residue since 1988. The loop that finds such places is the Burgers circuit that Nye and Berry drew round the zeros of waves in 1974. (Figures 2 and 12.)

7. **The current makes a loop that the instrument's confidence cuts in half.** Near disk centre, the current added up over the pixels where the instrument is confident of its arrows comes out upward in 30 of 41 fresh regions, in both hemispheres, and in 67 of 82 more read blind. Added up over each whole region, it does not. Taking out a turning of about 3 degrees shrinks it in 60 of those 82, but leaves about a quarter of it. The pipeline's own published mean current, the SHARP keyword MEANJZD, is added up over the same pixels and carries the same upward part: its sign agrees with this paper's net in all 123 regions, and its value is this paper's to a ratio of 1.000. The published sum of the two polarities' net currents, SAVNCP, carries the turning too, by about a sixth of its size in most regions near centre; the unsigned total does not. (Figures 13 and 14.)
8. **What failed.** The viewing geometry predicted that the turning would change sign beyond about 62 degrees from disk centre. It failed twice: 2 of 7 far regions in the first blind test, and 3 of 6 large far regions in the second, needed a negative turning. The turning fades toward the edge but does not cross. The 5.5-degree turning fitted on the first two regions is too large for the rest: taking it out puts a return current in weak field, where the maps show none. The edge of the instrument's confidence is not where the leftover current is made. A test that was to read the turning's cause directly from the light could not be read at all, because its premise was wrong, and the mis-modelled magneto-optical term that explains the turning's size and polarity does not explain its shape. These are kept, in section 13, as properties of the problem.

1. The pictures you need

The sideways field and its arrows. At each point of the surface the field has an upward part, Br, measured in gauss, and a sideways part lying flat along the surface. Draw the sideways part as an arrow at each pixel (Figure 1, A).

A loop, its rim and its count. A loop is a small square whose sides are 8, 16 or 32 steps from pixel to pixel, so its rim, the ring of pixels round its edge, is walked in 32, 64 or 128 steps. Walk round the rim and, at each step, note how far the arrow turned from one pixel to the next, taking the shorter way, so each step lies between half a turn back and half a turn forward. Add the steps. Because you come back to where you started, the total is a whole number of turns: the count. A count of +1 means the arrows turned once round, the same way as the walk; 0 means they did not turn round at all (Figure 1, B).




Figure 1

Figure 1. *The pictures a reader needs.* (A) The sideways field of NOAA 11166 at 23:24 TAI on 9 March 2011, drawn as arrows at every second pixel over the upward field (red where the field points up, blue where it points down). The orange loop sits round a spot and counts +1; the green loop beside it counts 0. (B) Walking once round each rim, 64 steps, the running total of how far the arrows have turned. The loop round the spot comes home having turned once; the other comes home where it started.

Why the count is a label, not a measure. The count is a holonomy, in the companion book's word: what an arrow carried round a closed path has turned by when it comes home. Squeeze or stretch the field smoothly and the count cannot change, because a whole number cannot creep. It can change only when something crosses the rim: a point where the sideways field vanishes, or a jump in the data. So a count on a rim that stays strong holds. A measure, such as the current, moves with every small change.

A readable pixel. A pixel is readable when three things hold: it lies inside the active region; the instrument's program was confident of which way its arrow points; and the sideways field there is at least three times its own measurement error. A loop is read only when every pixel of its rim is readable, and such a rim is called strong here: its sideways field stands well clear of its noise.

The current. The upward electric current through a patch is how much the sideways field circulates round it. Divide the circulation by a constant of nature and you get the current. Summed pixel by pixel, it gives a current density, in milliamperes per square metre.

Offset and twist. Split the pixels into bands by field strength, and into the two polarities, field up and field down. In each band, take the mean current over the pixels of each sign. Half their difference is the **twist**: it flips sign with the polarity, as a twisted bundle of field does. The word is used here for that part of the current only; it is not the Călugăreanu twist of a ribbon, the geometric twist that the companion book and the Field Guide use beside writhe, and the table at the end of this section keeps the two apart. Half their sum is the **offset**: it has the same sign in both polarities, which the Sun's own physics has no reason to make (Figure 4).

The turning. The instrument may report each sideways arrow turned a little from where it really points. A turning of a few degrees counterclockwise where the field points up, and clockwise where it points down, makes an offset. This paper puts such a turning into the maps on purpose, and takes it out again, to see what it does. The turning grows with the field's strength and reaches its full size at a set field, 250 gauss unless another is named. Section 6 distinguishes two kinds: a turning about the Sun's local vertical with the sign of Br, called kind V, and a turning about our line of sight with the sign of the field along the line of sight, called kind L. At disk centre the two are the same thing.

Viewing angle, μ . A region at the centre of the disk is seen straight on, and one near the edge is seen from the side. The number μ measures this. It is 1 at the centre and 0 at the edge. NEAR means μ at least 0.75, within about 41 degrees of disk centre. FAR means μ at most 0.40, at least about 66 degrees out.

Firm, even-with-steps and half. A loop is **firm** when every step round its rim is under a quarter turn. It is **half** when an odd number of steps are over a quarter turn, and **even-with-steps** when an even number, two or more, are. Section 8 explains why these three classes matter (Figure 2).




Figure 2

Figure 2. Firm, even-with-steps and half. Left: the instrument measures the line the sideways field lies along; turned half a turn, an arrow changes but a line does not, so which way along the line each arrow points is chosen afterwards. Right: one loop of each class, each counting +1, with the 32 steps round its rim. Dashed lines mark a quarter turn, dotted lines half a turn. A firm loop has no step past a quarter turn, and its headless count (section 8) is exactly twice its arrow count. Each step past a quarter turn moves the headless count by one, so one such step makes it odd: a half count. The examples are chosen by a fixed rule, the first loop of side 8 counting +1 in each class (with no more than two steps past a quarter turn), taking the four regions' maps in order; all three happen to lie in NOAA 11158.

Four words that cross a fence. Some of this paper's words are used differently on the two sides of a boundary it crosses, and the companion book asks that such words be pinned before anything else is said. Its own list, in the book's form:

word	on the maps' side	on the light's side, or the literature's
azimuth	the direction of the sideways field in the map's frame, counted from west toward north, made from (Bp, -Bt)	HMI's reported azimuth in the detector's frame, which sits a quarter turn from the textbook Milne-Eddington model's, because HMI counts Stokes Q with the opposite sign (section 6)
turning	the instrument's rotation of the sideways arrows, one way where the field points toward us and the other way where away (this paper)	the Faraday rotation of the light's polarisation as it leaves the magnetised gas; and, in the book, the whole turns a count adds up
twist	the polarity-odd part of the current in a band (this paper)	the helical twist of a flux rope, a property of the Sun's field; and, in the companion book and the author's Field Guide (Schulz 2026b), the Călugăreanu twist of a ribbon, which with writhe makes up its linking
helicity	not used here of the plasma	in the literature, the linkage of the field's own lines; in optics, the handedness of circular polarisation
holonomy	the companion book's word, used in this paper's subtitle: what an arrow carried round a closed path has turned by when it comes home; the loop count is one, read from the field's own arrows	in the book and in geometry, the same word covers the turning a carried arrow picks up along any path, closed or open; the light's polarisation is such an arrow, and its Faraday turning along an open path is the kind that has no natural zero (section 6)

2. The maps, and the light

The maps are SHARP maps from the Helioseismic and Magnetic Imager on NASA's Solar Dynamics Observatory: the series hmi.sharp_cea_720s, one map every 12 minutes of each active region, projected so that every pixel covers the same area of the Sun (Hoeksema and co-workers 2014; Bobra and co-workers 2014). Each map gives the upward field, the two sideways components, their errors, the program's confidence in its arrows, and a mask of the region. Six sets of maps are read, and for two moments the instrument's own light.

set	what	maps
Flare series	two regions through X-class flares, every 12 minutes for nearly three hours: NOAA 11166 (9 to 10 March 2011) and NOAA 11158 (15 February 2011)	30
Disk tracks	four regions followed across the disk, one map every day (NOAA 11158 and 11166) or every other day (NOAA 11429 and 12192)	38 read
One day	NOAA 11166 through one orbit of the satellite, every two hours	12
Blind moments	every numbered region at twelve moments fixed by the calendar: 00:00 TAI on the first of January, April, July and October, 2012 to 2014	68 read
Second blind moments	every numbered region at the other twenty-four first days of the month in those years: February, March, May, June, August, September, November and December; Stanford returned no records for 1 September 2014	143 read
Six tracked regions	the third blind test: six regions never read before, one map a day across the disk, from 2011, 2015, 2016, 2017 and two of 2024 (NOAA 11302, 12371, 12529, 12673, 13664, 13842)	85, 80 read
The light	the instrument’s polarised light, Stokes I, Q, U and V at each of its six filter positions (the series hmi.S_720s, 24 images, cut out round each region), for NOAA 11166 at 23:24 TAI on 9 March 2011 and NOAA 11158 at 00:00 TAI on 14 February 2011; and the pipeline’s own inversion of those same records in the detector’s frame (the series hmi.sharp_720s: field strength, inclination, azimuth, confidence and mask)	2 moments

Both sets of blind moments were chosen by a rule written before any of them was fetched. Each prediction was published with its fingerprint before its fetch. Stokes V, the circular polarisation, was fetched after I, Q and U; section 6 reads the light both ways.

Which code made them. Every map’s header names the code that made it. The headers of all 165 maps of the first four sets were read by script, and those of the second blind moments by the tool that read them. All of them were resolved for the arrowheads by the same code (AMBCODEV “disambig_v3 2013 Dec 06”), and none models scattered light (INVSTLGT “No”). They come from two versions of the inversion code, which turns the measured light into a field. Every blind-moment map from 1 January 2012 to 1 May 2013, and all seven maps of NOAA 11429, name versions of 2012 and March 2013 with no filter phase map. Every other map names “vfisv FD10 2013 Apr. 30; uses time-dependent HMI filter phase maps”. The pipeline’s overview says that inversions of data taken between 1 August 2011 and 23 May 2013 “use earlier, less accurate filter phase maps” (Hoeksema and co-workers 2014, section 4). Section 5 reads the offset under both.

Which maps the instrument itself doubts. Every map’s header also carries the pipeline’s QUALITY flags, whose 0x10000 bit marks a record the pipeline’s authors exclude from their own statistics, “typically missing filtergrams due to calibration sequences” (Hoeksema and co-workers 2014, section 7); a merged SHARP data set for flare forecasting excludes the same records (Kosovich and co-workers 2024, section 2.1). Read from the headers on the bench: all 158 records of the second blind moments and all 26 daily maps of NOAA 11158 and 11166 carry QUALITY 0; two of the twelve maps of the orbit day carry the bit, and the tool that read that day reported its numbers with and without them (section 5). The hourly tracks of section 13’s vane exploration are the exception: 18 of their 141 records carry the bit, at the 06:00 and 18:00 hours mostly, and were read with the rest; that exploration decided nothing, and the flag is one more reason it could not; read again without them, its numbers do not move (section 13). Seven maps of the third blind test’s NOAA 12529 carry QUALITY 128, another bit, from the day the inversion code’s name changes (section 11).

3. A count on a strong rim holds

The first question is the simplest: does a count behave like a label on the real Sun? A label is something a smooth push cannot change. The hardest push the Sun's surface gives in a few hours is a large flare, so the test is to follow counts through one.

Two active regions were followed for nearly three hours each, one map every 12 minutes, fifteen maps for each region. NOAA 11166 was followed from 22:00 TAI on 9 March 2011 to 00:48 on 10 March, through an X1.5 flare that began at 23:13. NOAA 11158 was followed from 00:36 to 03:24 TAI on 15 February 2011, through an X2.2 flare that began at 01:44. The regions sat near the centre of the disk, so they were seen almost straight on, and each series keeps one shape from map to map. A loop whose rim was readable in all fifteen maps was followed as the same square of pixels in every map, and its count was read in each.

Not one count changed. In NOAA 11166, 1,589 loops were readable in every map: 1,023 of side 8, 489 of side 16 and 77 of side 32. In NOAA 11158 there were 1,352 (891, 399 and 62). Not one of the 2,941 changed its count in any map. Most count 0. The 52 in the first region and 54 in the second that count a turn all count +1, the whole time, and they sit round the spots (Figure 3). That is what a spot should give: its field spreads out from its centre like the spokes of a wheel, or gathers into it, and either way the sideways arrows turn once round a loop that encloses the centre, the same way as the walk, whichever way the field points.




Figure 3

Figure 3. A count on a strong rim holds. Top: the upward field of each region in the first map of its series (white up, black down), with every loop readable in all fifteen maps outlined. Pale green loops count 0; orange loops count +1 and were firm in every map; purple loops count +1 and were not firm in every map. Bottom: map by map, how many of those loops were not firm (their rims held a step past a quarter turn). The dashed line marks the flare's start. No loop's count changed in any map.

How firm the ground was. The rim steps say more than the counts. A firm loop, one whose rim steps are all under a quarter turn, cannot have its count moved by any small push, as section 7 explains. In NOAA 11166, 1,573 of the 1,589 loops were firm in every map, and all 52 that count a turn were. In NOAA 11158, 1,343 of the 1,352 were firm in every map, and 51 of the 54 that count a turn. The other three are loops of side 32, overlapping one another inside that region's flare site. They held a step past a quarter turn in the first map, at 00:36 TAI, and were firm in all fourteen maps after it. So of the 106 counts of a turn, 103 rested on firm rims throughout.

The flares did not stir the rims that held. If a flare jolted the arrows on these rims, loops would change class across the flare's start. Across the pair of maps that brackets each flare's start, 2 loops changed class in NOAA 11166 and 1 in NOAA 11158. Two other pairs of maps in the first series, and three in the second, saw more. The biggest

change is a one-map blip at 00:36 TAI in NOAA 11166, nearly an hour and a half after its flare began, when 14 held loops were even-with-steps in that map only. The companion book, which read these same maps, found where the flares did their work: zeros of the sideways field came and went in the gaps between strong rims, never across one.

What this does and does not show. It shows that on the real Sun a count read round a strong, clearly measured rim behaves as a label should: a flare that rearranges the field nearby leaves it alone. It does not show that counts are unmoved everywhere. The loops followed are the ones readable in every map, which by construction avoids the weak, confused field along the line between the polarities, where a flare's change is concentrated. Two regions and three hours each are a first look, not a survey. What the result gives the rest of this paper is a label that can be trusted on strong rims, which is where sections 4 to 9 look for a false current.

4. Every spot's current seems to flow up

Around a sunspot, the current read from a rim that grows outward rises, peaks and falls back: the spot carries current one way through its middle and back round its edge. In both flare regions the book found all five spots' peak currents positive, whichever way the spot's field pointed. That is the puzzle this section pins down.

Split by polarity. A bundle of field twisted like a rope carries current along itself. Where it comes up through the surface, the current comes up with it; where it goes back down, so does the current. So the Sun's own twisting makes currents of opposite sign in the two polarities: the twist. Nothing in the Sun's physics prefers one sign in both polarities, because the laws that move the field treat a field and its mirror image alike. A part of the same sign in both, the offset, needs another cause. In the companion book's terms, the split is a parity decomposition: the Sun's equations respect the mirror, so whatever is even under it is the instrument's by construction, and the even part is an artifact meter.

Figure 4 splits the current near disk centre, band by band of field strength, in the six maps of the two flare regions seen most nearly straight on (μ at least 0.9, three maps each). In the two weaker bands, 200 to 1,000 gauss, the offset is small: between -0.59 and $+0.24$ mA/m². In strong field it is not. Each number here is the median over the three maps. In NOAA 11166, in the 1,500 to 2,000 gauss band, the current is $+4.88$ mA/m² where the field points up and $+3.93$ where it points down: an offset of $+4.42$ and a twist of only $+0.47$. NOAA 11158 carries a large twist, $+3.08$ to $+12.90$ in the three strong bands, so its down-field current comes out negative there; but the offset is positive in every strong band: $+3.25$, $+5.74$ and $+4.79$. The last of these is the median of three maps of which one, 13 February, has only 11 pixels of each sign in that band and an offset of $+22.14$; under the rule the other tools of this paper use, at least 50 confident pixels of each sign, that map does not count there, and the median of the other two is $+3.82$. Either way, the offset in the strong bands is 2 to 6 mA/m² in both regions.




Figure 4

Figure 4. Offset and twist near disk centre. The mean upward current where the field points up (red) and where it points down (blue), in five bands of field strength, as the median of each region's three maps nearest disk centre. The black bar is the offset: the average of the two currents in each map, and then the median over the maps. Medians do not add, so it need not sit exactly midway between the bars. In NOAA 11166 both polarities carry upward current in strong field. In NOAA 11158 a large twist pulls the two apart, but their average is still upward.

Across the disk. A current the Sun makes has no reason to care which way we look at it. A false one made by the instrument may. As the Sun turns, each region crosses the disk in about two weeks, so the same region is seen from many angles. In all four regions followed across the disk, the offset in the strong bands was largest near disk centre and fell toward the edge. In the two flare regions its correlation with μ ran from $+0.66$ to $+0.90$ across the three strong bands, each with a chance under 0.05 (from 10,000 shufflings). In the two others, with only five and six maps

each, it was positive in all six strong bands, +0.26 to +0.96, and passed that bar in four of them. In NOAA 11158 the twist, the part the Sun is expected to carry, stayed large all the way across and hardly followed the viewing angle. That was the cleanest contrast in the test.

The blind test. Four regions could share a quirk. So the question was put to regions never used in any fit, and the prediction was written first. It fixed the maps by a rule, every region carrying a NOAA number at twelve moments of the calendar; it fixed the arithmetic, the turning each region needs to cancel its offset in the strong bands; and it fixed what counts as a pass. Its fingerprint was published, and only then were the maps fetched. Of 83 records, 15 had no data yet, 68 were read, and 47 had enough strong field of both signs to need a turning at all (at least 50 confident pixels of each sign in a band of 1,000 gauss or more).

Near disk centre, 24 of 26 regions need a positive turning (Figure 5). If each sign were a coin flip, the chance of 24 or more of 26 would be about 5 in a million. It holds in both hemispheres: 16 of 16 regions in the south and 8 of 10 in the north. The mean turning needed near centre is about 3 degrees, smaller than in the four regions it was first seen in.




Figure 5

Figure 5. *The first blind test.* Left: every region read at the twelve moments, placed where it sat on the disk; an upward orange triangle needs a positive turning and a downward blue one a negative turning, and the size shows how much. Grey dots are regions with too little strong field to read. The dashed circles are $\mu = 0.75$ and $\mu = 0.40$. Right: the same 47 regions against the viewing angle. Near centre (blue band) 24 of 26 need a positive turning. Far out (red band) the prediction expected negative ones, and only 2 of 7 were; section 13 says why that part failed.

So the offset behind “every spot’s current flows up” is general near disk centre, not a property of four regions. A second blind test, on twenty-three more moments, found the same: 45 of 50 fresh regions near centre need a positive turning (section 10). The prediction had a second half, about the edge of the disk, which failed; it is taken up in section 13.

5. It behaves like the instrument

Which error can make it. Four kinds of error were put into the six near-centre maps on purpose, each of 10 degrees, and the maps were read again with the same arithmetic. Error A turns every sideways arrow one way where the field points up and the other way where it points down. Error B leans the field line toward the edge of the disk, as Rudenko and Dmitrienko (2018) report for these maps; error C leans the field vector the same way. Error D turns the whole map, as a control. Only error A left the offset’s fingerprint: current of one sign in both polarities, in the strong bands, with almost no twist. The two leans did next to nothing, because a lean points one fixed way across a spot, and its currents cancel round it. The whole-map turn made twist, not offset, as a control should. The change error A makes grows in proportion to its angle: on one map, 5, 10 and 20 degrees change the offset in the ratio 0.50 to 0.51 : 1 : 1.93 to 2.03.

How big, and from what strength. A turning that follows the field’s sign still needs a rule for how it grows with the field’s strength. Thirty-five ramps were tried, each switching on at one strength and reaching full size at a higher one. For each, the one angle that best matches the measured offset in five bands, in both regions at once, was found. The best ramps all start from zero and are at full size by 250 to 1,000 gauss, and all need 5.4 to 5.5 degrees. The misfit falls from 3.37 mA/m² with no turning to 0.76 to 0.89. Fitted one region at a time, the two need 5.2 to 5.6 degrees. So one turning of about five and a half degrees, with the field’s sign, reproduces the strong-field offset in both regions (Figure 6, A and B). The fit misses in the two weaker bands, where it puts current the maps do not show; section 9

comes back to that. Those numbers were fitted on the two regions' median offsets, ten numbers; the same ramp fitted on every one of the six maps' bands, 29 numbers, gives 5.30 degrees with a misfit of 1.05 (section 6). The two are different estimators of one thing, and each is named where it is used.




Figure 6

Figure 6. *It behaves like the instrument.* (A, B) The measured offset in five bands (dots) against the offset that a 5.5-degree turning with the field's sign, full from 250 gauss, makes when it is put into the same maps on purpose (bars); one angle fits both regions. (C) The turning each map of NOAA 11166 needs, through one day of the satellite's orbit (orange), against the satellite's speed away from the Sun (grey). The spectral line shifts by about one tuning step of the instrument every 12 hours, and the turning does not follow it.

It does not ride the orbit. The satellite's orbit carries it toward and away from the Sun each day, which shifts the spectral line back and forth across the instrument's six filter positions, "by about one tuning step every 12 hours" (Hoeksema and co-workers 2014, section 7.1.1). The pipeline's overview reports that the measured field strength in sunspots swings with that orbit, by about one per cent in umbrae (section 7.1.2). A turning made at a particular place in the line would be expected to swing too. It does not. Through one day of NOAA 11166, twelve maps two hours apart, while the speed ran from -2.52 to $+3.40$ km/s, the turning each map needs stayed between $+2.8$ and $+5.2$ degrees, mean $+4.36$, with no link to the speed (correlation $+0.28$, chance 0.38; Figure 6, C). It is also steady: the twelve maps' turnings have a spread of 0.72 degrees. Two of the twelve carry the pipeline's calibration-sequence flag (section 2); the ten clean maps alone give $+2.8$ to $+5.2$ degrees, mean $+4.41$ (run SUN_OFFSET_REMOVAL 20260927_193239).

Blurring does not make it. A region near the edge of the disk is seen slanted and so more blurred, which shrinks any current read from it. So an offset that fades toward the edge does not by itself point to the instrument. The test was run in house: near-centre maps were blurred as a slanted view would blur them and read again. Blurring shrank neither the turning nor the twist by anything like what the disk maps show, and blurring alone never made a map need a negative turning.

It is made under both versions of the inversion. Section 2 found the blind maps made by two versions of the inversion code, one with earlier, less accurate filter phase maps. The offset is in both. In the first blind test, looked at afterwards, 13 of 14 regions near centre inverted with the older code need a positive turning, and 11 of 12 with the newer. The second blind test fixed this as a test before its maps were fetched, and it passed: 17 of 18 older-code regions and 28 of 32 newer-code regions near centre need a positive turning. The book's two flare regions, where it was first seen, were inverted with the newer code. So the older phase maps do not make the turning, and the 2013 improvement does not remove it.

How big. The 5.5 degrees was fitted on two large regions. Fitted over every near-centre map in hand, 60 maps of the blind moments and the four regions, the best single angle is 2.97 degrees (2.40 on the moments alone, 3.76 on the four regions alone). That is a fit. The second blind test fixed a line before its maps were fetched, halfway between 2.97 and 5.5 degrees, and the median near centre came in below it, at $+3.24$ degrees.

Who met it before. A turning of this kind has been found in another kind of instrument. Gao, Su, Xu and Zhang (2008) measured the turning of the field's direction that Faraday rotation makes in a filter magnetograph at Huairou, found it growing with the field, and reported that uncorrected maps carry "false positive currents": a mean vertical current of 3.3 mA/m², the size of the offset here, reduced by 88.8 per cent by their correction (read in abstract and summary). This instrument's inversion models such magneto-optical effects. Centeno and co-workers (2014), describing its optimisation, report a coding error in them, "a missing factor of 2 in the computation of the imaginary part of the Taylor expansion of the Voigt function, which affected the magneto-optical effects in strong field regions", which led to "inhomogeneous patterns in the inverted magnetic fields inside sunspot umbrae". It is fixed in the version that made these maps. No report of the same signature in these maps was found in what was read.

6. Where the turning is made

The maps say what the turning does, not where it is made. With sections 4 and 5 written, three places remained and nothing decided between them. It could be a Faraday rotation the inversion leaves in. It could be scattered light, which the pipeline does not treat and which “can lead to systematically lower inferred field strengths and incorrect inferred angles of the magnetic vector” (Hoeksema and co-workers 2014, section 7.3.1, citing LaBonte 2004). Or it could be the inversion’s one-layer model, in which “all the parameters describing the atmosphere are constant along the line of sight (LOS) except for the source function” (Centeno and co-workers 2014), asked to fit a sunspot whose field changes with height. This section follows the turning as far as it can be followed from public data: first into the maps’ own frame, then into the instrument’s light. In the companion book’s words for a false signal, the hunted is the Sun’s field, the hunter is the inversion, and the substrate is the light; what follows moves the cause from three candidates to the hunter, on a light the hunter’s own family of models does not fully describe. Except where the text says otherwise, everything in this section is exploration, run once and printed, not a test fixed before its data were fetched; and every number names its run.

The turning’s frame. A turning made by the instrument should live in the instrument’s frame, about our line of sight, with the sign of the field along that line; a turning the Sun made would live in the Sun’s, about the local vertical, with the sign of Br. At disk centre the two are the same thing. Away from centre they part, and a map read both ways at once gives a ruler. Every map of both blind tests, the disk tracks and the orbit day, 292 records in all, was read twice: the paper’s way, by Br, and the instrument’s way, by the line-of-sight field; then a known turning of 4 degrees of each kind was planted and both readings were read again, so that the response of each reading to each kind is known map by map. The reading by Br reproduces the second blind test’s angles to 0.01 degree on all 143 of its records. Where the ruler separates the two kinds, which it does only at mid-disk, the maps carry the line-of-sight kind: for μ between 0.40 and 0.60, kind V -0.99 degrees and kind L $+3.99$; for 0.60 to 0.75, $+1.02$ and $+4.02$ (Figure 10, B). Nearer the centre the two kinds are too alike for the ruler to separate them (determinants 0.206 and 0.068, against 1.038 and 0.589 at mid-disk), which is as it should be. Read the instrument’s way, the median turning is $+3.05$, $+5.31$, $+4.77$ and $+5.52$ degrees across the four bins of μ from 0.40 to 1; read the paper’s way, -0.57 , $+1.25$, $+2.98$ and $+4.31$ (Figure 10, A). So the fade that section 10 describes is what a turning of about 4 degrees about the line of sight looks like when it is read about the vertical. That is description, and it was then read again as a test on six regions never read before, where the line-of-sight kind won both mid-disk bins again but at half the size, $+1.49$ against $+0.89$ and $+2.18$ against -0.34 , short of the 2-degree margin the prediction demanded in one bin (section 11); so it stays description, read twice with the same sign. A second ruler, which reads the twist of a spot’s own field on the two sides of the line where the line-of-sight field changes sign (the vanes), responds to the same planted kinds in the same order ($+2.37$ to 4 degrees of kind L, -0.11 to kind V) but scatters too widely to say anything on 59 readable spot-maps: its median step is -0.43 degrees with a bootstrap range of -3.13 to $+0.93$. Two hourly tracks of the two flare regions across the disk, 69 and 72 maps, were read to sharpen it, with each sector of a spot’s rim followed through time as the seam swept across it, so that the sector is its own control. Twenty-five sectors in four spots had both sides. Read over all their maps, the sectors’ median step is $+1.11$ degrees (bootstrap range -15.77 to $+1.76$); read over only the six maps of each side nearest the crossing, $+4.18$ (bootstrap $+1.58$ to $+5.83$), positive in 15 of 25, while a planted turning of 4 degrees of the line-of-sight kind moves the step by $+2.51$ and of the vertical kind by nothing. The local reading agrees in sign and size with the dead-on ruler and the whole reading does not, and the reason is in the sectors themselves: on one spot the step runs smoothly from $+25$ degrees on one side of the rim to -30 on the other, which is the spot’s own vanes changing by tens of degrees over three days, far more than any turning of a few degrees, so a sector is its own control only for hours. That is a wall with a number, and the vane test is kept as exploration on both readings. Runs of record: SUN_TURNING_DEADON 20260928_215526; the vanes, SUN_TURNING_VANES 20260928_220608 and 20260929_005842.

What the instrument records. HMI does not measure a field. It measures the polarisation of light in one spectral line of iron, at 617.3 nm, through six filter positions across the line, and a one-layer Milne-Eddington inversion turns the four polarisation signals at those six positions into a field strength, an inclination and an azimuth for each pixel (Hoeksema and co-workers 2014; Centeno and co-workers 2014). The azimuth is the direction of the linear polarisation, read from Stokes Q and U. The model of the line has magneto-optical terms, the Faraday terms, which rotate the polarisation as the light climbs out through magnetised gas, and a fit whose model of those terms is wrong in a plausible way will report an azimuth turned from the true one by an amount that depends on the field’s strength and sign. That is the first candidate above, stated as a question the light can answer. For two moments, one of each flare region, the light itself was fetched: Stokes I, Q and U at the six positions, and the pipeline’s inversion of those same records in the detector’s frame, so that the light and the arrows made from it can be set side by side pixel by pixel.

A test that could not be read. The first attempt was a test fixed in advance. Its premise was that the direction of the linear polarisation at any one filter position is the field’s azimuth plus a constant for that position, so that a Faraday rotation left in by the inversion would show as the line’s core turned against its wings, opposite ways for the two polarities. The prediction was fingerprinted before the light was fetched, and the tool had two gates it had to pass before anything was read. The first, that the light and the maps were the same pixels, passed (correlations of the

light's intensity with the maps' continuum 0.9991 and 0.9983). The second, that the polarisation's direction sat a constant angle from the reported azimuth at each position over the calibration pixels, with a spread under 10 degrees, did not: the spreads ran from 2.46 to 14.63 degrees in one region and 3.34 to 15.49 in the other, and by the rule written before the run nothing past that gate was read. A second, looser reading declared after the first had stopped also stopped at the same gate. Runs of record: SUN_STOKES_TURNING 20260928_194600 and 20260928_195923. The test is kept in section 13 as a property of the problem: the premise was wrong, and the next paragraph says how.

The light's mirror. Read as description, the direction of the polarisation minus the reported azimuth, position by position and band by band of the line-of-sight field, is not a constant with scatter round it. It is a curve that flips with the field's polarity, and the curve at the bottom of the line is the mirror image of the curve in the wings (Figure 7). In NOAA 11166, half the difference between the two polarities' medians in the strong band is +29.37 degrees at the bottom sample of the line and -5.68, -5.74, -8.58 and -11.62 at the other positions; in NOAA 11158, +30.29 at its bottom sample and -5.47 to -11.61 at the others. The two regions were observed a month apart and agree at every position. A one-layer Milne-Eddington model of the line, with its magneto-optical terms, given each pixel's own field from the inversion and sampled by six Gaussian filters of HMI's width and spacing, makes the same mirror: +15.57 and +12.36 at the two bottom samples and the other sign at the other positions, with a control of exactly zero when the terms are switched off (run SUN_STOKES_FORWARD 20260928_205957). So the gate that stopped the test was not noise. West and Hagyard measured the same thing in 1983, the azimuth's Faraday rotation changing with wavelength across a sunspot, and matched it by letting the line's strength vary over the spot. The polarisation's direction at one filter position is not the field's azimuth plus a constant; the Faraday terms turn it by an amount that changes with the field's strength and sign and changes sign across the line, and the model says so. The model's mirror is weaker than the data's, by half at the bottom and by several times in the wings, where the model assumes a line strength and width that the fits below find are not the line's. The line's width and strength were assumed for this run (a Doppler width of 34 mÅ and a line strength η_0 of 5), and every model number in this section names what it assumed.




Figure 7

Figure 7. *The light's mirror.* For each region, the direction of the light's linear polarisation minus the arrow the pipeline reported, wrapped into a quarter turn either way, as medians by band of the line-of-sight field: at the bottom of the spectral line (red) and in its wing (blue), the light itself (filled) and the one-layer model with magneto-optical terms (open). The curve at the bottom of the line is the mirror image of the curve in the wing, in both regions, and the model makes the same mirror with the wrong size. Run of record: SUN_STOKES_FORWARD 20260928_205957.

A fitter with weak Faraday terms. The question the mirror leaves is what a one-layer fit reports when its Faraday terms are wrong. It was put to a model: light was made with the same Milne-Eddington model, Faraday terms and all, for 520 pixels of each region drawn at random from the inversion's own confident pixels (up to 40 in each band of the line-of-sight field, so every band is represented), with each pixel's true field strength and inclination the inversion's own and its azimuth drawn at random; then a one-layer fitter with eight free numbers, damping fixed at 0.5 and filling factor fixed at 1, as HMI's, fitted the 24 sampled values back, and the fitted azimuth minus the true one is the turning. Every fit starts from the truth, so what a wrong model reports is the bias the wrong model forces, and a matched fitter reports exactly zero, as it must. What each variant would mean was written into the tool before it ran. A fitter whose Faraday terms are half the model's turns the strong-field arrows by 3.75 and 3.57 degrees in the two regions (half the difference between the polarities' medians, over 1,000 to 2,500 gauss), and one that ignores them by 9.20 and 9.08; the two regions agree to within 0.2 degrees at every variant, as they should, since only the pixels' fields differ. A wrong damping, a wrong line strength, or a line shifted by one filter position changes the number by less than a degree, and a wrong filling factor turns the arrows by 2.6 to 3.5 degrees with a shape of its own. So the size is right, the maps needing 3 to 7 degrees, and the sign flips with polarity, as the maps' does. What does not match is the shape

across field strength. Under-modelled Faraday terms turn the weak-field arrows the other way: with the terms at half, the medians by band in NOAA 11166 run +0.8, +1.3 and +0.9 degrees where the line-of-sight field runs from 100 to 1,000 gauss and then -2.0 , -4.3 and -4.9 from 1,000 to 2,500, crossing zero at 943 gauss, and at 947 in NOAA 11158 (Figure 8, A). Run of record: SUN_STOKES_FIT 20260928_225848, variants A to I.

Where the crossing sits, and whether it can move. The crossing sits where the Zeeman splitting of the line has grown past its Doppler width: the splitting is 44.5 mÅ at 1,000 gauss, and it equals one assumed width of 34 mÅ at 764 gauss. A narrower line, as in a cool umbra, should move the crossing down in proportion, and a wider one up. The expectation was written into the tool before the run, 690 gauss for a width of 25 mÅ and 1,390 for 50; the run gave 780 and 821 gauss at 25 mÅ and 1,162 and 1,240 at 50. It moved the same way and about half as far, because the fitter fits the width and takes up part of the change; the expectation is kept in the ledger beside the measurement as a bracket. Across every plausible width, then, the crossing moves from 687 gauss (a fitter that ignores the terms, at 25 mÅ, in NOAA 11158; 720 in NOAA 11166) to 1,240 (the half-term fitter at 50 mÅ, in NOAA 11158; 1,162 in NOAA 11166), and never below 500. Same run, variants J to M.

The curve planted into the maps. The maps can say whether their turning has that shape, because the offset they carry is measured band by band of field strength. Three curves were planted into the six near-centre maps of section 4 with the paper's own arithmetic, each with one free scale fitted by least squares over every map and band: the one-sign ramp of section 5, the half-Faraday curve above (its medians by band, averaged over the two regions), and the filling-factor curve. The ramp leaves an rms offset of 1.05 mA/m² (at 5.30 degrees), the Faraday curve 2.46 and the filling curve 2.95, against 3.30 with nothing planted; on the thirteen maps with μ of 0.75 or more, 1.89, 2.50 and 2.71 against 2.81. The place the Faraday curve breaks is the 1,000 to 1,500 gauss band, where the maps measure an offset of +2.53 mA/m² and the curve's best fit predicts -1.31 , because the model's turning changes sign near 1,000 gauss and the maps' offset is already large and positive there (Figure 8, B). The weaker bands decide nothing: the maps measure +0.08 and -0.02 there, and every curve predicts something small. So the maps' turning holds one sign from at least 1,000 gauss up, and the mis-modelled Faraday term of this model does not. Run of record: SUN_TURNING_PLANT 20260928_225126. One thing the planting also shows is that the curve fits the maps only turned over: its fitted scale is -1.10 . Whether that is a real opposition of sign or a convention between the model's frame and the maps' is not settled here. The maps' frame is the projected map's, and the model's is the detector's, and section 6 pins one convention between them below but not this one; the pattern across field strength is what is compared, and it is decisive on its own.




Figure 8

Figure 8. *The whale: right size, polarity flip, wrong shape.* (A) On model light, the turning a one-layer fitter puts into the arrows when its Faraday terms are half the model's (red), at Doppler widths of 25 (orange) and 50 mÅ (purple), and when it ignores them (blue): the fitted azimuth minus the true one, medians by band of the line-of-sight field, averaged over the two regions; vertical lines mark each curve's crossing of zero. (B) The measured offset in the six near-centre maps, band by band (black), against the offset each of three planted curves makes at its best scale: the one-sign ramp (blue), the Faraday curve of panel A (red) and the filling-factor curve (purple). The Faraday curve dips below zero at 1,000 to 1,500 gauss where the maps rise. Runs of record: SUN_STOKES_FIT 20260928_225848 and SUN_TURNING_PLANT 20260928_225126.

The idea behind the reverse play: the observer's phase. Everything the instrument knows about the field it learns from an arrow the light carries. A photon leaving the Sun through a sunspot's field carries the direction of its linear polarisation, and the field turns that arrow as the photon climbs out, by the magneto-optical effect, by an amount that grows with the field's strength along the line of sight and flips with its sign. The observer at the far end, the instrument, does not see the field; it sees the arrow at the end of the path, at six colours across the line, and reasons

back to the field with a model of the path. So the arrow's direction at the observer is the sum of two things: where the field pointed, and how much the path turned the arrow on the way out. A model that gets the path right takes the turning off and reports the field; a model that gets the path a little wrong leaves a little of the turning in, and calls it field. That is the whole of section 6 in one sentence, and it gives a way to ask the question without a second instrument. If two observers read the same arrow with two models of the path, they disagree about the field by exactly the difference between their models, pattern for pattern: the same size where the models agree and a curve with the field's strength and sign where they do not. The author framed this as a play run backwards, with the instrument as the defender who reads the arrow as the field's own direction and is drawn by the turning, and asked whether the turning is a price the light pays to leave the Sun or an optical effect of its passage. The reverse play is the tool built to ask it: fit the instrument's own light with a textbook model of the path, set our arrows against the instrument's pixel by pixel, and read the pattern of the disagreement. Two things follow from the idea before any fit is run. The pattern, not the absolute direction, is what carries meaning, because on an open path from the Sun to the observer the arrow has no natural zero: where the observer chooses to count the direction from is a convention, and a fixed offset between two observers is a convention and not a turning. And a disagreement between two observers is not a measurement of either one's error, because both may be wrong about the path in ways that cancel or add; only a third reading, by an instrument that sees the path differently, can say which. Both of those showed up in the run, and the paragraph after this one says how.

The reverse play. The model can be turned round: instead of asking what a wrong fitter does to model light, fit the instrument's real light with a textbook fitter and set its arrows against the instrument's. The same one-layer fitter, with its Faraday terms in full, at half, and ignored, was given HMI's own light at the six positions for 644 and 650 confident pixels of the two regions (up to 50 in each band of the line-of-sight field), starting every fit from HMI's own answer. It was run twice: first on Stokes I, Q and U alone, before the circular polarisation V had been fetched, with each pixel's inclination held to HMI's hemisphere because without V the light cannot fix the field's polarity on its own; then, with V, on all four Stokes with the inclination free, so that the light fixes the polarity itself. Two conventions showed themselves first. Our fitted azimuth sits a quarter turn from HMI's at every band and both polarities, 88.2 to 90.8 degrees over the weak pixels, because HMI counts Stokes Q and U with the opposite sign to the textbook model; and HMI's V is the model's V with its sign reversed (the correlation of the two, sample by sample over the strong pixels, is -0.988 and -0.981 in the two regions, positive in none). A fixed quarter turn and a fixed sign are conventions, not turnings, exactly as the paragraph above said they would be; the tool takes them out and prints them, and the quarter turn is the reason the light's mirror above is drawn wrapped into a quarter turn. With the Faraday terms in full and all four Stokes, our arrows differ from HMI's by -6.10 and -5.28 degrees in the strong band and $+1.81$ and $+2.10$ in the weak band, an S-shape crossing zero near 1,000 gauss, in both regions alike (Figure 9); on I, Q and U alone the same fitter gave -5.04 and -3.88 and $+2.28$ and $+1.75$, so adding V and freeing the polarity moved the reading by about a degree and changed nothing in its shape, and only one pixel of the 1,294 crossed into the other hemisphere from HMI's. Weakening the terms moves our arrows further from HMI's, not closer: -12.75 and -11.89 at half, -24.25 and -24.14 ignored. And the full-term fitter fits the light best, with the smallest residual of the three in both regions on either reading (on I, Q and U: 5.3 and 4.3 parts in a hundred thousand of the light's variance, against 6.7 and 5.9 at half and 8.2 and 9.1 ignored; on all four: 22 and 16 against 26 and 20 and 32 and 28); so the light itself demands the Faraday terms, and the mirror of Figure 7 is in the data, not only in the model. The fits also say what line they see: a line strength of about 15 to 17 and a Doppler width of about 22 mÅ (medians; the two trade off against the fixed damping), stronger and narrower than the 5 and 34 mÅ assumed above, which by the width run would put a Faraday crossing near 700 gauss rather than 940. Run of record: SUN_STOKES_REFIT 20260929_051925, whose first parts reprint the I, Q and U run 20260928_235356 and the four-Stokes run 20260929_022834 line for line.




Figure 9

Figure 9. *The reverse play.* HMI's own light, fitted with a textbook one-layer Milne-Eddington model: our fitted azimuth minus HMI's, less the quarter-turn convention, as medians by band of HMI's line-of-sight field, with the Faraday terms in full (red), at half (orange) and ignored (blue); solid with all four Stokes and the polarity free, dashed with I, Q and U only and the polarity held to HMI's. HMI's arrows are the zero line. Relative to any fitter of this family, HMI's arrows are turned as if its magneto-optical terms were the stronger, by an S-shaped few degrees that flips with polarity, and the circular polarisation changes nothing in that. Run of record: SUN_STOKES_REFIT 20260929_051925.

What that says, and what it does not. HMI's arrows lie on the far side of our full-Faraday fit: relative to any fitter of this one-layer family, HMI's inversion turns the arrows as if its magneto-optical terms were stronger than a textbook Milne-Eddington line with Gaussian filters gives, by a few degrees, flipping with polarity, with the S-shape of the Faraday mechanism. Which side is right, HMI's or the textbook's, this paper cannot say, because the difference between two fitters is not the difference from the Sun: two readers of one light are not a witness. The maps' verdict from the planting stands beside it. Whatever the turning relative to the truth is, it holds one sign from 250 gauss up, and neither fitter's S-shape is that. Put together: HMI's arrows and a textbook one-layer fit of the same light, all four Stokes, disagree by an S-shaped few degrees, the maps' offset says the true turning is one-signed, so at least one of the two, and possibly both, is off the truth by a one-signed turning plus part of an S. Two of the pipeline's own fitting choices can be copied from its papers and were: VFISV weights the four Stokes unequally, at about 1, 5, 5 and 3.5 for I, Q, U and V before its noise normalisation, chosen for the smoothest solution in umbrae (Centeno and co-workers 2014, section 2.1), and it adds to its merit function a term that pulls the line strength toward 5 (their section 2.4), where the textbook fits above weighted the four alike and let the line strength fall where the light put it, near 15. With the pipeline's weights, our arrows differ from HMI's by -5.62 and -5.07 degrees in the strong band; with its regularisation, -6.18 and -6.24 ; with both, -5.65 and -5.16 ; against -6.10 and -5.28 without them. Neither choice moves the arrows, so the disagreement is not made by them. What remains, and cannot be copied from public data, are the instrument's measured filter profiles with their side lobes in place of the Gaussians assumed here (their section 3.1), the two corrections the pipeline applies to Q and U before any inversion, a term proportional to I and a leak of I into the linear polarisation removed with a small kernel (Hoeksema and co-workers 2014, section 2.2), the pipeline's own path to its minimum, and whatever in the light a single layer cannot describe. The external witness for the turning is not a third fitter but a second instrument on the same region, a spectropolarimeter such as Hinode's, which this paper cannot run. That, and a run of the pipeline's own inversion on these cutouts against a textbook fit with the magneto-optical terms in full, are where to point a specialist.

The carried arrow. One reading of all this comes from the companion book, whose second section says that a winding counts and a carried turning measures. The light's linear polarisation is a carried arrow: it is what Tomita and Chiao's photon carried along a coiled fibre in 1986 and what Berry's phase (1984) describes, an arrow moved along a path and read at the end. HMI reads the carried arrow's direction and calls it the field's. The Faraday turning is what the arrow picks up along an open path out through the magnetised gas, and on an open path the absolute zero of such a turning is a convention, which is why the quarter turn found above is a convention and only the pattern across field strength and polarity is physics. The count of section 3 is the other thing: a winding, which the carried arrow cannot move except at a near-reversal, as section 7 shows. That reading is the book's, and the paper takes it from there.




Figure 10

Figure 10. *The dead-on ruler and the fade.* (A) The median turning needed, by band of viewing angle: the second blind test's reading by Br (black), the same maps and the disk tracks read by Br (blue) and by the line-of-sight field (red; bins of five maps or more). Read the instrument's way, the turning holds at 3 to 5.5 degrees from mid-disk to the centre; read the Sun's way it fades. (B) A turning of 4 degrees of each kind planted and read both ways gives a ruler,

and the ruler solved for how much of each kind the maps carry, in the two bins of μ where it separates them: about 4 degrees of the line-of-sight kind and about none of the vertical kind. Runs of record: SUN_TURNING_BLIND 20260928_181524 and SUN_TURNING_DEADON 20260928_215526.

7. The count reads through the turning

A turning that changes smoothly across the map cannot add a whole turn round a loop. It changes the step between two neighbouring arrows only by the difference of their two turnings, which for a turning of 5.5 degrees is at most 11 degrees. A step whose size stays under half a turn is counted the same way before and after, so the count cannot move. It can move only where a rim already holds a step close to half a turn, which a small push carries past, so that the shorter way round flips from one side to the other. On a firm loop every step is under a quarter turn, so no turning smaller than an eighth of a turn (45 degrees) can move its count.

Put in on purpose. The 5.5-degree turning was put into the 38 maps of the four regions followed across the disk, and every readable loop was counted again. Of 154,267 readable counts, 31 changed. At 20 degrees, 97 changed. No firm loop changed at either angle, as the arithmetic requires. None of the 31,209 readable loops in the day's twelve orbit maps changed at 5.5 degrees.

Where the 31 sit. All 31 sit on just five steps that were already in the maps, each between neighbouring pixels whose arrows point 172.7 to 179.2 degrees apart. Neighbouring loops share rim pixels, so one step can move many counts: the step in Figure 11 moved 15. Four of the five straddle the line where the field turns from up to down. The field on either side of them was +10 and -179 gauss, +270 and -939 , +20 and -110 , and +247 and -89 . The fifth sits within one polarity, between -842 and -79 gauss. The blind moments add one more of the same kind: in NOAA 11445 on 1 April 2012, 3 counts moved on one step of 176.2 degrees across the sign change. And in the flare series, the turning moved no count in NOAA 11166's fifteen maps and 13 in NOAA 11158's, all in one map and on one step of 176.6 degrees.




Figure 11

Figure 11. One near-reversal step. NOAA 12192 on 27 October 2014. (A) The upward field (red up, blue down) with the 15 loops whose counts the 5.5-degree turning changed; all of them pass through one step (star). (B) Close up, every arrow, with the two neighbouring pixels circled; their arrows point 179.0 degrees apart, one where the field is +20 gauss and one where it is -110 . (C) Measured, the step counts as +179.0 degrees; after the turning it counts as -178.1 , because the shorter way round has flipped. Every loop through it changes its count by one.

What it means. The current is a measure, and it takes the turning in full: that is how the offset was found. The count is a label, and it reads through the turning everywhere except at a handful of near-reversals, places where two neighbouring arrows already point almost opposite ways. There were six such steps in the 106 maps of the disk tracks and the blind moments, and one more in the thirty maps of the flare series. That is the old rule of phase unwrapping, found in radar and optics long before this paper: adding up wrapped steps is safe while every true step is under half a turn (Itoh 1982). It is also the reason for the classes of section 8. A near-reversal is exactly the kind of step that the headless count flags.

8. Two counts, with opposite blind spots

The choice of arrowhead. The instrument does not measure which way along its line the sideways field points. It measures the line, from the light's linear polarisation, and a program then chooses each arrowhead. This is the "180-degree ambiguity", resolved in these maps by a minimum-energy method (Metcalf 1994; compared with other methods by Leka and co-workers 2009). Where the choice goes wrong over a patch, a line of flipped arrows runs through the map, and a loop whose rim crosses it can have its count moved. The pipeline's overview shows one such failure, a checkerboard where neighbouring arrowheads alternate (Hoeksema and co-workers 2014, section 7.2.2).

A count that never sees the choice. Take twice each arrow's angle. That is the same for both arrowheads, since turning an arrow half a turn turns twice its angle by a whole turn. Add the steps of twice the angle round the rim, each taken the shorter way, and divide by a whole turn. This **headless count** reads only the measured line. It is the book's own headless read from its liquid crystals, met here on real maps. On a firm loop the headless count is exactly twice the arrow count. Each step past a quarter turn moves it by one:

$$\text{headless count} = 2 \times \text{arrow count} - (\text{the number of steps past a quarter turn going forward, minus the number going back}).$$

So a loop with an odd number of such steps has an odd headless count: a **half count**, a line that comes home turned half a turn, which no field with arrowheads can do (Figure 2).

Half counts are residues. The headless count reads only what the instrument measured, so a half count is a property of the measured line itself: a place where it turns by more than a quarter turn between neighbours, through noise or fine structure. Radar interferometry met the same thing in 1988, as "residues": small loops round which a wrapped phase fails to close, which "can readily be identified and avoided in phase estimation" (Goldstein, Zebker and Werner 1988). By the arithmetic of the pixel grid, a line of steps past a quarter turn can end only on a half count or at the edge of readable field. So the lines where the program's arrowheads flip end on half counts, and loops that such a line crosses twice are even-with-steps. Half counts are 0.90 per cent of the readable loops in the four regions, and 0.56 per cent in the blind moments.

Opposite blind spots. The two counts fail in opposite places. The arrowhead choice cannot move the headless count at all, but it can move the arrow count wherever a flipped patch crosses a rim. A smooth turning moves the arrow count only at steps near half a turn, but the headless count at steps near a quarter turn. Steps near a quarter turn are about three times as common: in the four regions' 38 maps, 997 steps between readable neighbours lie within 10 degrees of a quarter turn, and 336 within 10 degrees of half a turn, out of 4,689,666 (Figure 12, A). So under the same turning the headless count moves three to four times as often: 104 against 31 at 5.5 degrees, and 424 against 97 at 20 degrees (Figure 12, B). The book found the same price in pure mathematics, for lines with no head: the headless read has half the margin.




Figure 12

Figure 12. *Two counts, with opposite blind spots.* (A) The size of every step between neighbouring readable pixels in the 38 maps of the four regions (log scale). Steps near a quarter turn (purple) are about three times as common as steps near half a turn (orange). (B) How many counts a turning with the field's sign changes: the arrow count moves only at steps near half a turn, the headless count at steps near a quarter turn.

The trust mark. A count both agree on, a firm count, is the one to trust: no near-reversal sits on its rim, and no push of under an eighth of a turn at every pixel can move it. Firm counts are 88.8 per cent of the nonzero counts in the four regions (3,011 of 3,390), and 97.8 per cent in the blind moments (1,304 of 1,334). The share depends more on the

region than on where it sits on the disk. The half share of nonzero counts is 8.8 per cent in NOAA 11158, 7.3 in 11166, 1.8 in 11429 and 16.1 in 12192, against 1.8 per cent pooled over the blind moments, whose regions are mostly smaller.

What each reading sees. Four readings of the same maps appear in this paper, and each is blind to something the others see. Set out as the companion book sets out its instruments:

reading	sees	blind to
the arrow count (sections 3 and 7)	a zero of the sideways field enclosed by a strong rim; holds through a flare	a patch of arrowheads chosen the other way that crosses its rim; a near-reversal step pushed past half a turn
the headless count (section 8)	the measured line only; the arrowhead choice cannot touch it; flags residues	steps near a quarter turn, three times as common, where a smooth turning moves it
the current's offset (sections 4, 5 and 9)	the instrument's turning, in full, band by band, with its polarity flip; the vertical loop	the Sun's twist, which it sets aside by construction; anything even under the mirror that is not a turning
the light's fits (section 6)	the turning's mechanism in a model family; the convention between frames; which fitter the light prefers	the truth: a second fitter is not a second instrument

Where the arrowheads may have been chosen the other way. A smooth field sampled finely gives steps that thin out steadily toward half a turn. These do not: in both sets the steps stop thinning from about 130 degrees and run flat to 180 (the four regions' steps are in Figure 12, A). The last five bins, 130 to 180 degrees, hold 174, 174, 182, 169 and 167 steps in the four regions. Those flat tails are where to look for arrowheads chosen the other way inside readable field, and the half counts are their map.

The halves come in pairs. A half count is a residue of $+\frac{1}{2}$ or $-\frac{1}{2}$, and by the same arithmetic a residue cannot stand alone: a line of steps past a quarter turn has two ends, and whatever ends it, a seam of arrowheads chosen the other way or a real half-turn of the measured line, the two ends carry opposite halves. The author's picture of this, drawn before the run, is a pair of Savonius rotors, the S-shaped two-bladed wind rotor, of opposite bias, each unable to close on itself and the two balancing only together; the companion book gives a ring of phases that rotor's two properties, a handedness and a lagging drive, and finds that the ring then freezes with the sign the shape chooses, an engine that writes a count. So an exploration asked whether the maps' half residues sit in nearest-neighbour pairs of opposite sign, reading each 2-by-2 block of readable pixels as the smallest loop and merging touching residues of one sign (`sun_residue_pairs.py`, run `SUN_RESIDUE_PAIRS 20260929_214315` on the hourly tracks; the daily tracks and the blind moments as description). They do. On the 116 hourly maps with at least 10 half residues, the share whose nearest other residue is of the opposite sign has a median of 0.725 against 0.500 when the signs are shuffled among the same positions, above the shuffle in 107 of the 116; the nearest opposite partner sits a median of 3.4 pixels away and the nearest same-sign residue 12.9. The 12 such moment-maps and the 8 daily maps give 0.715 and 0.667 against 0.500, above the shuffle in every one. Half residues are rare, about one in 2,000 readable blocks, and about a quarter of them persist to the next hourly map within 3 pixels (442 of 1,676 and 253 of 1,164 on the two tracks), so most are a map's own. Two things this does not decide, stated as walls. The pairing cannot tell a real half-turn from a seam of the arrowhead choice, because a seam ends in an opposite pair too; what would tell them apart is whether the pair moves with the field from hour to hour, and a quarter persisting is neither answer. And the halves are not quite neutral: the net half charge over a map runs slightly positive, a median of +0.08 of the number of residues on the hourly maps and +0.02 on the moment-maps, an excess of one sign that is reported and not explained.

9. The current makes a loop

A current the Sun carries along its field goes up through one polarity and down through the other. Over a whole region its loop closes across the line between them, so the net current over both polarities is near zero, whatever the region's twist. Each polarity on its own may carry a net current, and Sun and Cheung (2020) show that a twisted tube meeting the surface at a slant projects one into each polarity, of opposite sign in the two; that is a part of the twist, and it too closes across the line. The instrument's turning is different. It pushes current the same way in both polarities where the field is strong. And a turning that fades to nothing at the map's edge adds nothing to the circulation round that edge. By Stokes' theorem, the net current through the whole map is that circulation, so the turning's false current must come back down inside the region, in the weaker field round each spot: its loop closes inside each polarity. The instrument is confident of its arrows mostly in strong field. In the median map near disk centre it is confident of 99.9 per cent of region pixels above 1,000 gauss, 68.5 per cent between 250 and 1,000, and 2.5 per cent below 250. So a sum over confident pixels should keep the false current's way up and lose much of its way down. This paper calls it the vertical loop.

Added up three ways. In the 41 blind-moment maps near disk centre, the net current over confident pixels comes out upward in 30, a chance of 0.0022 if each sign were a coin flip. It holds in both hemispheres, 13 of 18 in the north and 17 of 23 in the south. Its median is $+0.034$ of the unsigned current over the same pixels. Added up over the whole region, confident or not, the net is upward in only 16 of the 41 (median -0.006); over the whole map, in 19 (median -0.003). So each region as a whole closes its current, and the upward net lives in the confident pixels, its return running through the pixels whose arrowheads the instrument was not sure of (Figure 13, A). The bookkeeping holds too: over the 68 maps of the blind moments, taking out the turning moves the net over the whole map by a median of 0.028 of what it moves the net over confident pixels, as Stokes' theorem requires.




Figure 13

Figure 13. *The current makes a loop.* The 41 maps of the blind moments near disk centre. (A) Each map's net current, as a share of the unsigned current over confident pixels, added up over confident pixels only, over the whole region, and over the whole map; bars are medians. Only the confident pixels carry a net upward current. (B) Over confident pixels, the current of one sign in both polarities in each band of field strength (medians), as measured and after taking out three turnings with the field's sign.

Which turning accounts for it. Taking out a turning should shrink that upward net. The 5.5-degree turning fitted on the first two regions, full from 250 gauss, does not: it flips the net, leaving it upward in only 9 of 41 maps (median -0.029), and it puts a same-sign current of $+0.029$ into the weak band below 250 gauss, where the maps show none (Figure 13, B). The same angle reaching full size only at 1,000 gauss does shrink the net, in 29 of 41 maps (a chance of 0.0058), but leaves a same-sign current of $+0.031$ in medium field, 250 to 1,000 gauss. The edge of the confidence mask was tested as that current's source and is not: after taking out the turning, the medium-field offset was larger at the mask's edge than inside it in only 13 of 28 maps, a chance of 0.71. The smaller turning fitted on every near-centre map in hand, 2.97 degrees full from 250 gauss, leaves every band near zero: -0.010 of the unsigned current in medium field, $+0.011$ in weak field and $+0.001$ in strong field. It leaves the net near zero too, upward in 19 of 41 maps with a median of -0.003 . So the medium-field current left by the larger angle was most likely the larger angle overshooting, not a second effect. That was a fit on these same maps, so the second blind test fixed it as a test on new ones. On 82 fresh regions near centre, taking out 2.97 degrees shrank the net in 60, a chance of 0.00002, as predicted. It did not take all of it away: the median net fell from $+0.036$ to $+0.009$ of the unsigned current, and stayed upward in 53 of the 82. So on regions never used in any fit, the turning accounts for about three quarters of the upward net, and a quarter remains.

The published numbers carry it. The SHARP files carry summary numbers in their headers, computed by the pipeline "using the high-confidence values", the pixels whose arrowheads it is sure of, for studies of flare forecasting (Hoeksema and co-workers 2014, appendix C; Bobra and co-workers 2014). One of them, MEANJZD, is the mean upward current density over those pixels. In the 41 maps near disk centre it is positive in 30, the same 30 whose net is upward here, and its sign agrees with the net here in all 41. Its median is $+0.27$ mA/m². So the vertical loop's upward half is in the published number. Anyone who uses MEANJZD near disk centre reads that part with it, and any study that sums net currents over confident pixels would meet it too. This was read from the headers after the net had been measured here, so the second blind test fixed it as a test. On its 82 regions near centre, MEANJZD was positive in 67 and agreed in sign with the net here in all 82, with a median of $+0.28$ mA/m². Read again afterwards, as description, the agreement is closer than sign: this paper's mean current over confident pixels reproduces MEANJZD to a ratio of 1.000 in every one of the 82, and its unsigned total reproduces the keyword TOTUSJZ to 1.000, so the pipeline's keywords are the same arithmetic on the same pixels. The two other totals were then read the same way. Taking out the 2.97-degree turning changes the unsigned total by a median of -0.29 per cent, as a quantity with no sign should barely notice a turning that flips sign with polarity. The keyword SAVNCP, the sum of the sizes of the two polarities' net currents, is different: the turning adds a same-sign current to both polarities, so the sum of sizes changes wherever that part outweighs the Sun's own twist in one of them, and taking the turning out changes this paper's

polarity sum (a fixed 0.256 of SAVNCPP, a constant factor between the two definitions that is not explained here) by a median of -16.5 per cent, shrinking it in 56 of the 82 and growing it in 26. So on the 82 moment-regions SAVNCPP carried the turning by about a sixth of its size (run SUN_CURRENT_KEYWORDS 20260929_042026, exploratory). That reading was then fixed as a test on six regions followed across the disk, and it failed: the polarity sum shrank in 9 of 38 near-centre maps and grew in 29 (section 11). SAVNCPP is not a stable carrier of the turning, and the paper makes no general claim about it; what holds on both sets is that MEANJZD and TOTUSJZ are this paper's arithmetic to 1.000, on maps from 2011 to 2024. Many do: a 2024 review of the keywords' use lists SAVNCPP and TOTUSJZ among the best-performing predictors in the flare-forecasting literature, across more than fifty studies built on these keywords (Kosovich and co-workers 2024, section 4).

What it means. The upward current near disk centre that made “every spot’s current flows up” is one half of a loop that the instrument’s own confidence cuts in two. Its strong-field part is the instrument’s turning. Its return runs through pixels the instrument sets aside, so a sum over what it keeps comes out upward, in both hemispheres, and so does the pipeline’s own mean. The confident pixels are a window, and the upward net was a finding only until the window was widened to the whole region, where it vanished; the author’s method note on numerical coincidences (Schulz 2026a) states that rule, that a window is a finding until you widen it and the width at which the answer stops moving is itself the information, and section 9 is that rule met on the Sun. On regions never used in any fit, a turning of about 3 degrees accounts for about three quarters of it; the last quarter is open (section 12).

10. The second blind test

Findings 2, 3 and 7 were measured on maps that had already been read. So before any new map was fetched, a second prediction was written and its fingerprint published, with two supplements added while the maps were still downloading. It fixed the maps: every region with a NOAA number at 00:00 TAI on the first of the twenty-four months of 2012 to 2014 that the first test did not use. It left out every region already read. It fixed the arithmetic, the same as before, and six tests. The fetch of 1 September 2014 failed, and asked again, Stanford returned no records for it, so by the prediction’s rule that moment is reported empty and not replaced. The other twenty-three gave 158 records; 15 had no data yet, 143 were read, and 94 had enough strong field to need a turning. The tool that reads them was written, and its fingerprint published, before the maps arrived, and it was run once.

test, fixed before the fetch	result	verdict
P1. Near centre, the turning needed is positive	45 of 50; chance about 2 in a billion	passes
P2. Near centre, the median turning is below 4.24 degrees	+3.24 degrees	passes
P3. Near centre, taking out 2.97 degrees shrinks the net current over confident pixels	60 of 82; chance 0.00002	passes
P4. Far out, large regions (two strong bands) need a negative turning	3 of 6; chance 0.66	fails
P5. Near centre, positive under each version of the inversion code	older 17 of 18, newer 28 of 32	passes, both
P6. Near centre, SHARP’s published MEANJZD is positive	67 of 82; same sign as the net in all 82	passes

Chances are for at least that many of one sign if each sign were a coin flip. Near centre means μ at least 0.75; far out, μ at most 0.40.




Figure 14

Figure 14. *The second blind test.* (A) The turning each of the 94 fresh regions needs, against its viewing angle, coloured by the version of the inversion code that made it; squares are the far regions with two strong bands, which P4 tested. The dashed line is the 2.97 degrees fitted before the maps were fetched. (B) Near centre, each region's net current over confident pixels as measured (dots) and after taking out 2.97 degrees (crosses), sorted by the measured net. (C) SHARP's own MEANJZD against the same net: the published number carries it, sign for sign.

What the tests said, in the words written before the fetch. The positive turning near disk centre is confirmed on a second set of fresh regions. Its size is about 3 degrees, not the 5.5 fitted on the first two regions, which were larger than most. The upward net over confident pixels near disk centre is the turning's own loop, cut by the confidence mask. The turning appears under both versions of the inversion. The published mean current carries its upward part. And large regions far out do not need a negative turning, so the limb carries no sign change of this turning.

Beside the verdicts, as description. The turning fades toward the edge of the disk rather than crossing. Its median is +4.26 degrees for the 22 regions with μ of 0.9 or more, then +2.26, +0.78, -0.55 and -0.61 in bands of μ down to the edge. Section 6 reads the same maps the instrument's way and finds the fade to be what a steady turning about the line of sight looks like when read about the vertical. And taking it out leaves a quarter of the upward net near centre, as section 9 says.

11. The third and fourth blind tests: regions followed across the disk

Section 6's frame, section 9's SAVNCPD reading and the last quarter of the net were all read on maps already in hand. So a third prediction was written and fingerprinted (SUN_SPOKE_PREDICTION_3_2026-09-29, SHA-256 first 16 c7abbfc8079ddb89) before six regions the paper had never read were fetched as daily tracks across the disk, chosen by name from years the blind moments never touched: NOAA 11302 (2011), 12371 (2015), 12529 (2016), 12673 (2017), and two cycle-25 regions, 13664 and 13842 (2024), whose leading polarity is the reverse of cycle 24's. The tools were the ones already run, unchanged, and the rules were the second blind test's: confident pixels, near centre at μ 0.75 and above, the 2.97-degree turning full from 250 gauss. 85 records arrived, every one complete, five with no confident pixels at a limb and skipped.

test, fixed before the fetch	result	verdict
1. The frame: in each mid-disk bin of the dead-on ruler, kind L exceeds kind V by at least 2 degrees and is positive	μ 0.40 to 0.60: kind V +0.89, kind L +1.49 (10 maps); μ 0.60 to 0.75: kind V -0.34, kind L +2.18 (15 maps)	undecided: the line-of-sight kind wins both bins, by 0.60 and 2.52
2. The sign under the other cycle: positive by the line-of-sight field in more than two thirds of near-centre maps, and in more than half of the cycle-25 ones	18 of 27, two thirds exactly; cycle 25 (NOAA 13664, the one that reached centre) 4 of 6	undecided, by no maps at all
3. Near centre, taking out 2.97 degrees shrinks the net over confident pixels, at a chance below 0.05	23 of 38; chance 0.13	undecided
4. Gate: MEANJZD and TOTUSJZ reproduced to 1.000; then the polarity sum shrinks in more than half	gate held, 38 of 38; the sum shrank in 9 of 38 and grew in 29	gate passes; test fails

Runs *SUN_TURNING_DEADON*, *SUN_TURNING_BLIND* and *SUN_CURRENT_KEYWORDS* 20260929_192649. What each answer would mean was written before the run; the verdicts are applied as written.




Figure 15

Figure 15. *The third blind test.* Every map of the six regions: the turning it needs against its viewing angle, read by Br (the paper's reading) and by the field along our line of sight, with the medians in the blind tests' bins of mu. The same spot read day by day across the disk is a track through the figure.

What the tests said. Nothing the prediction asked for failed except the keyword, and nothing passed at the margins it asked for. The frame came out as before in sign, the line-of-sight kind winning both mid-disk bins, but at half the size of the first reading, 1.49 and 2.18 degrees against 3.99 and 4.02, so the frame stays a description read twice. The cycle-25 region that reached centre turned the same way as cycle 24's regions, 4 of 6, which is what a turning made in the instrument's frame should do; across all six regions the count sat on the line. The net shrank under the turning in 23 of 38 maps, against 60 of 82 in the second blind test, and the reason is in the angles: these tracked regions need a median of +1.76 degrees near centre, not 3, so taking out 2.97 overshoots in some of them (in NOAA 13842, whose maps need about +1 degree, the mean current over confident pixels goes from between -0.02 and $+0.16$ mA/m² to between -0.04 and -0.19 on every near-centre day). And SAVNCP, which shrank under the turning in 56 of 82 moment-regions, grew in 29 of these 38 maps. Its gate is the finding to keep: the pipeline's MEANJZD and TOTUSJZ are this paper's arithmetic to 1.000 on maps of 2011 to 2024, across three revisions of the inversion code (the 2011 maps name "vfisv FD10 2012 Jun. 23"; the maps to 13 April 2016, "vfisv FD10 2013 Apr. 30"; the maps from 14 April 2016 on, "vfisvcombine FD10 2013 Apr. 30", the day seven maps of NOAA 12529 begin to carry QUALITY 128, a bit that is not the calibration bit and is not explained here).

Beside the verdicts, as description. The turning is not one number. Two regions carry it plainly, NOAA 12529 at +1.8 to +4.8 through centre and NOAA 12673 at +7.6 to +11.1, the largest angles read anywhere in this work. One reads negative through its whole crossing, NOAA 12371, at -1.1 to -2.6 near centre by either reading. And two change sign from east to west: NOAA 13664 needs +4.9, +4.7 and +6.1 degrees on the east side of the meridian and +1.4, -2.7 and -3.9 on the west by Br, its line-of-sight reading following (+9.2, +8.4, +6.1, then +0.7, -5.0 , -9.1); NOAA 13842 needs +2.4, +1.7 and +2.1 east and about zero to -1.8 west. The daily tracks of NOAA 11158 and 11166 in 2011 showed no such change of sign across the meridian (section 4), and the blind moments, one map per region, could not have shown it. It is a wall with numbers (section 13) and an open question (section 12): whatever turns the arrows in these two regions knows which side of the meridian the region is on, which a turning fixed in the instrument's frame does not, and the Sun's own twist, fixed in the Sun's frame, does not either.

The fourth blind test. Two questions followed and both could be run, so a fourth prediction was fixed (SUN_SPOKE_PREDICTION_4_2026-09-29, SHA-256 first 16 5c975c85095e3f93) before ten more regions never read were fetched as daily tracks, five from 2015 to 2018 and five from 2021 to 2024, and judged on the author's computer by a tool written and checked before the fetch (sun_turning_tracks.py, run SUN_TURNING_TRACKS 20260929_135757; nine regions arrived; Stanford matched no record for NOAA 12567 in the window the cases file asked, and NOAA 12644's window began after it had crossed the meridian, both faults of the cases file's dates and not of the archive, run SUN_FETCH_SHARP 20260929_133821). The first question was whether the change of sign across the meridian is common: a region changes sign if its east and west medians at mid-disk have opposite signs and are each at least 1 degree; a quarter or fewer of the judged regions changing sign would pass for the instrument's frame, half or more for an east-west effect, and the rule asked for six judged regions. Five could be judged and none changed sign (NOAA 12297, 12699, 12887, 13590 and 13697); three were too small to say. **Undecided, one region short,** and as description the meridian sign change stays two regions in fifteen. The second question was whether a region's own turning, fixed from its first near-centre map east of the meridian, accounts for its net on its other near-centre maps better than the one number 2.97 degrees: the own angle shrank the net's size in 19 of 37 maps, a chance of 0.5. **Undecided.** The description beside it is the finding: 2.97 degrees shrank the net in 15 of 37, a chance of 0.91, so on these regions the one number does worse than a coin, and the median net as measured on their near-centre maps was +0.001 of the unsigned current, with the net upward in 26 of 47, a chance of 0.28. The sign of the turning held: 36 of the 47 near-centre maps with an angle are positive (chance 0.0002), median about +2.3 degrees. One region is the exception, and it is large: NOAA 12297 of March 2015 needs -10 to -38 degrees near centre, the largest angles of either sign in this work, while its net over confident pixels is upward on every near-centre day; it is left as description for a specialist to look at.

A fifth prediction, fixed before its fit and run on the same fifteen tracks (SUN_SPOKE_PREDICTION_5_2026-09-29, SHA-256 first 16 ba92f14ec479a7a8; run SUN_TURNING_TRACKS 20260929_142242), asked whether three maps fix a region's angle where one did not: the median of a region's first three near-centre angles, applied to its later near-centre maps, shrank the net in 25 of 39, a chance of 0.054 against the 0.05 the rule asked. **Undecided, by four thousandths.** Three maps do better than one (25 of 39 against 19 of 37) and better than the one number (21 of 39), and not well enough to call the size fixed; within a region the angle moves from one day to the next by more than itself (NOAA 12297: -37.6 , -15.5 , -10.3 on three consecutive days; 13697: +0.2, +6.7, +11.1), on the same strong bands read the same way. Read over all fifteen tracked regions as description, nine can be judged for the meridian sign and one changes it.

Taken together, the tracked tests bound the turning from both sides. Its sign near centre survives every test, 45 of 50, 25 of 36 and 36 of 47. Its size is a per-region quantity that nothing tried here pins: 2.97 degrees, fitted on two regions of 2011 and confirmed on 82 moment-regions, does worse than a coin on nine tracked regions; one early map is a coin; three maps are 25 of 39. And the upward net over confident pixels, 30 of 41 and 67 of 82 on the moment-regions, is 53 of 85 on the fifteen tracked regions (chance 0.014), which are larger on average; the regions where the net runs downward through the whole crossing are among the largest.

12. Open questions, and which can be run

- *The last quarter of the upward net.* After taking out 2.97 degrees, the net over confident pixels near centre is still upward in 53 of 82 fresh regions (median +0.009 of the unsigned current). It may be the turning's shape, which the fit pinned poorly, or a second effect, or the Sun. The shape was refitted on those 82 maps, as description (SUN_TURNING_SHAPE 20260929_173012): the ten ramps of the family all fit within 2 per cent of one another, the rule choosing the shortest, 0 to 250 gauss at +3.32 degrees, and under it the net falls from +0.036 to +0.006 of the unsigned current, upward in 48 of 82, but the weak band, which carried none of the same-sign current as measured, now carries +0.017 of it; a 5.5-degree turning drives the net below zero, upward in 31 of 82, and puts +0.032 into the weak band. No ramp of this family removes the last quarter without putting same-sign current where the maps show none, and these maps do not pin the shape between ramps. Either the turning's shape lies outside the family of ramps, or the last quarter is not the turning's. On the six tracked regions of section 11 the net was smaller to begin with (median +0.012 of the unsigned current) and 2.97 degrees overshot it in the regions that need less, so on those maps the question changes shape: not a last quarter left over, but a turning whose size differs from region to region. **Open; refitted, not resolved.**
- *The turning's frame, tested.* Section 6 finds, as description, that the maps at mid-disk carry about 4 degrees of turning about the line of sight and about none about the vertical, and that this accounts for the fade. The third blind test read it again on six regions never read before (section 11): the line-of-sight kind won both mid-disk bins, by 2.52 degrees in one and 0.60 in the other, against a demanded 2 in both. Same sign twice, the margin met once. A test with more maps at mid-disk, where the ruler is well conditioned, would settle the margin; that means daily tracks of more regions, which is a fetch of the kind just made. **Undecided; can be run again.**
- *The sign that changes across the meridian.* Two of the six tracked regions, both of 2024, need a positive turning on the east side of the disk and a negative one on the west (section 11), which neither a turning in the instrument's frame nor one in the Sun's frame does, and which the 2011 tracks did not show. The regions are the largest in this work; whether the change belongs to size, to the year, or to something in the pipeline's later revisions ('vfvsvcombine', from April 2016) can be asked of more tracks, chosen by size and year before the fetch. The fourth blind test asked it of ten more regions and judged five, none of which changed sign, one region short of the six its rule demanded (section 11), so it stays two regions in fifteen. **Undecided; a rerun with more regions would settle it.**
- *Whether the turning's size is one number.* The tracked regions need a median of +1.76 degrees near centre, the moment-regions +3.24, and one region needs -1 to -3 (section 11). The fourth blind test fixed each region's size from its first near-centre east-side map and tested it on the rest of the track: it shrank the net in 19 of 37 maps, no better than a coin, and the one number 2.97 shrank it in 15 of 37, worse than a coin (section 11). The fifth blind test tried three maps: 25 of 39, a chance of 0.054, undecided by four thousandths (section 11). Within a region the angle moves day to day by more than itself, so either the turning follows the region's own evolution or a single map's strong-band reading is noisier than the population medians suggested; a test of that would read the same region at several hours of one day, where the region cannot have changed, which the orbit-day maps of NOAA 11166 already do for one region (a spread of 0.72 degrees over twelve maps, section 5) and daily tracks cannot. **Open; can be run on hourly tracks of more regions.**
- *NOAA 12297.* One large region of March 2015 needs -10 to -38 degrees near centre, negative where every other region is positive or near zero, with its net upward on every near-centre day (section 11). Nothing here explains it. **Where to point a specialist.**
- *What in the inversion turns the arrows one way from about 250 gauss to the strongest umbrae, without the sign change the Faraday term alone would give.* Section 6 leaves this with the instrument's team, with the two runs that can be repeated on any bench: the pipeline's own inversion on these cutouts against a textbook fit with the magneto-optical terms in full, and the same fit with the instrument's real filter profiles in place of Gaussians. **For the HMI team; the cutouts and the fits ship.**
- *An external witness.* A second instrument on the same region, a spectropolarimeter such as Hinode's, read for the same turning. **Cannot be run here; where to point a specialist.**

13. What did not work, and why

Seen from the side, a fixed turning in the instrument's own frame changes sign in the Sun's frame at about 62 degrees from disk centre, by arithmetic. On the four regions it was drawn from, the offset did fall toward the edge. On 26 fresh regions near centre and 7 far out, read blind, it did not change sign: 2 of 7 far regions were negative, a chance of 0.94. Those far regions were small, read in one band each, so the second blind test asked the question of large far regions, those filling two strong bands. Of 6, 3 needed a negative turning, a chance of 0.66. The turning fades toward the edge but does not cross, for small regions or large. Section 6's ruler says why the arithmetic was wrong: the turning is about the line of sight with the sign of the line-of-sight field, not about the vertical with the sign of Br, and a turning of that kind read by Br fades without crossing. The 5.5-degree turning fitted on the first two regions puts its return current below 250 gauss when it is taken out, where the maps show none; a smaller angle does not. And the leftover medium-field current under that angle is not concentrated at the edge of the instrument's confidence: 13 of 28 maps had more there than inside, a chance of 0.71.

The test of the light that could not be read. A test fixed before the light was fetched was to read the turning's cause from the light directly, core against wings. It stopped at its second gate in both regions, twice, because its premise, that the polarisation's direction at one filter position is the field's azimuth plus a constant, is false by the physics of the line: the Faraday terms turn the polarisation by an amount that depends on the field's strength and sign and changes sign across the line, so that no constant fits to 10 degrees. The gate did its job. What the test was for, the light's own word on the mechanism, was then got by description and by models, in section 6, and not by a test.

The whale. A one-layer fit whose Faraday terms are too weak makes a turning of the right size, 3.6 to 9 degrees, flipping with polarity, in both regions alike. Planted into the maps it fails on shape: it crosses zero near 1,000 gauss, at any plausible width of the line, and the maps' offset in the 1,000 to 1,500 gauss band is +2.53 mA/m² where the curve gives -1.31. A test of size alone would have passed it. In the companion book's words, this kill is a rename: from "the mechanism" to "part of the mechanism", the part that gives the size and the polarity flip.

The vane ruler's scatter. A second ruler for the turning's frame, reading each spot's field on the two sides of the line where the line-of-sight field changes sign, responds to planted turnings of the two kinds in the right order but scatters from -12 to +14 degrees across 59 readable spot-maps, with a bootstrap range on its median of -3.13 to +0.93. Read on hourly tracks with each rim sector its own control, it scatters wider still, from -31 to +17 degrees across 25 sectors, because a spot's own vanes turn by tens of degrees over the days it takes the seam to cross a sector; only the maps within hours of the crossing give a step of the ruler's size (+4.18, bootstrap +1.58 to +5.83), and that reading rests on 15 of 25 sectors. The Sun's own vanes are the noise, and they are louder than the turning by a factor of ten. It decides nothing at that size and is kept as exploration. Read again without the 18 hourly maps that carry the pipeline's calibration flag, nothing moves: the local reading is +4.21 (bootstrap +1.87 to +5.83) on 24 sectors, the whole reading +1.24 (run SUN_TURNING_VANES 20260929_190750).

The sign that changes across the meridian. The third blind test's six tracked regions were meant to close the frame, and two of them opened something instead: NOAA 13664 needs +4.9, +4.7 and +6.1 degrees east of the meridian and +1.4, -2.7 and -3.9 west of it, and NOAA 13842 +2.4, +1.7 and +2.1 east and about zero to -1.8 west, by Br, with the line-of-sight reading following (section 11). The 2011 tracks of NOAA 11158 and 11166 are positive on both sides. A turning fixed in the instrument's frame, with the sign of the line-of-sight field, does not know which side of the meridian a region is on; the Sun's own twist, fixed in the Sun's frame, does not either; something that does is the satellite's daily motion, which section 5 found the turning does not ride, and the region's own evolution over the week it takes to cross. It is left as a wall with its numbers; the fourth blind test found no third region among nine more (five judged, none changing sign; section 11).

The shape, refitted. The family of ramps that section 5 fitted on two regions was refitted on the second blind test's 82 near-centre maps to see whether a better-pinned shape would take the last quarter of the upward net (section 12). It does not. The ten ramps fit the band offsets within 2 per cent of one another, so these maps do not choose among them, and the one the rule picks, 0 to 250 gauss at +3.32 degrees, trades the net for same-sign current in the weak band, +0.017 of the unsigned current where the maps as measured carry -0.005 (SUN_TURNING_SHAPE 20260929_173012). What was learned is a bound: whatever makes the last quarter is not a ramp of the turning's angle over field strength with the turning's sign.

14. Who came before

Nye and Berry (1974) found the zeros of waves, where amplitude vanishes and phase winds, in radio echoes from the bottom of the Antarctic ice sheet and in ultrasound. They found them this paper's way: the zero itself may be lost in noise, but a circuit through strong signal, counting the wavefronts crossed, shows that one is enclosed. They called it a Burgers circuit, after crystals. In two dimensions they point to the amphidromic points of the tides, after Whewell (1833, as they cite him). Radar interferometry met the half counts as residues (Goldstein, Zebker and Werner 1988), with Itoh's condition (1982) that unwrapping needs neighbouring steps under half a turn.

The same rule runs every quadrature encoder, the sensor that counts a turning shaft by two signals a quarter cycle apart. When the shaft moves half a cycle between two samples, both signals change at once and the direction cannot be told; decoders flag it. Texas Instruments' encoder module sets a phase error flag "when edge transition is detected simultaneously on the QEPA and QEPB signals", and Nordic Semiconductor's decoder counts such "double transitions" apart from the position. A move of more than half a cycle reads as a small move backwards, which the two signals cannot flag. What catches it is a closure: a third signal pulses once per revolution, and the count there must be the whole number of counts in one turn, or the module flags a position-counter error. The near-reversal steps of section 7 are the encoder's double transitions, and the counts held round a strong rim are its once-a-revolution check.

Counting a field's zeros with a topological index round loops is standard in the Sun's physics too (Longcope 2005, for the corona's topology), and the Sun's twist has been read on these same maps with winding-based quantities (Xiao, Prior and Yeates 2023). Currents round single sunspots were drawn with circles by Venkatakrisnan and Tiwari (2009). Net currents in active regions, and whether they come back, were studied by Georgoulis, Titov and Mikić (2012), Liu and co-workers (2017, in the *Astrophysical Journal Letters*) and Kontogiannis and co-workers (2017). The SHARP maps and their pipeline are described by Bobra and co-workers (2014), Hoeksema and co-workers (2014) and Centeno and co-workers (2014). The turning's nearest kin is Faraday rotation in a filter magnetograph (Gao and co-workers 2008). The light's mirror of section 6 has an older kin: West and Hagyard (1983) measured the Faraday rotation of a sunspot's transverse-field azimuth as a function of wavelength across the Fe I 5250 Å line, matched it with the analytical Stokes profiles by letting the line-formation parameter vary over the spot, and found it negligible below 1,800 gauss for their instrument (read in abstract). Whether net currents in a polarity come back has a literature of its own, which Sun and Cheung (2020) review and add to: Wheatland (2000) found the net current over a whole region consistent with zero while each polarity's was not, and Sun and Cheung show that a twisted tube meeting the surface at a slant projects a net current into each polarity, of opposite signs in the two, though the tube carries none. That is the polarity-odd part, this paper's twist; the polarity-even part, the offset, lies outside it by the same closed-loop argument they use (their equation 2), which is why section 4 can lay it at the instrument's door.

The light's side of section 6 rests on the standard theory of polarised radiative transfer in a magnetised atmosphere: the Unno-Rachkovsky solution for a Milne-Eddington atmosphere with the magneto-optical terms, as any modern account gives it, and the Voigt and Faraday-Voigt functions it needs. The pipeline's inversion is VFISV, described by Borrero and co-workers (2011) and, as run for HMI, by Centeno and co-workers (2014), which also describes the one-layer model, the fixed damping and the filling factor this paper's fitter copies. The apparent reversal of polarity in the wings of the line over off-centre sunspots, which the light's mirror of Figure 7 is a relative of, is a classic of the field, and its standard account is to be read and cited here before deposit. The reading of the light's polarisation as a carried arrow is the companion book's, which cites Berry (1984) for the geometric phase, Tomita and Chiao (1986) for its measurement on light in a coiled fibre, and Pancharatnam (1956) for its first appearance in polarised light. What was read at what depth is stated in the reference list.

Where this reading might be of use to them. The people whose work this paper leans on know these maps and this physics at every level, and this paper is written from outside their field by someone who does not. What it can offer is not a correction but a few readings taken from an odd angle, each with a tool that ships, in case one of them lands on a blind spot.

To anyone who sums the signed vertical current over a region's confident pixels near disk centre, whether as MEANJZD, as a net current per polarity, or as a feature for a forecast: that sum carries a part of one sign in both polarities, about +0.03 of the unsigned current on the moment-regions, that a sum over the whole region does not (section 9). The window makes it. The two-line check is to sum the current separately over the pixels where Br is positive and where it is negative and add the two: whatever the Sun does by projection or twist is odd under that mirror and cancels, and what survives is even and is not the Sun's (section 4). It costs nothing to run on any map, and it separates the two before any interpretation.

To those who study net currents and their neutralisation: the projection effect that Sun and Cheung describe is polarity-odd, and it explains why each polarity carries a net current of opposite sign. The polarity-even part is a different object, and this paper's reading is that it is the instrument's, made by a turning of the transverse-field azimuth of a few degrees whose sign follows the line-of-sight field. If that is right, a net current summed over confident pixels near centre carries an upward part of a few hundredths of the unsigned current, about three quarters of the upward net those sums show, and it fades toward the limb without changing sign (sections 6, 9 and 10). Liu and co-workers (2017) measure neutralisation as the ratio of direct to return current in each polarity, integrating the current of each sign separately over a chosen area, and report the two polarities' ratios apart: for NOAA 11158 before its X-class flare, 2.38 in one polarity and 3.48 in the other. A polarity-even part adds to the net of one polarity and takes the same amount from the other, so some of the difference between a region's two ratios is where such a part would show, and the mirror check above says how much of it is even before any area is chosen. The even part measured here is a few hundredths of the unsigned current, smaller than that difference, so it cannot be the whole of it; it is offered as one term to subtract, not as an account of the ratio. If the reading is wrong, the mirror test will say so on any set of maps, which is why the tool ships.

To the instrument’s team: three things are on the bench with the cutouts. HMI’s reported azimuths differ from a textbook Milne-Eddington fit of HMI’s own Stokes profiles by an S-shaped few degrees that flips with polarity and does not move under the pipeline’s weights or regularisation (section 6). A ruler that reads each map about the local vertical and about the line of sight at once finds the turning about the line of sight, twice, at 2 and 4 degrees (sections 6 and 11). And two large regions of 2024 turn one way east of the meridian and the other way west, which nothing here explains (section 11). Each may have a known cause; if so, a sentence naming it would close a section of this paper.

To anyone who wants a quantity from these maps that the azimuth systematics cannot move: a winding count round a strong rim is that quantity. It is a holonomy, a whole number the loop carries, and a smooth turning of a few degrees cannot change it except at the rare places where neighbouring arrows already point nearly opposite ways, which the count itself flags (sections 7 and 8). A current takes the turning in full; a count does not. In the encoder’s language above, the count is the once-a-revolution index pulse, and the current is the analogue reading between pulses.

15. Reproducing every number

Every number in this paper is printed in the ledger of a tool that ships with it. The tools read the maps and never change them. The maps are identified by fingerprints (the first 16 characters of the SHA-256 of their files, in a fixed order), printed in the ledgers. Figures 1 to 6 and 11 to 14 are drawn by one tool from the maps themselves, Figure 15 is the dead-on ruler’s own figure from its run on the six tracks, and each part of its ledger prints the numbers it draws beside those of the run of record; Figures 7 to 10 are drawn by a second tool from the ledgers of record of the runs they show, and its ledger prints every number it drew beside the run it took it from.

what	tool, version	run of record
Counts through the flares (section 3, Figure 3)	sun_half_count.py 2	SUN_HALF_COUNT 20260928_045400 (two flare series)
Offset and twist near centre; planted errors and the ramp scan (sections 4 and 5, Figures 4 and 6)	sun_offset_injection.py 2	SUN_OFFSET_INJECTION 20260927_132638
The offset across the disk (section 4)	sun_current_offset.py 2	SUN_CURRENT_OFFSET 20260927_114851 and 20260927_193239
The first blind test (section 4, Figure 5)	sun_offset_blind.py 1	SUN_OFFSET_BLIND 20260928_041724
The orbit day; counts under the turning on the disk tracks (sections 5 and 7)	sun_offset_removal.py 3	SUN_OFFSET_REMOVAL 20260927_193239
Blurring (section 5)	sun_offset_view.py 2	SUN_OFFSET_VIEW 20260928_025857
The two counts, classes, steps, near-reversals (sections 7 and 8, Figures 2, 11 and 12)	sun_half_count.py 2	SUN_HALF_COUNT 20260928_043627 (disk tracks, blind moments)
The half residues’ pairing (section 8), exploration	sun_residue_pairs.py 1	SUN_RESIDUE_PAIRS 20260929_214315 (hourly tracks; the same stamp’s run on the blind moments, and 20260929_214300 on the daily tracks, as description)
The vertical loop (section 9, Figure 13)	sun_current_budget.py 3	SUN_CURRENT_BUDGET 20260928_054447
The published totals MEANJZD, TOTUSJZ and SAVNCPD against the turning (section 9)	sun_current_keywords.py 1	SUN_CURRENT_KEYWORDS 20260929_042026
The turning’s size, fitted (sections 5 and 9); refitted on the second blind test’s maps (sections 12 and 13)	sun_turning_shape.py 1	SUN_TURNING_SHAPE 20260928_144658 and 20260929_173012 (the 82 near-centre maps, fingerprint ca22c0cf2af042e7 as in the keyword run; the tool’s slot for the disk tracks was empty on that run and prints as nan)
The second blind test (section 10, Figure 14)	sun_turning_blind.py 3	SUN_TURNING_BLIND 20260928_181524
The fourth and fifth blind tests (section 11)	sun_turning_tracks.py 2	SUN_TURNING_TRACKS 20260929_135757 (the fourth, version 1, on the nine tracks that arrived; prediction SUN_SPOKE_PREDICTION_4_2026-09-29,

		SHA-256 first 16 5c975c85095e3f93) and 20260929_142242 (the fifth, version 2, on all fifteen tracks; prediction SUN_SPOKE_PREDICTION_5_2026-09-29, ba92f14ec479a7a8), both run on the author's computer; version 2 reprints version 1's lines; the fidelity run here on the third test's six tracks, 20260929_200516, reproduces 80 of 80 angles and net shares of SUN_TURNING_BLIND 20260929_192649
The third blind test (section 11, Figure 15)	sun_turning_deadon.py 2, sun_turning_blind.py 3 and sun_current_keywords.py 1, unchanged	SUN_TURNING_DEADON, SUN_TURNING_BLIND and SUN_CURRENT_KEYWORDS 20260929_192649 (prediction SUN_SPOKE_PREDICTION_3_2026-09-29, SHA-256 first 16 c7abbfc8079ddb89)
Which codes made the maps (section 2)	sun_map_headers.py 1	SUN_MAP_HEADERS 20260928_155925
The turning's frame: the dead-on ruler (section 6, Figure 10)	sun_turning_deadon.py 2	SUN_TURNING_DEADON 20260928_215526
The vane ruler (sections 6 and 13)	sun_turning_vanes.py 3	SUN_TURNING_VANES 20260928_220608 (disk tracks and orbit day), 20260929_005842 (hourly tracks, sectors through time), 20260929_190750 (the same, then without the flagged maps; reprints version 2's run)
The test of the light that could not be read (sections 6 and 13)	sun_stokes_turning.py 2 and 3	SUN_STOKES_TURNING 20260928_194600 and 20260928_195923
The light's mirror, described (section 6)	sun_stokes_why.py 1	SUN_STOKES_WHY 20260928_202251
The light's mirror, modelled (section 6, Figure 7)	sun_stokes_forward.py 1	SUN_STOKES_FORWARD 20260928_205957
A fitter with weak Faraday terms; the width run (section 6, Figure 8)	sun_stokes_fit.py 2	SUN_STOKES_FIT 20260928_225848
The curves planted into the maps (section 6, Figure 8)	sun_turning_plant.py 1	SUN_TURNING_PLANT 20260928_225126
The reverse play (section 6, Figure 9)	sun_stokes_refit.py 3	SUN_STOKES_REFIT 20260929_051925 (all four Stokes, and the pipeline's weights and regularisation put in; its first parts reprint version 1's run 20260928_235356 and version 2's 20260929_022834)
Figures 1 to 6 and 11 to 14	sun_paper_figures.py 2	SUN_PAPER_FIGURES 20260928_182332
Figures 7 to 10	sun_paper_figures_b.py 3	SUN_PAPER_FIGURES_B 20260929_055346
Fetching the light from Stanford	sun_fetch_stokes.py 1 and 2	SUN_FETCH_STOKES 20260928_114916 (detector-frame inversions), 20260928_121558 (I, Q, U), 20260928_184127 (V)
Fetching the maps from Stanford	sun_fetch_sharp.py 3	SUN_FETCH_SHARP 20260927_113854, 20260927_192153 and 20260927_205224; the second blind test's fetch and its one-moment repeat, 20260928_111436; the hourly tracks, cases file SUN_FETCH_SHARP_CASES_DENSE 2026-09-28, fetched on the author's computer (two of its three cases arrived); the third blind test's six tracks, cases file SUN_FETCH_SHARP_CASES_BLIND3_2026-09-29, fetched on the author's computer, 20260929_110846 (all six cases complete); the fourth's ten, cases file SUN_FETCH_SHARP_CASES_BLIND4_2026-

The flare series are the companion book's maps, whose fingerprints match its runs of record. Earlier versions of the tools above are retired; each version reprints the runs of record of the one before it line for line. The files of Figures 11 to 14 are the files `sun_paper_figures.py` version 2 drew as `fig7` to `fig10` under an earlier numbering, renamed to their present numbers when section 6 was added and byte-identical to the run of record's; version 3 of that tool will draw them under the new names. A mistaken start of the width run, killed within a minute when a slip in its pre-written expectation was caught, is named in the shipping manifest and is not a run of record.

16. What a reader leaves with

This paper asked one plain question of the Sun's magnetic maps, whether a loop count is a label the instrument cannot move, and answered it, and on the way it found a second thing it had not gone looking for. Here is what to take away, in the order the paper found it, with the number that carries each.

The count holds. Round 2,941 loops through two X-class flares, not one count changed, and 103 of the 106 that counted a turn never had a rim step past a quarter turn (section 3). A count on a strong rim is a label on the real Sun, not only in the arithmetic.

Every spot's current seems to flow up, and that is the instrument. Near disk centre, the electric current in strong field has a part of the same sign in both polarities, which the Sun's own physics cannot make; it was predicted and found blind twice, 24 of 26 regions and then 45 of 50 (sections 4 and 10). A turning of the measured arrows of about 3 degrees, one way where the field points toward us and the other way where away, reproduces it, and the turning lives in the instrument's frame: about 4 degrees about the line of sight and about none about the Sun's vertical where the two can be told apart, which is why it seems to fade toward the limb (section 6).

The count reads through the turning; the current does not. Put the turning in on purpose and 31 of 154,267 counts move, all at five places where neighbouring arrows already point nearly opposite ways, while every current takes the turning in full (section 7). A second count that ignores the arrowheads has the opposite blind spot, and its half residues come in nearest-neighbour pairs of opposite sign, 0.725 against 0.500 shuffled (section 8).

The published numbers carry it. SHARP's MEANJZD is this paper's mean current to a ratio of 1.000, on maps of 2011 to 2024, and carries the turning's upward part; TOTUSJZ does not; SAVNCPP carried it on the 82 moment-regions and did not on six tracked regions, so no general claim is made for it (sections 9 and 11). Anyone who sums currents over confident pixels near disk centre meets it.

Where the turning is made was followed as far as public data go, and no further. The light itself carries a mirror of it, seen in 1983 by West and Hagyard and made by a one-layer model of the line. A one-layer fit whose Faraday terms are too weak gives the turning's size and its polarity flip but the wrong shape across field strength, at any width of the line. HMI's arrows disagree with a textbook fit of HMI's own light by an S-shaped few degrees, and putting the pipeline's own weights and regularisation into that fit moves nothing. Two conventions were found and taken out on the way, a quarter turn in Q and U and the sign of V (section 6).

Two more blind tests, on fifteen regions followed day by day across the disk, did not close what they were written to close and bounded the turning instead. Its sign near centre survives every test, 45 of 50, 25 of 36 and 36 of 47 maps positive. Its size is not one number: 2.97 degrees does worse than a coin on nine tracked regions, one early map is a coin, and three maps are 25 of 39. The frame kept its sign at half the margin; the cycle-25 regions turn the same way as cycle 24's; SAVNCPP failed its test; two regions of fifteen change the turning's sign across the meridian; and one large region of 2015 turns the other way by tens of degrees (section 11).

The walls, each with its number. The maps' offset is one-signed from at least 1,000 gauss up, +2.53 mA/m² where the Faraday model gives -1.31 (section 6). Two readers of one light are not a witness, so which side of the S-shape the Sun is on needs a second instrument (section 6). The vane ruler's own vanes turn by tens of degrees over the days a seam takes to cross a sector, ten times the turning (section 13). The 62-degree sign change the geometry predicted never came, twice (section 13). A quarter of the upward net near centre is not accounted for by a 3-degree turning, and no ramp of the turning over field strength takes it without putting current of the wrong sign into the weak band (sections 12 and 13). The turning's size is not one number: +3.24 degrees median on the moment-regions, about +2 on the tracked ones, -10 to -38 in NOAA 12297, and two regions of fifteen change its sign across the meridian (sections 11 and 13).

What conflicts with what. HMI's arrows and a textbook fit of the same light; the dead-on ruler and the vane ruler on the turning's frame; the paper's prediction of a limb sign change and the maps. Each is stated with its numbers and left standing, because the paper's rule is that a declared verdict is kept and an exploration never overturns one.

What is here for someone else. Every number is printed by a tool that ships with the paper and reads public maps; the two predictions and their fingerprints ship with them; and the questions the paper cannot answer are written for the people who can: the HMI team, with the cutouts and the fits, and anyone with a spectropolarimeter on the same region.

References

Each entry ends with the depth at which the work was read for this paper: in full, in part, abstract, or as cited by another work. Records were checked on the publisher's or an index page on 29 September 2026 except where marked.

- Auer, L. H., House, L. L., and Heasley, J. N. (1977). Solar Physics 55, 47–61. *As cited by Borrero and co-workers (2011), as the first Milne-Eddington inversion to take the magneto-optical effects fully into account.*
- Berry, M. V. (1984). Quantal phase factors accompanying adiabatic changes. Proceedings of the Royal Society of London A 392, 45–57. DOI 10.1098/rspa.1984.0023. *Abstract, from the Royal Society's scan; as cited by the companion book.*
- Bobra, M. G., Sun, X., Hoeksema, J. T., Turmon, M., Liu, Y., Hayashi, K., Barnes, G., and Leka, K. D. (2014). The Helioseismic and Magnetic Imager (HMI) vector magnetic field pipeline: SHARPs – Space-Weather HMI Active Region Patches. Solar Physics 289, 3549–3578. DOI 10.1007/s11207-014-0529-3. *Abstract.*
- Borrero, J. M., Tomczyk, S., Kubo, M., Socas-Navarro, H., Schou, J., Couvidat, S., and Bogart, R. (2011). VFISV: Very Fast Inversion of the Stokes Vector for the Helioseismic and Magnetic Imager. Solar Physics 273, 267–293. DOI 10.1007/s11207-010-9515-6. *In full, from the arXiv version (0901.2702).*
- Centeno, R., Schou, J., Hayashi, K., Norton, A., Hoeksema, J. T., Liu, Y., Leka, K. D., and Barnes, G. (2014). The Helioseismic and Magnetic Imager (HMI) vector magnetic field pipeline: optimization of the spectral line inversion code. Solar Physics 289, 3531–3547. DOI 10.1007/s11207-014-0497-7. *In full, from the arXiv version (1403.3677); the passages quoted were checked against it.*
- Gao, Y., Su, J., Xu, H., and Zhang, H. (2008). Statistical removal of Faraday rotation in vector magnetograms taken by filter-type magnetographs. Monthly Notices of the Royal Astronomical Society 386, 1959–1965. DOI 10.1111/j.1365-2966.2008.12848.x. *Abstract and summary.*
- Georgoulis, M. K., Titov, V. S., and Mikić, Z. (2012). Non-neutralized electric current patterns in solar active regions: origin of the shear-generating Lorentz force. The Astrophysical Journal 761, 61. DOI 10.1088/0004-637X/761/1/61. *Abstract.*
- Goldstein, R. M., Zebker, H. A., and Werner, C. L. (1988). Satellite radar interferometry: two-dimensional phase unwrapping. Radio Science 23, 713–720. DOI 10.1029/RS023i004p00713. *Abstract; the title from the Wiley listing.*
- Hoeksema, J. T., Liu, Y., Hayashi, K., Sun, X., Schou, J., Couvidat, S., Norton, A., Bobra, M., Centeno, R., Leka, K. D., Barnes, G., and Turmon, M. (2014). The Helioseismic and Magnetic Imager (HMI) vector magnetic field pipeline: overview and performance. Solar Physics 289, 3483–3530. DOI 10.1007/s11207-014-0516-8. *In full, from the arXiv version (1404.1881); the passages quoted and their section numbers were checked against it; the page range from the journal PDF's first page.*
- Itoh, K. (1982). Analysis of the phase unwrapping algorithm. Applied Optics 21, 2470. DOI 10.1364/AO.21.002470. *Title and record.*
- Kontogiannis, I., Georgoulis, M. K., Park, S.-H., and Guerra, J. A. (2017). Non-neutralized electric currents in solar active regions and flare productivity. Solar Physics 292, 159. DOI 10.1007/s11207-017-1185-1. *Abstract.*
- Kosovich, P. A., Kosovichev, A. G., Sadykov, V. M., Kasapis, S., Kitiashvili, I. N., O'Keefe, P. M., Ali, A., Oria, V., Granovsky, S., Chong, C. J., and Nita, G. M. (2024). Time series of magnetic field parameters of merged MDI and HMI Space-weather Active Region Patches as potential tool for solar flare forecasting. The Astrophysical Journal 972, 169. DOI 10.3847/1538-4357/ad60c3. *In full, from the journal PDF.*
- Landolfi, M., and Landi Degl'Innocenti, E. (1982). Solar Physics 78, 355. *As cited by Borrero and co-workers (2011), for the Milne-Eddington solution of the transfer equation with the magneto-optical terms.*
- Landi Degl'Innocenti, E., and Landolfi, M. (2004). Polarization in Spectral Lines. Astrophysics and Space Science Library 307, Springer, Dordrecht. DOI 10.1007/1-4020-2415-0. *Record only; the standard account of polarised radiative transfer with the magneto-optical terms, not read here.*
- LaBonte, B. (2004). The Imaging Vector Magnetograph at Haleakala: III. Effects of instrumental scattered light on Stokes spectra. Solar Physics 221, 191–207. *As cited by Hoeksema and co-workers (2014).*
- Leka, K. D., Barnes, G., Crouch, A. D., Metcalf, T. R., Gary, G. A., Jing, J., and Liu, Y. (2009). Resolving the 180° ambiguity in solar vector magnetic field data: evaluating the effects of noise, spatial resolution, and method assumptions. Solar Physics 260, 83–108. DOI 10.1007/s11207-009-9440-8. *Abstract.*

- Liu, Y., Sun, X., Török, T., Titov, V. S., and Leake, J. E. (2017). Electric-current neutralization, magnetic shear, and eruptive activity in solar active regions. *The Astrophysical Journal Letters* 846, L6. DOI 10.3847/2041-8213/aa861e. *In full, from the journal PDF (6 pages), which carries the volume, article number and date on its first page.*
- Longcope, D. W. (2005). Topological methods for the analysis of solar magnetic fields. *Living Reviews in Solar Physics* 2, 7. DOI 10.12942/lrsp-2005-7. *Abstract.*
- Metcalf, T. R. (1994). Resolving the 180-degree ambiguity in vector magnetic field measurements: the “minimum” energy solution. *Solar Physics* 155, 235–242. DOI 10.1007/BF00680593. *Abstract.*
- Nye, J. F., and Berry, M. V. (1974). Dislocations in wave trains. *Proceedings of the Royal Society A* 336, 165–190. DOI 10.1098/rspa.1974.0012. *In full.*
- Pancharatnam, S. (1956). Generalized theory of interference, and its applications. *Proceedings of the Indian Academy of Sciences A* 44, 247–262. DOI 10.1007/BF03046050. *Abstract; as cited by the companion book.*
- Rudenko, G. V., and Dmitrienko, I. S. (2018). The presence of a systematic error in SDO/HMI data. *Solar-Terrestrial Physics* 4(2), 3–7. DOI 10.12737/stp-42201801. Also as arXiv 1711.08156 (2017), “Examination of artifact in vector magnetic field SDO/HMI measurements”. *Abstract and a summary of the arXiv text.*
- Rachkovsky, D. N. (1962). Magnetic rotation effects in spectral lines. *Izvestiya Krymskoi Astrofizicheskoi Observatorii. As cited by Centeno and co-workers (2014) and Hoeksema and co-workers (2014), whose two reference lists give different volumes and pages (28, 259 and 27, 148); not resolved here.*
- Schulz, M. (2026a). A Coincidence Has No Single Price: the ladder underneath a numerical coincidence, and what it costs to climb. Zenodo, concept DOI 10.5281/zenodo.22053824 (version 2 of 22 August 2026). *In full, from the author’s copy; the concept DOI resolved from the Zenodo API on 29 September 2026. Cited for its check zero, the window rule of section 9.*
- Schulz, M. (2026b). Loop Walking: A Field Guide to Hacking the Universe with Holonomy. Zenodo, concept DOI 10.5281/zenodo.21811588 (version of 22 September 2026). *The author’s own; the concept DOI resolved from the Zenodo API on 29 September 2026. Cited for its vocabulary (twist and writhe) and its rules for reading.*
- Sun, X., and Cheung, M. C. M. (2020). Non-neutralized electric current of active regions explained as a projection effect. *Solar Physics*, DOI 10.1007/s11207-020-01742-9; arXiv 2011.11873. *In full, from the arXiv version.*
- Tomita, A., and Chiao, R. Y. (1986). Observation of Berry’s topological phase by use of an optical fiber. *Physical Review Letters* 57, 937–940. DOI 10.1103/PhysRevLett.57.937. *Abstract; as cited by the companion book.*
- Venkatakrishnan, P., and Tiwari, S. K. (2009). On the absence of photospheric net currents in vector magnetograms of sunspots obtained from Hinode (SOT/SP). *The Astrophysical Journal* 706, L114–L119. DOI 10.1088/0004-637X/706/1/L114. *Abstract.*
- Unno, W. (1956). Line formation of a normal Zeeman triplet. *Publications of the Astronomical Society of Japan* 8, 108. *As cited by Centeno and co-workers (2014).*
- West, E. A., and Hagyard, M. J. (1983). Interpretation of vector magnetograph data including magneto-optic effects. *Solar Physics* 88, 51–64. DOI 10.1007/BF00196176. *Abstract.*
- Whewell, W. (1833). *As cited by Nye and Berry (1974).*
- Wheatland, M. S. (2000). Are electric currents in solar active regions neutralized? *The Astrophysical Journal* 532, 616. *As cited by Sun and Cheung (2020).*
- Xiao, D., Prior, C. B., and Yeates, A. R. (2023). Computation of winding-based magnetic helicity and magnetic winding density for SHARP magnetograms in spherical coordinates. *Solar Physics* 298, 116. DOI 10.1007/s11207-023-02211-9. *Abstract.*
- Texas Instruments (2008). TMS320x280x, 2801x, 2804x Enhanced Quadrature Encoder Pulse (eQEP) Module Reference Guide, literature number SPRU790D, November 2004, revised December 2008; section 3.1.1 (the phase error flag) and section 4.1.1 (the position-counter error flag). *In part; the phrase quoted in section 14 checked against section 3.1.1.*
- Nordic Semiconductor. nRF52840 Product Specification, the QDEC (quadrature decoder) chapter: the ACCDBL register, “accumulating the number of detected double transitions”, and the rule that “double transitions (SAMPLE = 2) will not be accumulated in the ACC register”. *In part; the phrase quoted in section 14 checked against the chapter.*

Declaration

Status. A preprint. It has not been peer reviewed. Its search of the literature is limited: the papers it cites were read at the depth each passage states, and a paper read only in abstract or summary is marked so. It is shared freely under a Creative Commons Attribution licence (CC BY 4.0).

How it was made. The author directed the work and proposed several of its ideas, among them the current’s vertical loop (section 9), the encoder reading of the near-reversal steps (section 14), and the question of whether the turning is a price the light pays to leave the Sun or an optical effect of its passage, which section 6 follows. The companion book’s Sun chapter supplied the flare series and the first readings of them. The tools were written and run, the figures drawn and the text drafted with Claude, an AI model made by Anthropic, working under the author’s direction. Every number was printed by a tool run on the maps, and a separate check of the draft’s numbers and quotations against those runs was made before this version. Responsibility for the correctness of every number, quotation and citation rests with the author.

Provenance. The maps are public data from NASA’s Solar Dynamics Observatory, fetched from Stanford’s Joint Science Operations Center, as is the light. Two predictions, and two supplements to the second, were written and fingerprinted before the maps they test were fetched, and a third, with two supplements, before the light; all are shipped with the tools. Every citation states the depth at which the work was read: in full, in part, or in abstract or summary.

Companion work. This paper and the companion book, *The Latch Atlas: Holonomy and Locks that Can’t Be Picked* (Matthew “Bear” Schulz, written with Claude, Anthropic), are deposited together, each naming the other; the book’s record is cited here by its concept identifier, which resolves to its latest version and is entered at deposit, once the book’s record exists.

Dual use. The work reads public solar data and describes a measurement effect in a public instrument’s products. No risk category of concern was identified.

Data: SDO/HMI SHARP maps and Stokes data, courtesy of NASA/SDO and the AIA, EVE, and HMI science teams.