S.T. Yau High School Science Award

Research Report

The Team

Name of team member:	Edison Yang
School:	Lambert High School
City, Country:	Suwanee, Georgia, United States
Name of supervising teacher:	Rebecca Howell
Job Title:	Teacher
School/Institution:	Lambert High School
City, Country:	Suwanee, Georgia, United States

Title of Research Report

Testing the Temperature Scaling Law of Solar Flares and Inferring an Effective Heating Law from Ribbon Measurements

Date

September 2026

Testing the Temperature Scaling Law of Solar Flares and Inferring an Effective Heating Law from Ribbon Measurements

Edison Yang
Lambert High School, Suwanee, Georgia, United States

Abstract

Solar and stellar flare models often relate the temperature of the heated plasma to the magnetic field strength and the size of the reconnecting loops. The widely used Shibata–Yokoyama scaling law predicts $T_{\rm p} \propto B^{6/7}L^{2/7}$, but previous solar tests have relied on indirect estimates of the magnetic field or the loop geometry. In this work we test the law directly, using magnetic fields and loop lengths measured from flare ribbons.

We analyse 3056 C- to X-class solar flares using RibbonDB and GOES soft X-ray observations. A multithread hydrodynamic forward model, processed through the same GOES temperature inversion as the observations, predicts observable field and length exponents of 0.89 and 0.19. The observations instead give a field exponent of 0.35 ± 0.03. For the complete M- and X-class sample with measured loop lengths, we find

$$T_{\rm p} \propto \bar{B}_{\rm rbn}^{0.27\pm0.04}\, L^{-0.01\pm0.04}.$$

The temperature therefore depends much more weakly on magnetic field than predicted, and shows no measurable dependence on loop length. The same loop lengths predict flare rise, duration, and decay times, which confirms that they carry real information about the flare geometry.

To determine the origin of this discrepancy, we derive the thermodynamic response of a heated loop by coupling conductive energy transport to chromospheric evaporation. In the regime of strong heating the result is $T_{\rm p} \propto H^{2/7}L^{4/7}$, with the heating duration cancelling exactly. Hydrodynamic simulations recover the predicted exponents to within 0.005, and a two-fluid extension also reproduces the electron-temperature response without fitted parameters.

Inverting this verified loop response gives an effective solar heating law

$$H_{\rm eff} \propto \bar{B}_{\rm rbn}^{0.92}\, L^{-1.85},$$

close to a BL^{-2} coronal-heating scaling and inconsistent with the B^3L^{-1}-type heating assumed by the standard flare law. Independent measurements show that the reconnection rate itself retains approximately the expected magnetic dependence. These results locate the discrepancy in how magnetic energy is converted into heat; the subsequent thermal response of the loop follows classical physics.

Keywords: solar flares; magnetic reconnection; flare ribbons; scaling laws; soft X-ray temperature; coronal magnetic field; loop length

Acknowledgement

This project grew from my interest in solar and stellar flares, and from a question about the Shibata–Yokoyama scaling laws: whether their predicted temperature dependence would hold when the magnetic field and the loop geometry were measured directly from flare ribbons.

I developed the research question and the analysis independently. I assembled and processed the RibbonDB [1] and SolarErupDB [2] data, analysed the GOES/XRS observations, constructed the ribbon-field and loop-length measurements from SDO/AIA and SDO/HMI images, implemented the statistical and hydrodynamic analyses, performed the numerical parameter sweeps, and developed the mathematical derivations presented in this report.

I thank Ms. Rebecca Howell, my physics teacher at Lambert High School, for serving as the supervising instructor for this competition and for supporting my study of physics.

This work used publicly available data from RibbonDB, SolarErupDB, NOAA/NCEI GOES/XRS, and NASA's Solar Dynamics Observatory, together with the open-source software cited in the report.

Use of artificial intelligence. No AI was used.

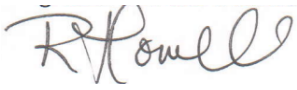
Commitments on Academic Honesty and Integrity

We hereby declare that we

1. are fully committed to the principle of honesty, integrity and fair play throughout the competition.
2. actually perform the research work ourselves and thus truly understand the content of the work.
3. observe the common standard of academic integrity adopted by most journals and degree theses.
4. have declared all the assistance and contribution we have received from any personnel, agency, institution, etc. for the research work.
5. undertake to avoid getting in touch with assessment panel members in a way that may lead to direct or indirect conflict of interest.
6. undertake to avoid any interaction with assessment panel members that would undermine the neutrality of the panel member and fairness of the assessment process.
7. observe the safety regulations of the laboratory(ies) where we conduct the experiment(s), if applicable.
8. observe all rules and regulations of the competition.
9. agree that the decision of YHSA is final in all matters related to the competition.

We understand and agree that failure to honour the above commitments may lead to disqualification from the competition and/or removal of reward, if applicable; that any unethical deeds, if found, will be disclosed to the school principal of team member(s) and relevant parties if deemed necessary; and that the decision of YHSA is final and no appeal will be accepted.

Name of team member: Edison Yang Signature: 

Name of supervising teacher: Rebecca Howell Signature: 

Contents

1 Introduction **1**
1.1 Motivation and Prior Tests . . . 1
1.2 Aim and Main Results . . . 2

2 Physical Hypothesis and Observable Predictions **2**
2.1 Derivation of the Shibata–Yokoyama Temperature Scaling . . . 2
2.2 From Structural to As-Observed Exponents . . . 3

3 Measurements from the Reconnection Geometry **3**
3.1 The Ribbon Magnetic Field . . . 3
3.2 Loop Length and Its External Validation . . . 4
3.3 GOES Temperature and the Two Flare Samples . . . 4

4 Forward Model and Analysis Design **5**
4.1 The Multithread Hydrodynamic Forward Model . . . 5
4.2 Errors-in-Variables Inference . . . 6
4.3 Analysis Protocol . . . 6

5 Results **6**
5.1 Forward-Model Recovery of the Predicted Law . . . 6
5.2 Field Dependence of the Peak Temperature . . . 7
5.3 The Joint Field–Length Test . . . 8
5.4 Loop Length, Flare Timescales, and Peak Temperature . . . 9

6 Nonlinear Loop Response and the Effective Heating Law **10**
6.1 Exact Conduction–Evaporation Reduction . . . 10
6.2 The Two-Fluid Correction and Its Parameter-Free Tests . . . 11
6.3 The Conductive Relaxation Spectrum . . . 12
6.4 Inverse Determination of the Heating Law . . . 12
6.5 Implications for Transport and Stellar Inference . . . 14
6.6 The Reconnection-Rate Test . . . 14

7 Reconnection Sequence and Flare Duration **14**

8 Limitations and Next Tests **15**

9 Conclusion **15**

A Supporting Methods and Robustness **16**
A.1 Thermal Inversion and Forward-Model Details . . . 16
A.2 Numerical Verification of the Conductive Theory . . . 17
A.3 Measurement Construction and Validation . . . 17
A.4 Statistical and Variable-Identification Robustness . . . 17

A.5 Robustness Tables 18

B The Conductive Relaxation Spectrum **19**

C Two-Fluid Partition Tests **20**

1 Introduction

1.1 Motivation and Prior Tests

Solar flares release magnetic free energy stored in the corona, producing hot plasma, radiation, bulk motion, and accelerated particles. Newly reconnected magnetic loops contain much of the soft X-ray emitting plasma, so their peak temperature provides a direct test of how a flare-heating model converts magnetic energy into heat. Large events release more than 10^{32} erg and reach tens of millions of kelvin, and they are the solar counterparts of the more energetic flares seen on other stars.

Flare ribbons give a direct view of the reconnection geometry. They are bright, elongated structures in the lower atmosphere that mark the footpoints of field lines which have recently reconnected, and the magnetic flux swept by their motion gives the reconnection flux $\Phi_{\rm rbn}$ [3]. RibbonDB shows that this flux is closely related to the peak soft X-ray flux measured by the Geostationary Operational Environmental Satellites (GOES), which defines the standard C, M, and X flare classes in order of increasing peak flux [1, 4]. Reconnection flux alone leaves that output underdetermined: at fixed $\Phi_{\rm rbn}$ the peak X-ray flux still varies by more than an order of magnitude, and it is essentially uncorrelated with the total magnetic flux of the host active region [1]. Relevant physics must accordingly depend not only on how much flux reconnects, but also on the magnetic field and the geometry of the reconnecting region.

The Shibata–Yokoyama scaling laws [5–7] make a specific prediction. Balancing reconnection heating against field-aligned thermal conduction [8], together with chromospheric evaporation, gives

$$T_{\rm p} \propto B^{6/7}\, n_0^{-1/7}\, L^{2/7}, \tag{1}$$

where B is the reconnecting coronal field, L is the loop half-length, and n_0 is the pre-flare density. Because the density dependence is weak, the field and length exponents provide a particularly direct test of the model. These same relations are also widely inverted to estimate magnetic fields and loop sizes for stellar flares that cannot be spatially resolved [7, 9].

Existing solar tests remain inconclusive. Namekata et al. [9] tested inverse relations using photospheric field proxies and lengths inferred from ribbon area. Their field dependence, with power-law indices of ≈ 0.37–0.46, was shallower than the model predicts, while their length relation held. Aschwanden [10] tested related scalings using nonlinear force-free field estimates but found that the available M- and X-class sample had too little dynamic range to settle it. Reep and Knizhnik [11] studied flux–energy relations using RibbonDB and GOES thermal parameters, but their volume estimate required an assumed relationship between ribbon area and loop length. These studies left open whether the disagreement came from the flare physics or from the geometric proxies used in the tests.

Ribbon observations make a more direct forward test possible. Our flux-weighted mean ribbon field

$$\bar{B}_{\rm rbn} = \frac{\Phi_{\rm rbn}}{S_{\rm rbn}} \tag{2}$$

provides an intensive measure of the field swept by reconnection, largely independent of overall flare size, and differs from the extensive flux $\Phi_{\rm rbn}$, which grows with total ribbon area. Separation of the two ribbons gives an estimate of the newly reconnected loop length. That same ribbon sequence also records the rate at which magnetic flux is reconnected, and with it the reconnection electric field, so the energy input can be measured without passing through the temperature law being tested. We use these quantities to test the predicted temperature law directly across C-, M-, and X-class flares, and then use the pattern of any disagreement to identify which part of the physical model fails.

1.2 Aim and Main Results

This study tests the Shibata–Yokoyama temperature law using magnetic field and loop length measured from the reconnection geometry itself. We compare the observations with a multi-thread hydrodynamic forward model passed through the same GOES temperature inversion as the data, and then use ribbon measurements of the reconnection rate, together with an analytical treatment of the loop thermodynamics, to locate where the observed scaling departs from the standard model.

The observed field dependence is far weaker than predicted, and the peak temperature is essentially independent of loop length. Across 3056 C- to X-class flares the field exponent is 0.35 ± 0.03 against a forward-model value of 0.89, and on the complete M- and X-class sample with measured loop lengths

$$T_{\rm p} \propto \bar{B}_{\rm rbn}^{0.27\pm0.04}\, L^{-0.01\pm0.04}, \tag{3}$$

where the model predicts 0.89 and 0.19. The same loop-length measurements predict the flare rise, duration, and decay times with exponents between 0.7 and 1.4, which shows that the null temperature dependence is physical and not a failed geometric proxy.

We then derive the loop response analytically and verify it numerically. Coupling the pressure driven by the heat input to the conduction-driven evaporation gives a system integrable in quadrature, in which the pulse duration cancels and

$$T_{\rm p} \propto H^{2/7} L^{4/7}, \tag{4}$$

with no leading dependence on heating duration. Simulations recover all three exponents to within 0.005, and extending the derivation to separate electron and ion temperatures gives a closed two-fluid correction that needs no fitted coefficient.

Inverting this response against the solar census yields an effective heating law

$$H_{\rm eff} \propto \bar{B}_{\rm rbn}^{0.92}\, L^{-1.85}, \tag{5}$$

near B/L^2, while the independently measured reconnection rate retains approximately the magnetic scaling an Alfvénic inflow would give. The loop thermodynamics therefore follow classical conduction–evaporation physics, and the discrepancy lies in the law converting magnetic energy into heat.

2 Physical Hypothesis and Observable Predictions

2.1 Derivation of the Shibata–Yokoyama Temperature Scaling

The Shibata–Yokoyama temperature law follows from a balance between reconnection heating and thermal conduction in a loop of half-length L [5, 6]. Reconnection delivers magnetic energy at the rate at which the field flows into the current sheet at the Alfvén speed, giving a heating rate per unit volume $H \propto B^3 n_0^{-1/2} L^{-1}$, which is also the heating function of the forward model (Section 4.1). Balancing it against field-aligned conduction, $C \simeq \kappa_0 T^{7/2}/L^2$, at the temperature peak gives

$$T^{7/2} \propto HL^2 \propto B^3 n_0^{-1/2} L, \qquad T \propto B^{6/7}\, L^{2/7}\, n_0^{-1/7}, \tag{6}$$

the exponents of the Shibata–Yokoyama temperature law. Appendix A.1 gives the intermediate steps. With the evaporation closure of Shibata and Yokoyama [7], and with B in gauss, n_0 in cm^{-3}, and L in cm, the normalized form is

$$T_{\rm p} \simeq 1.8 \times 10^4 B^{6/7} n_0^{-1/7} L^{2/7}\ \mathrm{K}. \tag{7}$$

where the coefficient is the evaporation-closed form (Appendix A.1).

Two features of the derivation matter for the test. The field exponent 6/7 is inherited from the B^3 dependence of the heating rate, so it is independent of the heating normalization, and the forward model recovers it for any efficiency (Section 5.1). The length exponent 2/7 is positive because a longer loop loses heat by conduction less efficiently: at fixed heating per unit volume, a loop twice as long reaches a temperature higher by $2^{2/7} \simeq 1.2$. Over the 1.9 dex of loop lengths in the sample for the joint test this predicts a factor $10^{1.9\times 2/7} \simeq 3.5$ in temperature, or $10^{1.9\times 0.19} \simeq 2.3$ after the GOES inversion, which is the contrast the joint test finds to be absent.

Corresponding laws for the emission measure depend strongly on the pre-flare density and are reproduced only poorly by the loop hydrodynamics used below, so no trustworthy target can be computed for them. We set them aside and test the temperature law, which depends on density only through $n_0^{-1/7}$.

2.2 From Structural to As-Observed Exponents

The exponents in Equation (7) describe single loop apex quantities. The GOES X-Ray Sensor (XRS) reports one response-weighted temperature, an isothermal summary, for what is in reality an arcade of loops at many temperatures, and the two-channel inversion compresses that distribution. Comparison must therefore be made against as-observed exponents: the structural values of Equation (7) after the arcade and the instrument have acted on them. Section 4.1 builds that transfer function from multithread loop models driven by the reconnection Poynting input, with a heating normalization fixed on a pilot sample and no reference to the measured thermal response, so the resulting targets are an independent standard for the data.

3 Measurements from the Reconnection Geometry

3.1 The Ribbon Magnetic Field

The mean field measured from the ribbons is

$$\bar{B}_{\rm rbn} = \frac{\Phi_{\rm rbn}}{S_{\rm rbn}}, \tag{8}$$

computed per magnetic polarity and averaged. Figure 1a shows the two ribbons and the underlying magnetogram for a representative M7.2 flare. We propagate the uncertainty from the dispersion across the cutoff family of the ribbon masks, and exclude events with flux imbalance greater than 20% from the full catalog.

Because $\bar{B}_{\rm rbn}$ is constructed from $\Phi_{\rm rbn}$ and $S_{\rm rbn}$, we also test whether the data demand the intensive ratio, by fitting $\Phi_{\rm rbn}$ and $S_{\rm rbn}$ with free coefficients and asking whether they concentrate on the ratio direction (Appendix A.4).

The field in Equation (7) is the coronal field reconnecting at the current sheet, the thin layer in the corona where reconnection occurs, and is of order tens of G, while $\bar{B}_{\rm rbn}$ is a flux-weighted photospheric mean of order hundreds of G. We relate them by $B_{\rm cor} = \kappa\, \bar{B}_{\rm rbn}$. A constant κ rescales the field and changes only the normalization of the law, leaving the exponents alone, so only a κ that varies systematically with flare class or active-region field concentration could flatten the measured slope. Our test is conditional on its approximate constancy, and Appendix A.4 examines the departures directly.

We assume no numerical conversion in the main analysis. Namekata et al. [9] adopted an explicit $B_{\rm obs} \simeq B/3$ reduction. We fix κ on a pilot subsample without tuning it to the thermal response, and drive the forward model of Section 2.2 with $\bar{B}_{\rm rbn}$, so the target and the measured slope use the same field coordinate.

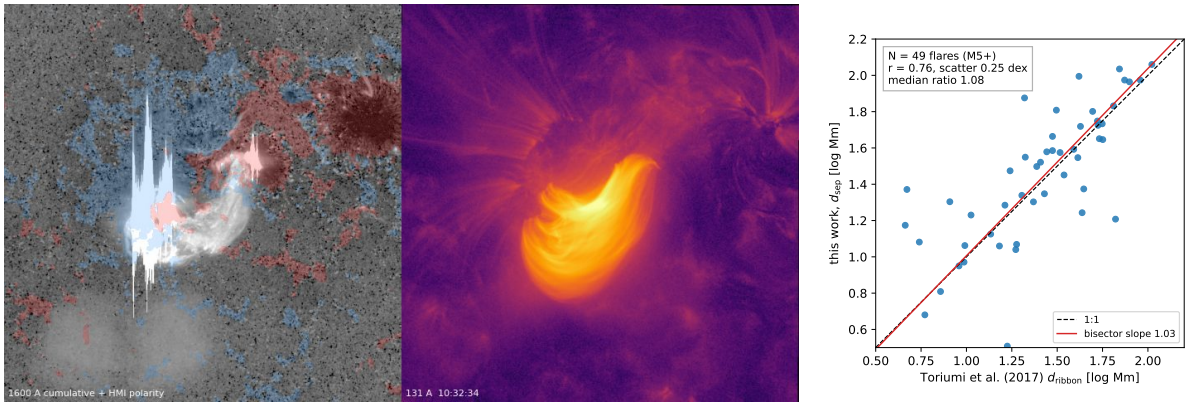


Figure 1: Both predictors are measured directly from the flare ribbons. Example for the M7.2 flare of 2014 January 7 (NOAA active region 11944). **(a)** AIA 1600 Å cumulative brightening over the flare window, showing the two flare ribbons, with the sign of the HMI line-of-sight magnetic field shaded (red positive, blue negative). The ribbons trace the footpoints of the reconnected loops; the flux they sweep divided by their area gives the mean field measured from the ribbons $\bar{B}_{\rm rbn}$, and the separation of the two ribbon centroids gives the loop half-length L. **(b)** the same field of view in AIA 131 Å during the decay phase, showing the hot post-reconnection loops that connect the two ribbons and whose emission sets the GOES peak temperature. **(c)** External validation of the loop-length estimator: our ribbon separation against the independently published ribbon distances of Toriumi et al. [12] for the 49 flares in common (dashed: one-to-one; red: bisector fit).

3.2 Loop Length and Its External Validation

The loop half-length is derived from the separation of the two flare ribbons, following the definition of Toriumi et al. [12]. For each event in the joint test we build a ribbon composite from AIA 1600 Å frames between 3 min before the GOES start and 3 min after the GOES peak, so that the composite comprises the footpoints formed by the temperature peak. We split it by the sign of the co-registered HMI line-of-sight field. Ribbon distance $d_{\rm sep}$ is then the great-circle separation of the two area-weighted centroids on the solar sphere, which removes foreshortening. Half-length follows from a semicircular loop,

$$L = \frac{\pi}{4} d_{\rm sep}. \tag{9}$$

Appendix A.3 gives the mask construction. We vary the brightness threshold over $k \in \{4, 6, 8, 12\}$ quiet-Sun levels. For each event we take the median over that family, and the dispersion across it (median 0.09 dex) is the threshold component of the error budget.

We validate the estimator externally. Toriumi et al. [12] published ribbon distances, measured by the same definition with independent software, for 51 flares of class M5.0 and above, 50 of which are in the sample for the joint test. On the 49 with a valid composite the two agree with a median ratio of 1.08, a bisector slope of 1.03, and a correlation of $r = 0.76$ (Figure 1c). Their scatter of 0.25 dex includes their measurement error as well as ours, and we adopt half of that variance, 0.18 dex, as the external error term per event, added in quadrature to the threshold term. Appendix A.3 reports the further comparisons and the cross-check estimators we use, which include cooling-time lengths [13–15] and footpoint-to-polarity-inversion-line lengths [16].

3.3 GOES Temperature and the Two Flare Samples

The primary catalog is RibbonDB [1]: 3137 flares of class C1.0 and greater observed within 45 degrees of central meridian, avoiding the strongly foreshortened regions near the solar limb, be-

tween 2010 April and 2016 April. It provides flare-ribbon masks from the 1600 Å channel of the Atmospheric Imaging Assembly (AIA) on the Solar Dynamics Observatory (SDO), the ribbon areas, and reconnection fluxes derived from the line-of-sight magnetograms of SDO's Helioseismic and Magnetic Imager (HMI). SolarErupDB [2] supplies vetted ribbon fluxes and eruptive or confined labels for the higher classes. Peak temperature and emission measure come from a background-subtracted two-channel GOES/XRS inversion [17] with the admissibility criteria of the Temperature and Emission measure-Based Background Subtraction method (TEBBS) [18]. Appendix A.1 gives the background treatment and the flux recalibration.

Our two predictors differ in measurement cost, so we use two samples. $\bar{B}_{\rm rbn}$ is catalog arithmetic and is available for the full catalog, whereas loop length requires measuring each event on the ribbon images. Our full-catalog test is the marginal test dominated by the field, and includes all RibbonDB events passing thermal quality control, the flux imbalance cut, and a background quality threshold fixed on a pilot subsample. Its slope of $\log T_{\rm p}$ on $\log \bar{B}_{\rm rbn}$ is a marginal coefficient, absorbing the effect of the unmeasured length (which may covary with $\bar{B}_{\rm rbn}$), so its target is the marginal slope from the forward model.

For the joint field–length test we use all M- and X-class events with measured loop lengths, plus a C-class sample stratified by class decade, rise-time tercile, and eruptivity (Appendix A.3). Population-level fits use inclusion weights so that the estimates refer to the parent catalog. Because every M- and X-class event is included, that subsample is a census with inclusion probability one, and its unweighted fit is reported alongside the weighted estimate (Section 5.3).

4 Forward Model and Analysis Design

4.1 The Multithread Hydrodynamic Forward Model

We produce the targets of Section 2.2 with a stochastic multithread-arcade model evaluated over the measured parameter ranges. Each strand follows the enthalpy-based thermal evolution of loops (EBTEL) formulation [19, 20] in its two-fluid ebtelplusplus implementation [21], driven by a single triangular heating pulse. This spatially averaged treatment reproduces the density and temperature evolution of resolved hydrodynamic models at far lower cost [22, 23]. Pulse amplitude is the reconnection Poynting input of the Shibata–Yokoyama balance,

$$H_0 = \epsilon \, \frac{B_{\rm cor}^2}{4\pi} \, \frac{v_{\rm A}}{L}, \qquad v_{\rm A} = \frac{B_{\rm cor}}{\sqrt{4\pi\rho}}, \qquad \rho = 1.4 \, m_{\rm p} \, n_0, \tag{10}$$

so that $H_0 \propto B_{\rm cor}^3 \, n_0^{-1/2} \, L^{-1}$, with $B_{\rm cor} = \kappa \bar{B}_{\rm rbn}$ (Section 3.1). Twelve strands with staggered onsets and log-normal amplitude jitter are summed into one differential emission measure, folded through the GOES-15 temperature-response functions [17], and inverted by the identical background-subtraction and two-channel pipeline applied to the data (Section 3.3). Our targets are the exponents of the recovered $\log T_{\rm p}$ on the input $(\log \bar{B}_{\rm rbn}, \log L, \log n_0)$ over a grid spanning the observed ranges, with the model scatter taken from the spread between stochastic realizations. Appendix A.1 gives the configuration.

Because the heating law is fixed before any comparison with the observed temperatures, the forward-model target is independent of the measured result. In Equation (10) the cube of the field is the reconnection Poynting scaling, set by the analytic balance alone, and the loop integration then yields $T_{\rm p} \propto B_{\rm cor}^{6/7}$ at fixed (L, n_0) without that exponent being imposed. Every configuration choice was fixed on a pilot subsample, and the grid was run before the thermal response was examined.

4.2 Errors-in-Variables Inference

The primary estimator is a multivariate errors-in-variables regression. Random measurement error in a predictor biases an ordinary fitted slope toward zero, an effect known as regression dilution, or attenuation. We correct for it with $\beta = (S_{xx} - D)^{-1} S_{xy}$ [24], generalized to the correlated $(\bar{B}_{\mathrm{rbn}}, L)$ predictors, where D is the mean measurement-error variance per predictor. Standard errors come from a block bootstrap over active regions, which keeps repeated flares from one active region grouped together. This estimator carries no latent predictor for each event and is thus free of the slope bias that a fully hierarchical latent-variable model incurs when the predictors are correlated. Theil–Sen [11] and ordinary least squares are reported alongside as uncorrected references, and agreement across the estimator set is required for a primary claim (Appendix A.4).

Three properties of the model bound the interpretation. Our temperature law depends on density only as $n_0^{-1/7}$, so the unmeasured pre-flare density leaves the test unaffected. Catalog selection operates on GOES flux while the truncation term acts in temperature space, so the flux-to-temperature mapping of the C1.0 floor is calibrated from the forward model independently of the data. A misplaced floor would bias slopes near threshold, where the C-class sample concentrates, and Section 5.2 shows the primary slope insensitive to it. Our marginal full-catalog field target also inherits the $(\bar{B}_{\mathrm{rbn}}, L)$ covariance, which is propagated as a band over its plausible range (Section 5.2).

4.3 Analysis Protocol

Principal selection criteria, forward-model targets, estimators, and decision rules were fixed on a pilot subset before the confirmation tests. The full-catalog field test and the joint field–length test were then evaluated on the remaining data. The subsequent hydrodynamic and timescale analyses were used to investigate the resulting discrepancy.

5 Results

5.1 Forward-Model Recovery of the Predicted Law

The forward model is independent of the measured thermal response, so its outputs are reported before the primary outcome (Section 2.2). Its peak temperatures span the C–X range. Across the model grid the structural exponents are recovered in the loops themselves: the field exponent comes back at 0.89 against the Shibata–Yokoyama 6/7, and the length exponent at 0.21, somewhat weaker than the structural $2/7 \simeq 0.286$.

Passing the same loops through the GOES/XRS synthesis and the identical inversion returns $b'_B = 0.89 \pm 0.01$ and $b'_L = 0.19 \pm 0.01$, with the spread between stochastic realizations setting the consistency band of 0.013 in log-slope. Those two exponents are the targets used throughout. Per-loop isothermal compression is appreciable ($T_{\mathrm{obs}}/T_{\mathrm{true}} \approx 0.93$) and nearly independent of B, so it shifts the normalization and leaves the field *slope* intact ($b'_B/b_B^{\mathrm{true}} \simeq 1.0$): the GOES inversion preserves the field exponent. Its length target falls somewhat below the structural 2/7, and follows from how the hydrodynamic model scales with loop length.

That result bears on the shortfall reported from proxy studies. A recovered field slope near 0.4 [9] is not produced by GOES compression of a true 6/7 law, even with a multithermal arcade, so the historical shallow slope reflects either the proxies or the scaling itself, and the tests below decide between them. Figure 2 shows the grid.

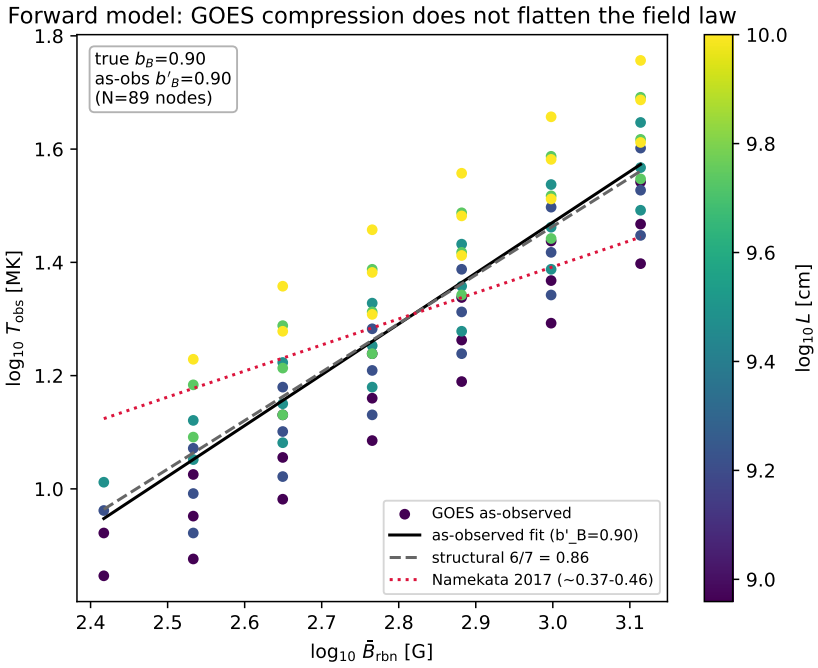


Figure 2: The GOES inversion preserves the predicted field exponent, so the observed shortfall is not instrumental. Forward-model calibration of the field target. Each point is one node of the hydrodynamic model grid: the temperature recovered by the GOES two-channel inversion from the synthetic light curve of a multithreaded loop arcade, against the input ribbon field, colored by the input loop half-length. The recovered slope (0.90, black) coincides with the structural 6/7 (dashed): the inversion preserves the field exponent. The proxy-based index of Namekata et al. [9] (dotted) is far shallower than either.

5.2 Field Dependence of the Peak Temperature

After the catalog cuts and the requirement of an admissible background-subtracted inversion, the sample comprises 3056 flares (2758 C, 279 M, 19 X) from 508 active regions. Peak temperature rises with the mean field measured from the ribbons. The two correlate at $r = 0.32$ in the logarithm, and across quintiles of ribbon field from 350 to 790 G the mean peak temperature climbs monotonically from 12.3 to 14.7 MK (millions of kelvin, Figure 3a). Even so, the field accounts for only $R^2 \simeq 0.10$ of the variance in $\log T_{\rm p}$, so field strength alone is a weak predictor of flare temperature.

Peak temperature responds to the intensive ratio. Fitting $\log \Phi_{\rm rbn}$ and $\log S_{\rm rbn}$ with free coefficients returns $+0.19$ and -0.14 (Figure 3b). Those opposite signs support a dependence on the intensive ratio $\Phi_{\rm rbn}/S_{\rm rbn}$ in place of either extensive quantity alone: temperature rises with reconnection flux and falls with ribbon area, as it must if it responds to the mean field. An extensive dependence repackaged in flux and area would instead carry coefficients of the same sign. Their magnitudes differ a little from the exact ratio direction, by an amount Appendix A.4 attributes partly to measurement attenuation.

We measure a field exponent of $b_B = 0.35 \pm 0.03$, far below the predicted 6/7. In other words, stronger ribbon fields do produce hotter flares, but much more weakly than the standard scaling requires. Correcting for measurement error and changing the regression estimator does not remove the discrepancy (Appendix A.4). Exclusion of the structural exponent holds even if the error budget is somewhat wrong, because the same field error sets both the attenuation correction and the rotation of the ratio direction, and the measured rotation bounds the cor-

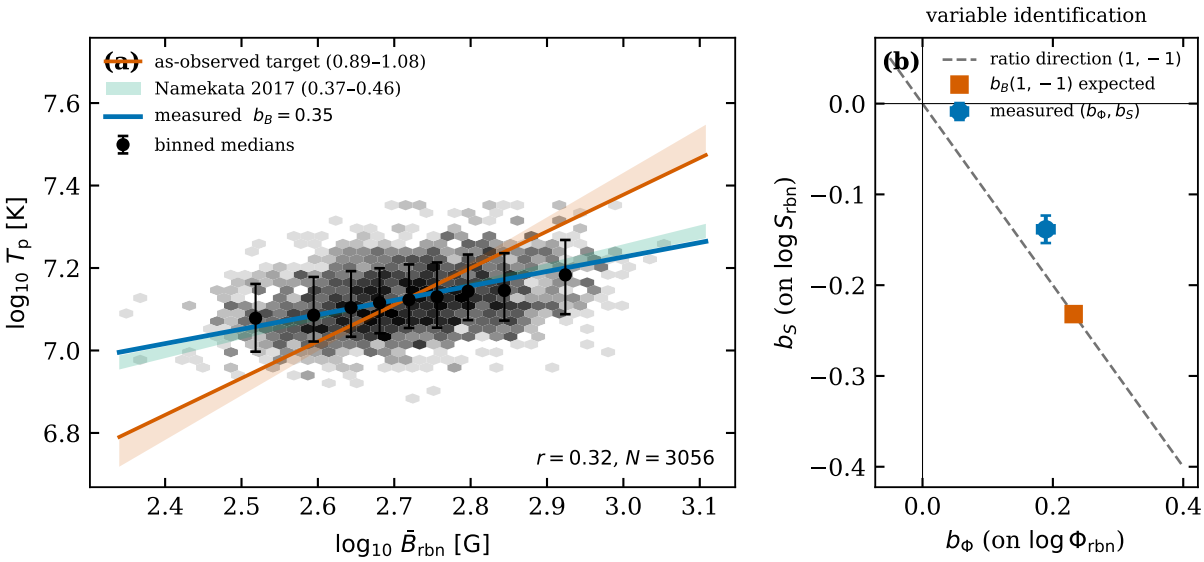


Figure 3: Peak temperature rises with ribbon field far more weakly than the standard law predicts. Full-catalog field–temperature test. **(a)** Peak temperature against the mean field measured from the ribbons for the 3056 flares (grey: log-density; black: binned medians with 16–84% ranges). The measured slope ($b_B = 0.35$, blue; error-corrected for the attenuation in $\bar{B}_{rbn}$) lies in the Namekata et al. (2017) proxy band (green) and far below the target (orange band, 0.89–1.08 over the plausible field–length covariance). All slope lines are pinned at the sample median, so the figure compares slopes alone. **(b)** Variable-identification test: the unrestricted coefficients of $\log T_p$ on $\log \Phi_{rbn}$ and $\log S_{rbn}$ (blue) are of opposite sign, the signature of dependence on the ratio $\bar{B}_{rbn} = \Phi_{rbn}/S_{rbn}$ alone, and lie near the ratio direction $(1, -1)$. The residual offset from the expected $b_B(1, -1)$ point (orange) is part measurement attenuation and part a small but significant departure from the exact ratio (Section 5.2).

rected slope at 0.6 (Appendix A.4). Our systematic range is therefore $b_B \in [0.35, 0.6]$, every value of which excludes 6/7.

Our forward model supplies the target. We compare the marginal slope on the full catalog with the field exponent $b'_B \simeq 0.89$ marginalized over the empirical field–length covariance, $b'_{B,\mathrm{marg}} = b'_B + \gamma\, b'_L$, where γ is the population regression of $\log L$ on $\log \bar{B}_{rbn}$ (Section 4.2). Our sample for the joint test measures $\gamma = 0.30 \pm 0.45$, and larger active regions host both stronger ribbon fields and longer loops, so the plausible range $\gamma \in [0, 1]$ places the marginal target in $[0.89, 1.08]$. No plausible covariance closes the gap: reaching 0.35 would need $\gamma \simeq -2.8$, contrary to the expected trend, and the measurement is excluded at $z > 15$ across the whole admissible range, including the field exponent at fixed loop length, 0.89, which needs no covariance assumption at all. Our measured value coincides instead with the index of 0.37–0.46 obtained by Namekata et al. [9], so the ribbon-measured field reproduces the departure from proportionality seen in the proxy era. We measure the length term directly in Section 5.3.

The shallow field dependence is stable across flare class, epoch, eruptivity, quality cuts, and progressively higher flux thresholds (Appendix A, Table 3).

5.3 The Joint Field–Length Test

The sample for the joint test comprises 422 events: all 280 M- and 19 X-class flares with an admissible thermal inversion, and 123 stratified C-class flares. Of these, 416 pass the measurement quality cuts, giving 277 M, 19 X, and 120 C flares from 147 active regions. Half-lengths span 1.9 dex with a median of 19.8 Mm, and the length error of each event is 0.20 dex. Inclusion

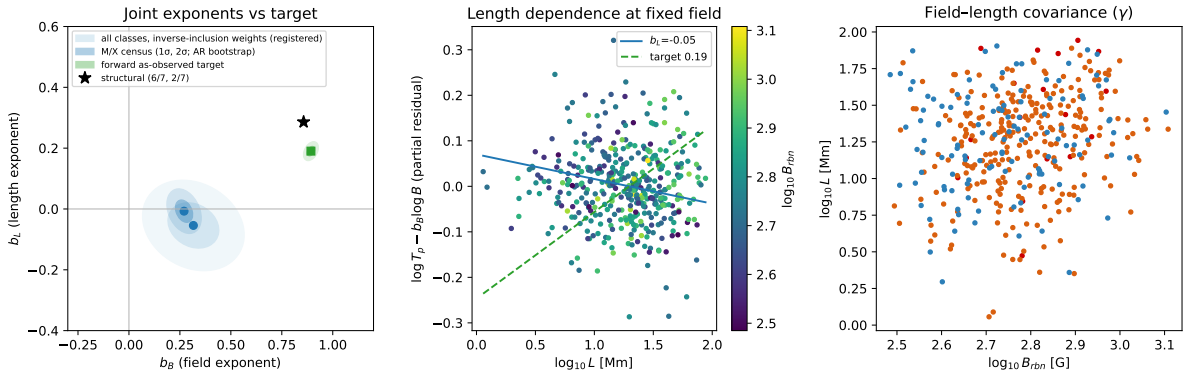


Figure 4: Loop length has no measurable effect on peak temperature at fixed ribbon field. Joint field–length test. **(a)** The measured exponents (b_B, b_L) for the weighted all-class fit (light blue, 1σ and 2σ from the active-region bootstrap) and the M/X census (blue) against the target (green) and the structural exponents (star). **(b)** Partial residual of $\log T_{\rm p}$ after removing the field term, against $\log L$, colored by $\log \bar{B}_{\rm rbn}$; the fitted length slope (blue) is flat, while the target slope (dashed green) is not. **(c)** The field–length covariance of the sample ($\gamma = 0.30 \pm 0.45$).

weights sum to the parent population of 2986 flares, but the effective sample size under them is 82 (Kish), so the M/X census carries the statistical precision and we report both.

Figure 4 presents the joint test. The fit of $\log T_{\rm p}$ on $(\log \bar{B}_{\rm rbn}, \log L)$, corrected for attenuation and using inclusion weights with a block bootstrap over active regions, returns $b_B = 0.32 \pm 0.13$ and $b_L = -0.05 \pm 0.08$. Against the target $(b'_B, b'_L) = (0.89 \pm 0.01, 0.19 \pm 0.01)$, the field exponent is displaced by 4.5σ and the length exponent by 3.0σ. Their joint Mahalanobis distance, from the bootstrap covariance of the two exponents together with the model bands, is 6.2 ($p = 5 \times 10^{-9}$ for two degrees of freedom). On the M/X census the same estimator gives $b_B = 0.27 \pm 0.04$ and $b_L = -0.01 \pm 0.04$ ($N = 296$, 97 active regions). The length exponent sits 5.4σ below the target and 8.4σ below the structural 2/7, and at fixed ribbon field the partial correlation of length with temperature is zero. Figure 4 shows the two exponent estimates against the target, and the partial residual of temperature against length. Neither panel shows a length trend. Every alternative estimator returns the same pair of exponents (Appendix A.4). Peak temperature does not depend on loop length at fixed ribbon field.

Our length exponent stays consistent with zero across every alternative length definition, threshold, measurement-error assumption, weighting scheme, and peak-temperature definition we tested (Appendix A, Table 4). Our field–length covariance is $\gamma = 0.30 \pm 0.45$, which enters the marginal target on the full catalog (Section 5.2).

5.4 Loop Length, Flare Timescales, and Peak Temperature

The absence of a temperature–length dependence raises a separate question: does the length from ribbon separation carry any other physical information about the flare? It does, and the same flares show it. Figure 5 and Table 1 regress each GOES timescale of the M/X census on $\log L$ with the same estimator, controlling for ribbon field and flare class through the peak flux. Rise time, full-width-at-half-maximum duration, e-folding decay time, time to the temperature maximum, and total duration all scale with loop length, with exponents between 0.7 and 1.4 and partial correlations between 0.3 and 0.5. On the identical rows the peak-temperature exponent is -0.03 ± 0.03. Adding reconnection flux as a further control leaves the timescale correlations positive and the temperature correlation at zero.

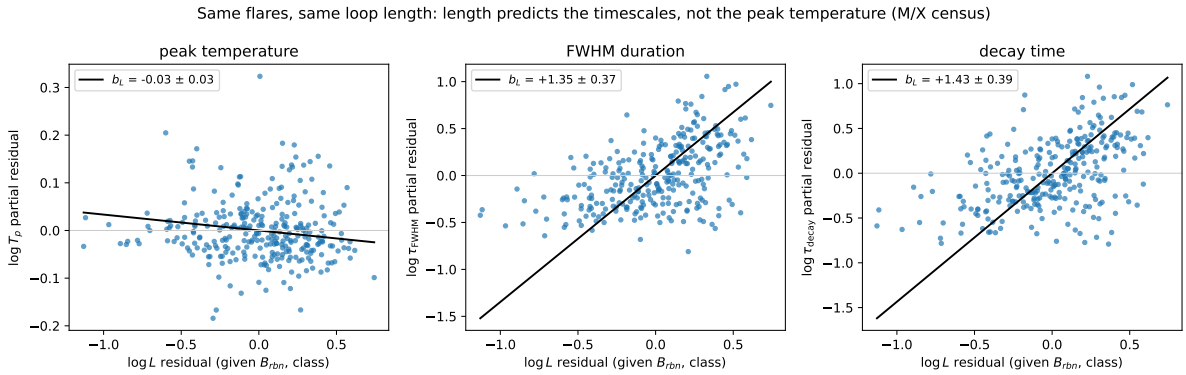


Figure 5: Same flares, same loop length. Partial residuals, after removing the dependence on $\bar{B}_{\rm rbn}$ and flare class, of **(a)** peak temperature, **(b)** GOES full-width-at-half-maximum duration, and **(c)** GOES decay time against loop length, for the M/X census. The fitted exponents (attenuation-corrected) are shown in each panel. Loop length predicts the timescales and not the peak temperature.

Table 1: Loop length against flare timescales on the M/X census. Each row gives the exponent corrected for attenuation of the response on $\log L$, alone and with controls for $\log \bar{B}_{\rm rbn}$ and log peak flux (class), with active-region block-bootstrap errors, and the partial correlation of the response with $\log L$ given the controls. The peak-temperature row is the joint test of Section 5.3 on the identical flares.

Response	N	b_L (L only)	b_L ($\bar{B}_{\rm rbn}$, class controlled)	partial r
Peak temperature $T_{\rm p}$	296	$+0.04 \pm 0.04$	-0.03 ± 0.03	-0.05
GOES rise time	296	$+0.87 \pm 0.26$	$+0.95 \pm 0.31$	$+0.37$
GOES FWHM duration	295	$+1.22 \pm 0.31$	$+1.35 \pm 0.37$	$+0.50$
GOES decay time	294	$+1.31 \pm 0.33$	$+1.43 \pm 0.39$	$+0.49$
Time from start to $T_{\rm p}$	294	$+0.65 \pm 0.22$	$+0.70 \pm 0.24$	$+0.30$
GOES total duration	296	$+0.87 \pm 0.26$	$+0.94 \pm 0.29$	$+0.40$

Toriumi et al. [12] established the duration–length relation for flares of M5 and above, and a conductively cooling loop predicts it. What is new is the contrast within the same sample: loop length predicts how long a flare lasts but not how hot it becomes. These timescale correlations also show that the same length measurement does carry physical information, so the null temperature result cannot be blamed on it.

6 Nonlinear Loop Response and the Effective Heating Law

6.1 Exact Conduction–Evaporation Reduction

The observations tell us that the standard law fails, but not why. There are two broad possibilities: either a heated loop responds differently than the theory assumes, or the magnetic field supplies heat differently than the theory assumes. We test the first here, deriving what a loop does with a given heating law before asking what the heating law itself must be. During the heating phase the pressure is driven by the heat input while the density is driven by the conductive flux that evaporates chromospheric material,

$$\frac{1}{\gamma-1}\frac{\mathrm{d}p}{\mathrm{d}t} \simeq Q(t), \qquad \frac{\mathrm{d}n}{\mathrm{d}t} \simeq -\,C\,\frac{F_{\rm c}}{TL}, \qquad F_{\rm c} \simeq -\frac{2}{7}\kappa_0\frac{T^{7/2}}{L}, \tag{11}$$

which are the enthalpy-balance equations that the reduced loop model integrates. Writing $p = 2nk_{\mathrm{B}}T$ turns the density equation into $\mathrm{d}n/\mathrm{d}t = Kp^{5/2}n^{-5/2}L^{-2}$, and multiplying by $n^{5/2}$ makes the pair integrable in quadrature,

$$n^{7/2}(t) = n_0^{7/2} + \frac{K'}{L^2}\int_0^t p^{5/2}\,\mathrm{d}t', \qquad p(t) = p_0 + (\gamma-1)\int_0^t Q\,\mathrm{d}t'. \tag{12}$$

For a pulse of amplitude H and shape fixed in $x = t/\tau_H$, the strong-heating limit gives $p \propto H\tau_H$, so the integral carries $H^{5/2}\tau_H^{7/2}$, whence, with $T_{\mathrm{p}} = p/(2nk_{\mathrm{B}})$,

$$\boxed{\,n \propto H^{5/7}\tau_H\,L^{-4/7}, \qquad T_{\mathrm{p}} \propto H^{2/7}\,\tau_H^{0}\,L^{4/7}\,}. \tag{13}$$

The duration cancels exactly, because $(5/2+1)\times(2/7) = 1$. This is a physical statement and not a dimensional one: a longer pulse does inject more energy, and in the evaporative regime that energy goes into loading more plasma into the loop, leaving its temperature alone.

This prediction carries no free parameter, and the simulations meet it. Sweeping single loops across two decades of heating rate, a factor of sixteen in pulse duration and a factor of eight in length, with one temperature enforced as Equation (11) assumes, the measured response is

$$T_{\mathrm{p}} \propto H^{0.285}\,\tau_H^{-0.005}\,L^{0.576}, \tag{14}$$

against the predicted $(2/7,\ 0,\ 4/7) = (0.286,\ 0,\ 0.571)$: every exponent within 0.005, with a residual of 0.0024 dex. Agreement is asymptotic in the one term the reduction drops, radiative loss: restricting to the runs where it is smallest tightens every exponent toward the prediction and halves the residual. Across that duration range the density rises by a factor of fourteen while the temperature moves by 1.4 per cent, so the cancellation is carrying the entire difference. Temperature is reproduced an order of magnitude more tightly than the pressure or the density separately, as Equation (13) requires of their ratio.

6.2 The Two-Fluid Correction and Its Parameter-Free Tests

The instrument measures an electron temperature, and in a real loop the electrons and ions need not share one. We extend the reduction of Section 6.1 to two temperatures and obtain the correction in closed form. Where electrons and ions separate, the exact relation $T_{\mathrm{p}} = p/[nk_{\mathrm{B}}(1+R)]$ with $R = T_i/T_e$ replaces the last step of Equation (13). Conduction now runs on both species, so the flux that drives the evaporation is

$$F_{\mathrm{c}} = -\frac{2}{7L}\left(\kappa_e T_e^{7/2} + \kappa_i T_i^{7/2}\right) = -\frac{2\kappa_e}{7L}\,T_e^{7/2}\left(1 + \kappa_\star R^{7/2}\right), \qquad \kappa_\star \equiv \frac{\kappa_i}{\kappa_e}, \tag{15}$$

which is where the $R^{7/2}$ of the correction comes from: the ion channel is small in $\kappa_\star$ but carries the seventh power of the temperature ratio. Mass conservation contributes a further prefactor $1/(1+R)$ and a $1/T_e$, and substituting $T_e = p/[nk_{\mathrm{B}}(1+R)]$ into Equation (15) gives

$$\frac{\mathrm{d}}{\mathrm{d}t}\,n^{7/2} = K\,\Psi(R)\,\frac{p^{5/2}}{L^2}, \qquad \Psi(R) = \frac{1+\kappa_\star R^{7/2}}{(1+R)^{7/2}}, \qquad \kappa_\star = \frac{\kappa_i}{\kappa_e} = 0.0410, \tag{16}$$

with $R = 1$ the single-fluid case. Carrying this through $T_{\mathrm{p}} = p/[nk_{\mathrm{B}}(1+R)]$, the factors of $(1+R)$ cancel identically and the whole correction collapses to

$$\log\frac{T_{\mathrm{p}}^{2\mathrm{F}}}{T_{\mathrm{p}}^{1\mathrm{F}}} = -\frac{2}{7}\,\log\frac{1+\kappa_\star R^{7/2}}{1+\kappa_\star}. \tag{17}$$

Something useful cancels here. Separating the ion and electron temperatures lowers the electron temperature, but it also suppresses the evaporation, and at the peak those two effects offset

exactly. If the ions did not conduct, the single-fluid law would hold for the electron temperature at *any* R. What remains is the ion conductive channel alone. Ion conductivity is intrinsically weak, but conduction scales as the seventh power of the square root of temperature, so hot enough ions still matter: at the measured split the amplification reaches 17, and the channel carries 41 per cent of the driver.

Nothing in Equation (17) is fitted; $\kappa_\star$ is a ratio of Spitzer coefficients. On the two-fluid grid the predicted and measured exponents agree to within 0.005, and the correction cuts the run-to-run residual sevenfold. The relation also predicts what happens when the heating partition is changed, a test the derivation never uses. Depositing all of the energy in the electrons drives the split the other way, where Equation (17) changes sign and predicts almost no shift, because cold ions conduct nothing; the measured shifts match to within 0.002. Depositing it all in the ions instead breaks the approximation, and breaks it informatively: the electron peak is then delayed and the similarity variables are sampled at a different point of the same trajectory. Restoring that timing term closes the decomposition at all three partitions. Appendix C gives the grids and the numerical shifts.

6.3 The Conductive Relaxation Spectrum

The duration cancellation has a complementary dynamical explanation. Linearizing the conductive problem about its equilibrium gives a weighted Sturm–Liouville system whose first two eigenvalues are $\lambda_1 = 3.207$ and $\lambda_2 = 22.300$, so structure beyond the fundamental mode decays about seven times faster. A loop consequently loses memory of the detailed heating history within a fraction of its own cooling time, which is why the pulse duration drops out of Equation (13). The derivation, the Rayleigh bound, the orthogonality relation, and the numerical verification are given in Appendix B.

6.4 Inverse Determination of the Heating Law

The previous sections fix how a loop responds to a given heating rate. We can now reverse the problem and ask what dependence of heating on B and L would produce the exponents the Sun actually shows.

Both measured exponents differ from the reconnection-heating prediction by several standard errors, and the constraint on the length dependence is the stronger: over the observed 1.9 dex in length the law predicts a factor of 2.3 in temperature, and the data show no measurable change.

Under the classical conductive balance of Section 2.1, a heating rate $H \propto B^a L^c$ gives $T_\mathrm{p} \propto B^{2a/7} L^{2(c+2)/7}$, so the measured census exponents require $a \simeq 0.96 \pm 0.13$ and $c \simeq -2.04 \pm 0.12$. The pair is close to the heating law inferred for steady active-region loops, for which a volumetric rate proportional to B/L^2 reproduces the observed emission [25]: the exponents the flare data require are ordinary for coronal heating and unexpected only for an instantaneous reconnection balance. The inversion is conditional, since it assumes that thermal transport and density retain their classical forms, and Section 6.1 shows that they do: the coupled conduction and evaporation of a loop returns $T_\mathrm{p} \propto H^{2/7} L^{4/7}$ with the pulse duration cancelling, and the simulations meet that prediction to within 0.005 in every exponent. A heating rate of the coronal form thus cancels the length almost exactly, because the L^2 of the conductive balance meets the L^{-2} of the rate, and with the ideal exponents the cancellation is exact. Reading the census through them directly gives $a = \frac{7}{2} b_B = 0.95$ and $c = \frac{1}{2}(7 b_L - 4) = -2.04$.

The measured single loop response tells us how a small change in the heating law changes the observable temperature, so we can invert that relation. Because the instrument responds to the electrons, we carry the two-fluid response of Equation (17) through the inversion in place of

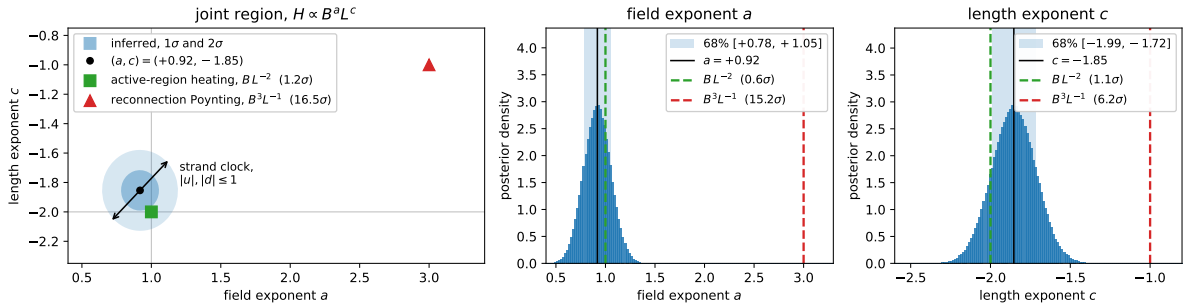


Figure 6: The heating law the Sun requires is close to B/L^2, and excludes the reconnection Poynting form assumed by the standard law. The volumetric rate is written $H \propto B^a L^c$ and inferred by inverting the measured single loop response of Section 6.1 through Equation (18), propagating the uncertainties of all five measured exponents. **(a)** The joint region, with the shaded areas the 1σ and 2σ contours. The arrow spans the displacement caused by an unknown strand clock $\tau_H \propto B^u L^d$ over $|u|, |d| \le 1$, which moves the inference by only $q/p = 0.21$ per unit. The heating law inferred for quiescent active-region loops lies 1.2σ from the centre; the reconnection Poynting rate that the Shibata–Yokoyama balance assumes lies 16.5σ away. **(b)** and **(c)** The two marginals, on which the coronal form sits 0.6σ and 1.1σ from the centre while the Poynting rate sits 15.2σ and 6.2σ away.

the classical exponents. Writing the unknown volumetric rate and strand clock as $H \propto B^a L^c$ and $\tau_H \propto B^u L^d$, the census exponents satisfy $b_B = pa + qu$ and $b_L = pc + qd + r$, and therefore

$$a = \frac{b_B - qu}{p} = 0.918 - 0.214\,u, \qquad c = \frac{b_L - r - qd}{p} = -1.854 - 0.214\,d. \tag{18}$$

Physically, this means the heating rate rises roughly linearly with field strength and falls almost as the inverse square of loop length. The clock enters with a leverage of $q/p = 0.214$, so even if the strand-clock exponents are uncertain by order unity the inferred heating-law exponents move only modestly. A clock independent of geometry gives $(a, c) = (+0.92, -1.85)$, one proportional to the loop gives $(+0.92, -2.07)$, and an Alfvénic one gives $(+1.13, -2.07)$. Propagating the uncertainties of all five measured exponents places the field exponent at 0.92 with a 68 per cent interval of $[0.78, 1.05]$ and the length exponent at -1.85 with $[-1.99, -1.72]$, the two being essentially uncorrelated. The coronal form $\bar{B}_{\rm rbn} L^{-2}$ lies inside that region, at 0.6 and 1.1 standard errors. The reconnection Poynting rate, which carries the cube of the field, lies 16 standard errors outside it (Figure 6).

Transport is the other alternative, and the simulations close it: driving single loops with the reconnection Poynting rate and then suppressing conduction with a flux limiter moves both exponents away from the census, the field exponent rising from +0.89 to +1.55 and the length exponent from +0.22 to +0.45, while the peak temperature climbs past 500 MK. No conduction regime we can construct turns the canonical heating law into the observed response. The observations therefore disagree with the assumed law converting magnetic energy to heat, not with the conductive physics that law is fed into.

The inferred $\bar{B}_{\rm rbn} L^{-2}$ relation should therefore be interpreted as an effective thermal heating law; its underlying microscopic mechanism remains unidentified.

6.5 Implications for Transport and Stellar Inference

The observations reveal a distinction that the standard law does not capture: loop length predicts how long a flare lasts but not how hot it becomes. Any revised heating model has to reproduce both halves, timescales that grow with loop length and a peak temperature that does not. The consequence for stellar work is direct. The Shibata–Yokoyama relations are routinely inverted to infer coronal fields and loop lengths of unresolved stellar flares from their X-ray temperatures and emission measures [7, 9]. Because an inferred field scales as a high inverse power of temperature, a forward slope of 0.35–0.6 in place of 6/7 compresses a wide range of true fields into a narrow inferred one, so such fields are best interpreted at order-of-magnitude precision. With no measurable temperature dependence on length, a loop length inferred this way has no empirical support in the solar sample at all.

6.6 The Reconnection-Rate Test

One alternative to changing the heating law remains: the assumed law might fail simply because the reconnection rate itself has the wrong dependence on magnetic field. Ribbon motion tests that directly, and independently of the temperature law. Writing ℓ for the mean extent of the two ribbons along their principal axes and $\tau_{\rm rise}$ for the GOES rise time, the mean reconnection electric field is

$$E_{\rm rec} = \frac{\Phi_{\rm rbn}}{c\,\tau_{\rm rise}\,\ell}, \tag{19}$$

a quantity set by a Petschek-type inflow [26] and related by flux conservation to the field of the newly connected loops [27]. Its median is 10.6 V cm^{-1} on the M/X census, within the range reported from Hα and ultraviolet kernel motions [27–29]. It scales steeply with the ribbon field,

$$E_{\rm rec} \propto \bar{B}_{\rm rbn}^{1.34\pm0.24} L^{-0.34\pm0.08}, \tag{20}$$

So the reconnection rate itself becomes substantially stronger where the field is stronger, close to what an Alfvénic inflow requires. The peak temperature, by contrast, does not follow it: regressing $\log T_{\rm p}$ on $\log E_{\rm rec}$ at fixed field and length gives $+0.01 \pm 0.02$, against the 2/7 that conductive balance requires. The reconnection rate therefore retains approximately the magnetic dependence expected for an Alfvénic inflow, while the peak temperature does not track it. The reconnection-rate scaling is therefore not the source of the discrepancy; the mismatch appears later, in how that magnetic energy becomes heat.

7 Reconnection Sequence and Flare Duration

Loop length predicts every timescale of these flares and no temperature (Figure 5). Section 6.1 settles half of that separation: the pulse duration cancels out of the peak temperature exactly, and more heating in this regime loads more plasma into the loop and leaves it no hotter. The arcade, furthermore, sets the other half: the clock.

A flare arcade is a sequence of loops formed and heated in turn, so its soft X-ray light curve is the response of one strand convolved with the sequence in which strands are lit [30, 31],

$$F(t) = \int R(t')\, g_{L(t')}(t - t')\, {\rm d}t'. \tag{21}$$

A strand emits for tens of seconds and an arcade forms over hundreds, so R sets the width of the observed light curve. Ultraviolet footpoint brightenings make that sequence measurable: summing the magnetic flux under newly brightened 1600 Å pixels gives the cumulative reconnected flux of each flare [32, 33], whose 10-to-90 per cent interval is a reconnection duration of median

661 s, against an observed full width of about 720 s. That duration grows with loop length, as $\tau_{\rm rec} \propto L^{0.67\pm0.25}$, and driving an arcade with the measured $\tau_{\rm rec}(L)$, with nothing fitted to the GOES light curve, returns a width of 670 s at $L = 20$ Mm. Longer arcades simply take longer to assemble.

8 Limitations and Next Tests

Three limits bound the measurement. The loop length is a ribbon footpoint separation converted to a half-length under a semicircular geometry, so sheared or non-circular loops change its normalization while leaving its scaling intact; the external validation bounds the per-event error at the 0.25 dex inter-method scatter, within which the null length dependence is stable (Table 4). The inclusion weights make the all-class estimate imprecise, so the population statement rests on the M/X census. The thermal pipeline searches the whole light-curve window for the temperature maximum, and restricting that search to the catalog flare interval moves the slope on the full catalog from 0.35 ± 0.03 to 0.42 ± 0.03, with the joint exponents unchanged within their errors; we retain the whole-window procedure as primary.

Three conditions attach to the inferred heating law. The reduction of Section 6.1 is verified against a spatially averaged two-fluid loop model, which stops short of resolved radiative hydrodynamics [22, 23]. The closed-form two-fluid correction assumes that the one- and two-fluid calculations peak at the same time, which holds for equal and electron-dominated heating and fails for pure ion heating; restoring the timing term closes the decomposition in every case, so only the closed form is restricted (Appendix C). The inversion into a heating law is conditional on the strand clock, which this census does not measure, though it enters with a leverage of only 0.21. What remains open is the microscopic magnetic origin of the law the census infers.

Three experiments would narrow what remains. Measuring the hot component above 20 MK on this census, from hard X-ray spectroscopy, would test whether any part of the flare retains the reconnection-heating scaling; that component is distinct from the 10 to 20 MK plasma GOES reports and its temperature behaves differently [34]. Assembling each arcade from its own reconnection history, in place of a median $\tau_{\rm rec}(L)$, would show how much of the duration exponent comes from later loops being longer. Repeating the analysis on Solar Cycle 25 flares would test every number here out of sample.

9 Conclusion

1. **The direct measurements reject the canonical exponents.** On observables taken from the reconnection geometry itself, neither exponent of the Shibata–Yokoyama temperature law is recovered: $(b_B, b_L) = (0.27 \pm 0.04, -0.01 \pm 0.04)$ on the complete M- and X-class sample against a predicted $(0.89, 0.19)$, and $b_B = 0.35 \pm 0.03$ over 3056 flares. The forward model preserves the field exponent through the GOES inversion, so the shortfall is not instrumental. The same loop lengths, on the same flares, predict the rise, duration and decay times with exponents near unity.
2. **Classical loop thermodynamics explain the response.** Coupling the pressure driven by the heat input to the density driven by the conductive flux gives a system integrable in quadrature, in which the pulse duration cancels and the peak temperature follows $T_{\rm p} \propto H^{2/7} L^{4/7}$ with no free parameter. A sweep across two decades of heating rate, a factor of sixteen in duration and a factor of eight in length reproduces every exponent to within 0.005. The two-temperature correction follows from the same reduction with nothing fitted, and is confirmed out of sample by the sign change it requires when the heating is given to the electrons alone.

3. **The failure lies in the magnetic-to-thermal heating law.** Inverting the measured response, the effective volumetric rate the Sun requires is $\bar{B}_{\rm rbn}^{0.92} L^{-1.85}$, and the coronal form $\bar{B}_{\rm rbn} L^{-2}$ lies inside that region, while the reconnection Poynting rate lies 16 standard errors outside it. The independently measured reconnection rate retains approximately the magnetic dependence an Alfvénic inflow requires, so the classical thermal physics holds and the law feeding it does not. The microscopic mechanism behind that law remains open.

Loop geometry sets the flare clock, but not its thermometer.

A Supporting Methods and Robustness

A.1 Thermal Inversion and Forward-Model Details

Reconnection releases magnetic energy at a rate per unit volume set by the magnetic energy density carried into the current sheet at the Alfvén speed,

$$H \simeq \frac{B^2}{4\pi} \frac{v_A}{L}, \qquad v_A = \frac{B}{\sqrt{4\pi\rho}} \propto B\, n_0^{-1/2}, \tag{22}$$

so that $H \propto B^3 n_0^{-1/2} L^{-1}$. Heated plasma loses energy by field-aligned thermal conduction, whose rate per unit volume for a loop with apex temperature T is

$$C \simeq \kappa_0 \frac{T^{7/2}}{L^2}, \tag{23}$$

with $\kappa_0 \simeq 10^{-6}$ in cgs units (Spitzer conductivity). Setting $H = C$ gives Equation (6).

The normalization of Equation (7) is the pressure-balance form of Shibata and Yokoyama [7]: the coefficient 1.8×10^4 is the evaporation-closed version, which sets the evaporated density by pressure balance. The larger coefficient 5.3×10^4 quoted by Shibata and Yokoyama [6] is the bare reconnection–conduction balance before the chromospheric-evaporation closure, with the identical exponents.

All pre-GOES-16 fluxes are recalibrated to the true scale before inversion [35]. We report the median peak temperature across the admissible background windows and take the spread across the admissibility criteria as the uncertainty, extracting $T_{\rm p}$ and $EM_{\rm p}$ at their own maxima.

The forward model is an arcade of twelve strands with onsets staggered over three pulse durations and log-normal jitter in heating amplitude, evaluated on the grid of Table 2. The GOES synthesis uses coronal abundances and the 0.5–4 and 1–8 Å channels. The heating efficiency was fixed on the pilot at $\epsilon \simeq 5 \times 10^{-3}$ so that peak temperatures span the observed C–X range; because it multiplies the heating rate uniformly, it sets the normalization of the temperature law and leaves its exponents alone. Thread count, onset stagger, amplitude jitter, abundance set, heating duration, efficiency, and density grid were all fixed before the thermal response was examined.

Passing a heating law of the coronal form through this arcade and the instrument costs part of the accuracy with which the length cancels. Folded through the GOES response and inverted as the data are, the model returns a length exponent of -0.15 at the temperature maximum and -0.26 at the flux peak. The first is recovered exactly by adding the transfer the forward model already carries, -0.10, to the single-loop value, so the arcade and the instrument account for the difference. Letting the strand heating time grow with the loop restores the agreement at the flux peak, at $(+0.26, -0.00)$, though the sweep of Section 6.1 shows the duration itself has little leverage on a single loop.

Table 2: Forward-model configuration. All entries were fixed before the thermal response was examined.

Quantity	Value
Strands per arcade	12, onsets staggered over $3\tau_H$
Heating pulse	triangular, $\tau_H = 100$ s
Amplitude jitter	log-normal, 0.15 dex
Ribbon field $\bar{B}_{\rm rbn}$	$[200, 1300]$ G, 8 nodes
Loop half-length L	$[5 \times 10^8, 10^{10}]$ cm, 8 nodes
Pre-flare density n_0	$[3 \times 10^8, 3 \times 10^9]$ cm^{-3}, 5 nodes
Realizations per node	4 independent stochastic draws
GOES synthesis	coronal abundances, 0.5–4 and 1–8 Å
Heating efficiency	$\epsilon \simeq 5 \times 10^{-3}$ (normalization only)

A.2 Numerical Verification of the Conductive Theory

The eigenvalues quoted in Section 6.3 come from a cell-centred finite-volume discretization of Equation (25) on $[0, 1]$, with zero flux at the loop apex and the Dirichlet condition applied half a cell beyond the last node, and with the weight integrated over each cell so that the endpoint singularity is averaged over each cell. With the symmetric operator the constant-weight control reproduces the closed-form eigenvalues $(\pi/2)^2$, $(3\pi/2)^2$ and $(5\pi/2)^2$ to five decimal places. The weighted problem then converges to $\lambda_1 = 3.2074$ and $\lambda_2 = 22.2999$, stable in the fourth decimal between 3200 and 12800 cells, and the computed fundamental respects the Rayleigh bound of Equation (26).

A.3 Measurement Construction and Validation

For each event we use AIA 1600 Å cutouts ($400''$, native 24 s cadence) and one HMI 720 s line-of-sight magnetogram. Each frame is normalized by exposure, and frames containing saturated pixels are discarded. The ribbon composite is the set of pixels exceeding k times the quiet-Sun level in any retained frame, with isolated pixels removed. It is split by the sign of the co-registered HMI field, and the two largest connected components in each polarity are kept. An event enters the joint test if it has at least three valid thresholds and at least five pixels in each polarity.

The comparison with Toriumi et al. [12] improves from $r = 0.76$ to 0.81 when our composite is extended to 35 min after the peak. The estimator also passes a physical check on our own sample: the ribbon distance correlates with the GOES full-width-at-half-maximum duration at $r = 0.69$ for the flares of M5 and above, where Toriumi et al. report 0.83, while our durations reproduce theirs at $r = 0.98$. As cross-checks we compute cooling-time and footpoint-to-polarity-inversion-line lengths, and evaluate the ribbon composite at the GOES flux peak and at the end of the extended window (Table 4).

The C-class stratification draws 60 events per class decade, evenly in time within cells defined by class decade, GOES rise-time tercile (edges from the M/X reference distribution), and eruptivity. The resulting inclusion weights range up to 59 for the sparsest cell.

A.4 Statistical and Variable-Identification Robustness

One question underlies the three checks collected here: could the construction of the variables, or error in them, have flattened the measured slopes?

Because $\bar{B}_{\rm rbn}$ is constructed from $\Phi_{\rm rbn}$ and $S_{\rm rbn}$, the unrestricted regression

$$\log T_{\rm p} = a + b_\Phi \log \Phi_{\rm rbn} + b_S \log S_{\rm rbn} \tag{24}$$

is fit with both coefficients free. Equation (7) predicts that they concentrate on the ratio direction $(b_\Phi, b_S) \approx b_B(1, -1)$. Coefficients lying elsewhere would mean the data fail to support the variable identification of the law. The measured pair is $(+0.19, -0.14)$, a regression angle of -36° against -45° for the exact ratio. A Monte Carlo attributes about a third of that difference to measurement attenuation and the remainder to a small real departure from the exact ratio, and neither affects the field slope.

The reported exponents come from the errors-in-variables estimator of Section 4.2, with ordinary least squares and Theil–Sen quoted alongside as uncorrected references.

The attenuation bound of Section 5.2 follows from the field error entering in two places at once. The dispersion $\sigma_{\log \bar{B}_{\rm rbn}}$ that sets the attenuation correction also sets how far the unrestricted flux–area regression rotates away from the exact intensive direction. Attributing the entire -36° offset of that angle (Appendix A.4) to attenuation requires a field error 1.35 times the nominal dispersion, which raises the corrected marginal slope from 0.35 to 0.6. Reaching 6/7 would need 1.5 times the nominal dispersion together with a ratio angle of -29°, which the measured angle excludes.

On the joint field–length fit the alternative estimators agree with the corrected one. The unweighted attenuation-corrected fit gives $(b_B, b_L) = (0.36 \pm 0.05, -0.04 \pm 0.03)$, weighted ordinary least squares gives $(0.21, -0.02)$, and the Theil–Sen slope of the residual of the field fit on $\log L$ is -0.02.

The κ-residual check of Section 3.1 asks whether the photospheric $\bar{B}_{\rm rbn}$ tracks the coronal field closely enough for the measurement to bear on the law. A constant κ shifts the target normalization by ≈ 0.3 dex and the target slope by less than 0.01, so flattening the measured slope requires a conversion that varies with field strength or environment. The residual analysis detects such a contribution and finds it too small to restore the predicted exponent. The residual correlates only weakly with projection ($|r| \approx 0.16$) and most strongly with the mean active-region field $B_{\rm AR}$ ($r = +0.28$, and $+0.24$ at fixed flare size), which carries temperature information comparable to that in the ribbon field and independent of flare size. The variable-identification residual of Appendix A.4 gives a consistent indication. The coronal field relevant to peak temperature may therefore act over a larger scale than the instantaneous ribbon footprint, as expected if the reconnecting current sheet is supplied by the broader active-region field [36]. None of the photospheric field measures considered here reaches the predicted exponent, so the data do not distinguish completely between an imperfect photosphere-to-corona mapping and a breakdown of the analytic heating balance.

A.5 Robustness Tables

Table 3 collects the full-catalog field-slope checks and Table 4 the joint field–length checks. The shallow field dependence and the null length dependence hold across every subsample, threshold, estimator, error budget, weighting scheme and temperature definition in them.

Transport is the one alternative to changing the heating law that can be tested directly here. Driving single loops with the reconnection Poynting rate and suppressing conduction with a flux limiter moves the predicted response further from the measurement, not closer, across the whole range of limiter strengths we can construct.

Table 3: Full-catalog field-slope robustness. Each row gives the uncorrected Theil–Sen slope b_B of $\log T_{\rm p}$ on $\log \bar{B}_{\rm rbn}$, and the Pearson correlation r, for a subsample or cut. The slopes are uncorrected and scale to their counterparts corrected for attenuation by subsample-specific factors (near 1.5 for the full sample, larger where the $\bar{B}_{\rm rbn}$ range is restricted); none approaches the Shibata–Yokoyama exponent $6/7 \simeq 0.857$, corrected or not.

Subsample or cut	N	b_B	r
Full sample	3056	0.24	0.32
C only	2758	0.20	0.28
M and X	298	0.23	0.36
2011–2013	1598	0.25	0.35
2014–2016	1438	0.22	0.29
Confined	324	0.23	0.33
Eruptive	151	0.29	0.40
$\sigma_{\log T_{\rm p}} \leq 0.15$ dex	2650	0.21	0.34
Flux above 10th percentile	2708	0.24	0.32
M and X, $\sigma_{\log T_{\rm p}} < 0.1$	192	0.18	0.39
Above C5	699	0.25	0.35
Above M1	298	0.23	0.36
Above M5	51	0.18	0.38
Within active region	3056	0.14	–

Table 4: Joint field–length exponents. Rows give the attenuation-corrected joint slopes (b_B, b_L) of $\log T_{\rm p}$ on $(\log \bar{B}_{\rm rbn}, \log L)$ with standard errors from a block bootstrap over active regions; the primary rows use the inclusion weights, the census row weights every M- and X-class event equally. The target is $(0.89, 0.19)$ and the structural exponents are $(6/7, 2/7)$.

Fit	N	b_B	b_L
Primary (weighted, ribbon-centroid L)	416	0.32 ± 0.13	-0.05 ± 0.08
M/X census (unweighted)	296	0.27 ± 0.04	-0.01 ± 0.04
Weights capped at 10	416	0.43 ± 0.08	-0.06 ± 0.05
L from composite at GOES flux peak	414	0.31 ± 0.12	-0.04 ± 0.04
L from cumulative mask	415	0.30 ± 0.12	-0.05 ± 0.04
L from footpoint-to-PIL distance	405	0.27 ± 0.13	-0.04 ± 0.10
L from GOES cooling time	413	0.34 ± 0.15	-0.07 ± 0.08
Threshold family, $k = 4$ to 12 (range)	416	0.28–0.33	-0.06 to -0.03
$\sigma_{\log L} = 0.25$ dex for all events	416	0.32 ± 0.12	-0.05 ± 0.06
M and X, $\sigma_{\log T_{\rm p}} < 0.1$	190	0.24 ± 0.04	-0.04 ± 0.04
Flare-interval $T_{\rm p}$, weighted	415	0.53 ± 0.13	0.04 ± 0.08
Flare-interval $T_{\rm p}$, M/X census	295	0.22 ± 0.04	-0.04 ± 0.05

B The Conductive Relaxation Spectrum

The duration cancels because a loop relaxes toward its conductive equilibrium quickly, and we can show that analytically. Substituting $u = \theta^{7/2}$ in the fixed-density problem makes the stationary equation linear, $1 + \frac{2}{7} u_{xx} = 0$, with the exact profile $u_0 = \frac{7}{4}(1 - x^2)$. First-order perturbations obey $\eta_\xi = u_0^{5/7} \eta_{xx}$, and separating variables gives a weighted Sturm–Liouville problem with an integrable endpoint singularity,

$$-\phi''(x) = \lambda \left[\tfrac{7}{4}(1 - x^2)\right]^{-5/7} \phi(x), \qquad \phi'(0) = 0, \quad \phi(1) = 0, \tag{25}$$

whose fundamental eigenvalue has the variational characterization

$$\lambda_1 = \min_{\phi(1)=0} \frac{\displaystyle\int_0^1 (\phi')^2 \,\mathrm{d}x}{\displaystyle\int_0^1 \left[\tfrac{7}{4}(1-x^2)\right]^{-5/7} \phi^2 \,\mathrm{d}x}, \tag{26}$$

which the trial function $\cos(\pi x/2)$ bounds at $\lambda_1 \le 3.223$. The problem is self-adjoint in the weight $w(x) = [\tfrac{7}{4}(1-x^2)]^{-5/7}$, so its eigenfunctions are orthogonal,

$$\int_0^1 w(x)\, \phi_m(x)\, \phi_n(x) \,\mathrm{d}x = 0 \qquad (m \neq n), \tag{27}$$

and any initial departure from equilibrium relaxes as

$$\eta(x,\xi) = \sum_{m=1}^{\infty} a_m\, \phi_m(x)\, e^{-\lambda_m \xi}, \qquad a_m = \frac{\int_0^1 w\, \eta(x,0)\, \phi_m \,\mathrm{d}x}{\int_0^1 w\, \phi_m^2 \,\mathrm{d}x}. \tag{28}$$

A cell-centred finite-volume solution, verified against a constant-weight problem with known eigenvalues, gives $\lambda_1 = 3.207$ and $\lambda_2 = 22.300$. Because $\lambda_2/\lambda_1 = 6.95$, the second mode decays 6.95 times faster than the first in Equation (28): after one fundamental relaxation time $\xi = 1/\lambda_1$, the first overtone has fallen by $e^{-5.95}$ relative to the fundamental, a factor of 385. The residual is fundamental-mode dominated almost immediately, and that is why a finite pulse leaves so little trace of its duration.

C Two-Fluid Partition Tests

A heating partition (which the derivation never uses) supplies a test out of sample. Depositing all of the energy in the electrons drives R to 0.37, a split of similar size in the opposite sense, where Equation (17) changes sign and predicts almost no shift at all because cold ions conduct nothing: the measured shift is $(-0.001, -0.011, +0.003)$ against a predicted $(-0.001, -0.013, +0.005)$, and the electron exponents stay within 0.016 of the single-fluid law despite a factor of 2.7 between the species. Depositing all of it in the ions instead breaks the approximation, and the way it breaks is informative. There the electron peak is delayed by 74 per cent, and because $p(t)$ is nearly common to every run the similarity variables are then sampled at a different point of the same trajectory. That timing term, at most 0.002 in the other two cases, grows to 0.18 and dominates. Restoring it closes the decomposition to 0.003 in every exponent at all three partitions. This reduction therefore holds throughout. Only Equation (17) (which assumes a common peak time) is restricted. Since the instrument responds to the electrons, Section 6.4 carries the two-fluid response through to the solar inference.

References

[1] M. D. Kazachenko, B. J. Lynch, B. T. Welsch, and X. Sun. A Database of Flare Ribbon Properties from the Solar Dynamics Observatory. I. Reconnection Flux. *Astrophys. J.*, 845: 49, 2017. doi: 10.3847/1538-4357/aa7ed6.

[2] M. D. Kazachenko. A Database of Magnetic and Thermodynamic Properties of Confined and Eruptive Solar Flares. *Astrophys. J.*, 958:104, 2023. doi: 10.3847/1538-4357/ad004e.

[3] J. Qiu, J. Lee, D. E. Gary, and H. Wang. Motion of Flare Footpoint Emission and Inferred Electric Field in Reconnecting Current Sheets. *Astrophys. J.*, 565:1335, 2002. doi: 10.1086/324706.

[4] J. Tschernitz, A. M. Veronig, J. K. Thalmann, J. Hinterreiter, and W. Poetzi. Reconnection Fluxes in Eruptive and Confined Flares and Implications for Superflares on the Sun. *Astrophys. J.*, 853:41, 2018. doi: 10.3847/1538-4357/aaa199.

[5] T. Yokoyama and K. Shibata. A Two-dimensional Magnetohydrodynamic Simulation of Chromospheric Evaporation in a Solar Flare Based on a Magnetic Reconnection Model. *Astrophys. J. Lett.*, 494:L113, 1998. doi: 10.1086/311174.

[6] K. Shibata and T. Yokoyama. Origin of the Universal Correlation between the Flare Temperature and the Emission Measure for Solar and Stellar Flares. *Astrophys. J. Lett.*, 526:L49, 1999. doi: 10.1086/312354.

[7] K. Shibata and T. Yokoyama. A Hertzsprung-Russell-like Diagram for Solar/Stellar Flares and Corona: Emission Measure versus Temperature Diagram. *Astrophys. J.*, 577:422, 2002. doi: 10.1086/342141.

[8] R. Rosner, W. H. Tucker, and G. S. Vaiana. Dynamics of the quiescent solar corona. *Astrophys. J.*, 220:643–665, 1978. doi: 10.1086/155949.

[9] K. Namekata, T. Sakaue, K. Watanabe, A. Asai, and K. Shibata. Validation of a Scaling Law for the Coronal Magnetic Field Strength and Loop Length of Solar and Stellar Flares. *Publ. Astron. Soc. Japan*, 69:7, 2017. doi: 10.1093/pasj/psw111.

[10] M. J. Aschwanden. Global Energetics of Solar Flares. XII. Energy Scaling Laws. *Astrophys. J.*, 903:23, 2020. doi: 10.3847/1538-4357/abb946.

[11] J. W. Reep and K. J. Knizhnik. What Determines the X-Ray Intensity and Duration of a Solar Flare? *Astrophys. J.*, 874:157, 2019. doi: 10.3847/1538-4357/ab0ae7.

[12] S. Toriumi, C. J. Schrijver, L. K. Harra, H. Hudson, and K. Nagashima. Magnetic Properties of Solar Active Regions That Govern Large Solar Flares and Eruptions. *Astrophys. J.*, 834:56, 2017. doi: 10.3847/1538-4357/834/1/56.

[13] S. Serio, F. Reale, J. Jakimiec, B. Sylwester, and J. Sylwester. Dynamics of Flaring Loops. I. Thermodynamic Decay Scaling Laws. *Astron. Astrophys.*, 241:197, 1991.

[14] F. Reale, R. Betta, G. Peres, S. Serio, and J. McTiernan. Determination of the Length of Solar Flare Loops from the X-Ray Light Curves. *Astron. Astrophys.*, 325:782, 1997.

[15] F. Reale. Diagnostics of stellar flares from X-ray observations: from the decay to the rise phase. *Astron. Astrophys.*, 471:271, 2007. doi: 10.1051/0004-6361:20077223.

[16] W.-J. Liu, J. Qiu, D. W. Longcope, and A. Caspi. Determining Heating Rates in Reconnection Formed Flare Loops of the M8.0 Flare on 2005 May 13. *Astrophys. J.*, 770:111, 2013. doi: 10.1088/0004-637X/770/2/111.

[17] S. M. White, R. J. Thomas, and R. A. Schwartz. Updated Expressions for Determining Temperatures and Emission Measures from GOES Soft X-Ray Measurements. *Solar Phys.*, 227:231, 2005. doi: 10.1007/s11207-005-2445-z.

[18] D. F. Ryan, R. O. Milligan, P. T. Gallagher, B. R. Dennis, A. K. Tolbert, R. A. Schwartz, and C. A. Young. The Thermal Properties of Solar Flares over Three Solar Cycles Using GOES X-Ray Observations. *Astrophys. J. Suppl.*, 202:11, 2012. doi: 10.1088/0067-0049/202/2/11.

[19] J. A. Klimchuk, S. Patsourakos, and P. J. Cargill. Highly Efficient Modeling of Dynamic Coronal Loops. *Astrophys. J.*, 682:1351, 2008. doi: 10.1086/589426.

[20] P. J. Cargill, S. J. Bradshaw, and J. A. Klimchuk. Enthalpy-based Thermal Evolution of Loops. II. Improvements to the Model. *Astrophys. J.*, 752:161, 2012. doi: 10.1088/0004-637X/752/2/161.

[21] W. T. Barnes, P. J. Cargill, and S. J. Bradshaw. Inference of Heating Properties from Hot Non-flaring Plasmas in Active Region Cores. I. Single Nanoflares. *Astrophys. J.*, 829:31, 2016. doi: 10.3847/0004-637X/829/1/31.

[22] S. J. Bradshaw and P. J. Cargill. The Influence of Numerical Resolution on Coronal Density in Hydrodynamic Models of Impulsive Heating. *Astrophys. J.*, 770:12, 2013. doi: 10.1088/0004-637X/770/1/12.

[23] F. Reale. Coronal Loops: Observations and Modeling of Confined Plasma. *Living Rev. Solar Phys.*, 11:4, 2014. doi: 10.12942/lrsp-2014-4.

[24] M. G. Akritas and M. A. Bershady. Linear Regression for Astronomical Data with Measurement Errors and Intrinsic Scatter. *Astrophys. J.*, 470:706, 1996. doi: 10.1086/177901.

[25] C. J. Schrijver, A. W. Sandman, M. J. Aschwanden, and M. L. DeRosa. The coronal heating mechanism as identified by full-sun visualizations. *Astrophys. J.*, 615:512–525, 2004. doi: 10.1086/422404.

[26] H. E. Petschek. Magnetic field annihilation. In *The Physics of Solar Flares*, volume 50 of *NASA Special Publication*, pages 425–439, 1964.

[27] J. Qiu, D. E. Gary, and G. D. Fleishman. Evaluating mean magnetic field in flare loops. *Solar Phys.*, 255:107–118, 2009. doi: 10.1007/s11207-009-9316-y.

[28] J. Qiu, H. Wang, C. Z. Cheng, and D. E. Gary. Magnetic reconnection and mass acceleration in flare-coronal mass ejection events. *Astrophys. J.*, 604:900–905, 2004. doi: 10.1086/382122.

[29] J. Hinterreiter, A. M. Veronig, J. K. Thalmann, J. Tschernitz, and W. Pötzi. Statistical Properties of Ribbon Evolution and Reconnection Electric Fields in Eruptive and Confined Flares. *Solar Phys.*, 293:38, 2018. doi: 10.1007/s11207-018-1253-1.

[30] J. W. Reep and S. Toriumi. The direct relation between the duration of magnetic reconnection and the evolution of goes light curves in solar flares. *Astrophys. J.*, 851:4, 2017. doi: 10.3847/1538-4357/aa96fe.

[31] J. W. Reep, H. P. Warren, C. S. Moore, C. Suarez, and L. A. Hayes. Simulating solar flare irradiance with multithreaded models of flare arcades. *Astrophys. J.*, 895:30, 2020. doi: 10.3847/1538-4357/ab89a0.

[32] J. Qiu, W.-J. Liu, and D. W. Longcope. Heating of flare loops with observationally constrained heating functions. *Astrophys. J.*, 752:124, 2012. doi: 10.1088/0004-637X/752/2/124.

[33] J. Qiu, Z. Sturrock, D. W. Longcope, J. A. Klimchuk, and W.-J. Liu. Ultraviolet and extreme-ultraviolet emissions at the flare footpoints observed by the atmospheric imaging assembly. *Astrophys. J.*, 774:14, 2013. doi: 10.1088/0004-637X/774/1/14.

[34] A. Caspi, S. Krucker, and R. P. Lin. Statistical properties of super-hot solar flares. *Astrophys. J.*, 781:43, 2014. doi: 10.1088/0004-637X/781/1/43.

[35] Janet Machol, Rodney Viereck, Courtney Peck, and James Mothersbaugh, III. GOES X-ray Sensor (XRS) Operational Data. NOAA/NCEI documentation, Version 1.5, 2022. `https://www.ngdc.noaa.gov/stp/satellite/goes/doc/GOES_XRS_readme.pdf` (accessed 2026 June 9).

[36] S. Toriumi and H. Wang. Flare-productive active regions. *Living Rev. Solar Phys.*, 16:3, 2019. doi: 10.1007/s41116-019-0019-7.