







The Origins of Hot Jupiter Inflation from Causal Discovery

ZEHAO JIN (金泽灏) ^{1,2} MOHAMAD ALI-DIB ¹ YUJIA ZHENG (郑雨嘉) ³ MARIO PASQUATO ^{4,5}
BENJAMIN L. DAVIS ¹ AND ANDREA V. MACCIÒ ¹

¹Center for Astrophysics and Space Science (CASS), New York University Abu Dhabi, PO Box 129188, Abu Dhabi, UAE

²Center for Astronomy and Astrophysics and Department of Physics, Fudan University, Shanghai 200438, People’s Republic of China

³Carnegie Mellon University, PA, USA

⁴INAF IASF-Milano, Via Alfonso Corti 12, 20133 Milano, Italy

⁵Ciela, Computation and Astrophysical Data Analysis Institute, Montreal, Quebec, Canada

(Received June 12, 2026)

Submitted to AAS Journals

ABSTRACT

Hot Jupiters often have radii larger than predicted by standard cooling–contraction models, but it remains unclear which process supplies or preserves the extra internal heat. We analyze 328 short-period giant planets with measured M_p , R_p , P_{orb} , and host-star T_{eff} using causal discovery, a statistical framework that asks which observed properties remain directly connected to planet radius after the others are accounted for. As a check, the same pipeline recovers the expected mass–radius connection for a super-Earth control sample. For hot Jupiters, the preferred graph links R_p directly to P_{orb} and T_{eff} , but not to M_p . Since incident flux increases with T_{eff} and decreases with P_{orb} at fixed stellar properties, this paired dependence is naturally interpreted as a population-level signature of irradiation-regulated inflation. Comparing the graph with analytic radius-excess scalings, we find that Gold–Soter thermal tides provide the closest single-mechanism match, while kinetic/mechanical heating and ohmic dissipation remain viable contributors. Purely period-controlled gravitational tides are disfavored as the sole explanation because they lack a leading dependence on stellar temperature. Distinguishing thermal tides, kinetic/mechanical heating, ohmic dissipation, and mixed scenarios will require radius-excess measurements that control for incident flux, age, composition, stellar properties, and selection effects. More broadly, this work shows how causal discovery can turn population-level exoplanet data into physically interpretable tests of hot-Jupiter inflation.

Keywords: [Astrostatistics \(1882\)](#) — [Exoplanet evolution \(491\)](#) — [Exoplanet formation \(492\)](#) — [Exoplanet structure \(495\)](#) — [Hot Jupiters \(753\)](#)

1. INTRODUCTION

Transiting hot Jupiters, gas giants of $M \sim 0.3$ to $2 M_{\text{J}}$ on $P \lesssim 10$ d orbits, show $R \gtrsim 1.2 R_{\text{J}}$ and reach $\sim 1.9 R_{\text{J}}$, exceeding cooling-contraction predictions for their ages and compositions (D. Charbonneau et al. 2000; A. Burrows et al. 2003; J. J. Fortney et al. 2007; G. Laughlin et al. 2011). Surveys by *HAT*, *WASP*, *Kepler*, and *TESS* show significant inflation mainly above $F_{\text{inc}} \gtrsim 2 \times 10^8 \text{ erg s}^{-1} \text{ cm}^{-2}$, or $T_{\text{eq}} \gtrsim 1000 \text{ K}$ for zero Bond albedo/full redistribution, with radius excess ris-

ing with F_{inc} and falling with planet mass (B.-O. Demory & S. Seager 2011; L. M. Weiss et al. 2013; M. Sestovic et al. 2018; D. P. Thorngren & J. J. Fortney 2018). Persistence over Gyr timescales, including possible “re-inflation” as host stars brighten, implies ongoing cooling suppression and/or energy deposition rather than primordial heat retention (E. D. Lopez & J. J. Fortney 2016; S. K. Grunblatt et al. 2017; T. D. Komacek et al. 2020; D. P. Thorngren et al. 2021).

Models either slow entropy loss or deposit absorbed stellar power at depth (S. Ginzburg & R. Sari 2015). Delayed cooling invokes enhanced atmospheric/interior opacities, compositional gradients, and double-diffusive

(“layered”) convection (A. Burrows et al. 2007; G. Chabrier & I. Baraffe 2007; J. Leconte & G. Chabrier 2012); three-dimensional circulation may raise the radiative-convective boundary pressure and insulate the interior (P. Tremblin et al. 2017), though delayed cooling alone struggles with re-inflation. Active channels include:

- **ohmic dissipation** from ionized winds, efficient near $T_{\text{eq}} \sim 1500\text{K}$ but weakened by low ionization or magnetohydrodynamic drag (K. Batygin & D. J. Stevenson 2010; K. Menou 2012; T. M. Rogers & T. D. Komacek 2014);
- **thermal tides** from stellar-heating-driven density perturbations that sustain asynchronous rotation and dissipate orbital/spin energy (P. Arras & A. Socrates 2009; A. Socrates 2013; D. P. Thorngren & J. J. Fortney 2018);
- **kinetic/mixing heating** by circulation, turbulence, or wave breaking at depth (A. N. Youdin & J. L. Mitchell 2010; F. Sainsbury-Martinez et al. 2019); and
- **gravitational tides** from maintained eccentricity, obliquity, or asynchronism, otherwise transient under circularization and synchronization (P. Bodenheimer et al. 2001; B. Jackson et al. 2008; S. Millholland et al. 2020).

Empirical inversions require anomalous powers of 10^{26} – 10^{27} erg s^{−1}, with efficiencies peaking at intermediate irradiation and tapering for the hottest planets (D. P. Thorngren & J. J. Fortney 2018; P. Sarkis et al. 2021). Distinguishing mechanisms requires constraints on wind speeds, magnetic fields, heat-deposition depths (D. Thorngren et al. 2019), and tidal responses via atmospheric characterization and structural probes such as Love numbers (U. Kramm et al. 2012).

Causal discovery offers a complementary, data-driven route for this discrimination (P. Spirtes et al. 2000). Rather than fitting one inflation model at a time, it asks which variables retain direct conditional dependence on the target observable after the remaining variables are accounted for. Built on the foundation laid out by P. Spirtes et al. (2000) and J. Pearl (2009), causal discovery has been widely used in epidemiology, genomics, condensed matter physics, and economics (N. Friedman 2004; K. Sachs et al. 2005; J. Runge et al. 2019). In recent years, a number of applications in astrophysics are also emerging. For example, to infer causal structure in supermassive black-hole–galaxy coevolution (M. Pasquato et al. 2023; Z. Jin et al. 2024, 2025b; B. L.

Davis et al. 2026), to test whether trans-Neptunian-object colors are primordial (B. L. Davis et al. 2025), to study the effect of environment on star formation (S. Mucesh et al. 2024), to model stellar mass estimation (W. Zhang et al. 2025), and to decipher the stellar migration history in the Milky Way (Z. Jin et al. 2025a).

Here we apply the same logic to hot-Jupiter inflation: we compare the direct graph parents of R_p identified from the observational data with the expected signatures of delayed cooling, ohmic dissipation, thermal tides, kinetic/mechanical heating, and gravitational tides. We start by deriving the radius-excess scalings used to translate each physical mechanism into predicted dependencies on M_p , P_{orb} , and stellar T_{eff} ; we then describe the planet samples and causal-discovery pipeline, validate the approach on super-Earths, and interpret the hot-Jupiter graph in light of these theoretical scalings.

The remainder of this Letter is organized as follows. §2 presents the analytic radius-excess scalings, planet samples, and causal-discovery pipeline. §3 validates the method on a super-Earth control sample and applies it to the hot-Jupiter sample, interpreting the recovered graph in the context of delayed cooling, ohmic dissipation, thermal tides, kinetic/mechanical heating, and gravitational tides. §4 summarizes the implications for irradiation-regulated inflation and outlines the follow-up radius-excess tests needed to separate the remaining viable mechanisms. Appendices A and B give the detailed theoretical derivations and causal-discovery formalism.

2. METHODS

2.1. Analytic Radius-excess Scalings

We compare the causal-discovery graph to theoretical power-law scalings for the fractional radius excess, $\Delta = (R_p - R_0)/R_0$ where R_0 is the radius expected from a non-inflated cooling track. Here we show only the resulting dependence on M_p , P_{orb} , and stellar T_{eff} ; the full derivation, assumptions, and mechanism-by-mechanism table are moved to Appendix A. The scalings use $F_{\text{inc}} \propto T_{\text{eff}}^4 P_{\text{orb}}^{-4/3}$, $T_{\text{eq}} \propto T_{\text{eff}} P_{\text{orb}}^{-1/3}$, $R_0 \propto M_p^{-0.06}$, $T_{\text{int}} \propto M_p^{0.50}$, and an active-heating response exponent $\chi = 0.20$. For active heating, deposited power is converted to radius excess through $\Delta_i \propto (P_i/L_{\text{int}})^\chi$; delayed-cooling models are treated separately as low-irradiation benchmarks. For compactness, the mechanism scalings used below are summarized as

$$\Delta_i \propto M_p^{\alpha_M} P_{\text{orb}}^{\alpha_P} T_{\text{eff}}^{\alpha_T}, \quad (1)$$

with $(\alpha_M, \alpha_P, \alpha_T)$ listed in Table 1.

The key comparison is whether a mechanism predicts direct radius-excess dependence on irradiation, period, and/or mass in Eq. (1). The mass exponent is especially

Table 1. Numerical Local Theoretical Scalings for the Radius-excess Law $\Delta_i \propto M_p^{\alpha_M} P_{\text{orb}}^{\alpha_P} T_{\text{eff}}^{\alpha_T}$.

Scenario	α_M	α_P	α_T	Dominant Causal Implication	Comparison to Data-driven Causal Graph
Ohmic dissipation, pre-turnover	-0.62	-0.33	+1.00	$M_p \rightarrow R_p, T_{\text{eff}} \rightarrow R_p$	Dominant $M_p \rightarrow R_p$ not identified in observation
Ohmic dissipation, near efficiency peak	-0.62	-0.27	+0.80	$M_p \rightarrow R_p, T_{\text{eff}} \rightarrow R_p$	Dominant $M_p \rightarrow R_p$ not identified in observation
Thermal tides, Gold-Soter	-0.42	-0.60	+0.60	$P_{\text{orb}} \rightarrow R_p, T_{\text{eff}} \rightarrow R_p$	Cleanest match
Thermal tides, dynamical	-0.44	-0.80	0.00	$M_p \rightarrow R_p, P_{\text{orb}} \rightarrow R_p$	Missing $T_{\text{eff}} \rightarrow R_p$
Kinetic / mechanical heating	-0.40	-0.27	+0.80	$M_p \rightarrow R_p, T_{\text{eff}} \rightarrow R_p$	Dominant $M_p \rightarrow R_p$ not identified in observation
Eccentricity / obliquity tides	-0.44	-1.00	0.00	$M_p \rightarrow R_p, P_{\text{orb}} \rightarrow R_p$	Missing $T_{\text{eff}} \rightarrow R_p$
Radiative opacity / delayed cooling	-0.78	-0.08	+0.25	$M_p \rightarrow R_p, T_{\text{eff}} \rightarrow R_p$	Missing $P_{\text{orb}} \rightarrow R_p$
Layered convection, fixed layer strength	-1.06	0.00	0.00	$M_p \rightarrow R_p$	Missing $P_{\text{orb}} \rightarrow R_p, T_{\text{eff}} \rightarrow R_p$
Potential-temperature advection	-1.06	-0.33	+1.00	$M_p \rightarrow R_p, T_{\text{eff}} \rightarrow R_p$	Dominant $M_p \rightarrow R_p$ not identified in observation

NOTE— The entries are the exponents in Eq. (1): α_M exponentiates M_p , α_P exponentiates P_{orb} , and α_T exponentiates stellar T_{eff} . The values assume $R_0 \propto M_p^{-0.06}$, $T_{\text{int}} \propto M_p^{0.50}$ at fixed age for active-heating channels, $\chi = 0.20$ for active deep heating, fixed stellar mass and radius, fixed age and composition, fixed B , fixed ϵ or obliquity, and fixed tidal Q'_p . The first ohmic row applies on the rising, pre-turnover branch; the second applies near the empirical heating-efficiency maximum. The dominant causal implication is the qualitative expected parent set; the mapping from exponent magnitude to graph edge is heuristic because causal discovery explores conditional dependence, not the local power-law slope.

151 useful because it separates mechanisms whose inflation
 152 efficiency is mainly limited by higher surface gravity or
 153 faster intrinsic cooling from those whose observed sig-
 154 nature is mostly set by irradiation. Negative α_M means
 155 that, at fixed P_{orb} and T_{eff} , lower-mass Jovian planets
 156 should show a larger fractional excess; values near -0.4
 157 to -0.6 are comparatively moderate, while $\alpha_M \simeq -1$
 158 would imply a much stronger residual mass dependence.
 159 Irradiation-dominated channels such as ohmic dissipa-
 160 tion and kinetic/mechanical heating predict strong posi-
 161 tive T_{eff} dependence and weaker negative period depen-
 162 dence; Gold-Soter thermal tides (T. Gold & S. Soter
 163 1969; A. Socrates 2013) predict both T_{eff} and P_{orb} de-
 164 pendencies of more comparable strength; purely grav-
 165 itational eccentricity/obliquity tides predict a strong
 166 period dependence but no leading stellar-temperature
 167 dependence. Delayed-cooling models serve as weak-
 168 irradiation-response benchmarks, although some have
 169 steep mass exponents because radius perturbations are
 170 suppressed by surface gravity. The numerical exponents
 171 used in §3 are listed in Table 1.

172 2.2. Data & Causal Discovery

173 We constructed the hot-Jupiter sample from the
 174 Encyclopaedia of exoplanetary systems (J. Schnei-
 175 der et al. 2011) using the same cuts described in Ap-
 176 pendix B: $0.5 M_{\oplus} \leq M_p \leq 2.0 M_{\oplus}$, $P_{\text{orb}} \leq 10$ days, and
 177 $0.5 R_{\oplus} < R_p \leq 3.0 R_{\oplus}$, yielding 328 planets with com-
 178 plete values for the filtered quantities. A sanity-check
 179 super-Earths dataset of 179 objects was also constructed
 180 from the same source, taking $R_p \leq 1.8 R_{\oplus}$.

181 Uncovering causal structures from passively observed
 182 data is possible since different causal structures lead to
 183 distinct sets of conditional independence relationships
 184 among data. For example, $A \leftarrow C \rightarrow B$ implies $A \not\perp\!\!\!\perp B$

185 and $A \perp\!\!\!\perp B \mid C^6$, but $A \rightarrow C \leftarrow B$ gives $A \perp\!\!\!\perp B$
 186 and $A \not\perp\!\!\!\perp B \mid C$. Therefore, one can go over possible
 187 causal graphs made by variables of interest, checking
 188 the likelihood that the observed data is compatible with
 189 the conditional independence relationships implied by
 190 the causal graph, until arriving at a causal graph with
 191 the highest likelihood.

192 We infer the causal structure among M_p , R_p , P_{orb} ,
 193 and T_{eff} using a search strategy called the Best Order
 194 Score Search (BOSS; B. Andrews et al. 2023) with a
 195 likelihood defined by the Fourier Feature Marginal Like-
 196 lihood (FFML) score (J. D. Ramsey 2026). All four
 197 variables were continuous and standardized before ker-
 198 nel scoring, making the analysis insensitive to physical
 199 units.

200 In brief, BOSS searches over variable orderings and
 201 projects each ordering to a directed acyclic graph
 202 (DAG, a graphical representation of causal structures)
 203 by adding or removing candidate parents when doing
 204 so improves the decomposable FFML score. FFML
 205 evaluates nonlinear conditional relationships by approx-
 206 imating an RBF-kernel Gaussian-process marginal like-
 207 lihood with random Fourier features (C. E. Rasmussen
 208 & C. K. I. Williams 2006; A. Rahimi & B. Recht 2007),
 209 allowing each candidate edge to be tested by its improve-
 210 ment to the child variable’s conditional model. Graph
 211 edges are interpreted as conditional dependencies, not
 212 complete causal histories. The full dataset description,
 213 pair-plot figure, ordering search, grow-shrink parent se-
 214 lection, and FFML equations are given in Appendix B.

⁶ “ $A \rightarrow C$ ” denotes “ A causes C ”, or A is the causal parent of C .
 “ $A \not\perp\!\!\!\perp B$ ” reads as “ A is dependent on B ,” and “ $A \perp\!\!\!\perp B \mid C$ ”
 reads as “ A is independent of B when conditioned on C .”




Figure 1. The causal structure found from super-Earth data using BOSS with FFML score. An undirected edge $A - B$ suggests both directions are possible (either $A \rightarrow B$ or $A \leftarrow B$), as long as no new colliders are created. Here, the graph means that $P \rightarrow M_p \rightarrow R_p$, $P \leftarrow M_p \leftarrow R_p$, or $P \leftarrow M_p \rightarrow R_p$ (but **not** $P \rightarrow M_p \leftarrow R_p$), plus an isolated node T_* , are equally possible statistically and are together the most likely causal structure.

3. RESULTS

3.1. Super-Earths baseline

As a sanity-check we first run our pipeline on the super-Earths dataset. Results are shown in Fig. 1. The recovered graph is consistent with the expected behavior of compact, mostly rocky or atmosphere-poor planets. In this regime, R_p is expected to be governed primarily by mass and bulk composition, while irradiation mainly affects the probability of envelope loss rather than the radii of already stripped cores (S. Seager et al. 2007). The M_p - P edge is also meaningful, reflecting the role of formation and migration in setting which planet masses are delivered to close-in orbits and survive there. In contrast, the isolated T_{eff} node is consistent with the sample selection. Once planets have lost their primordial envelopes, stellar effective temperature no longer strongly controls their observed radii; irradiation matters mainly through atmospheric erosion (J. E. Owen & Y. Wu 2013), not through the structure of bare rocky cores. With background knowledge above, we can further orient the DAG as: $P \leftarrow M_p \rightarrow R_p, T_*$.

3.2. Inflation of hot Jupiters

For the hot-Jupiter sample, the preferred four-variable graph in Fig. 2 contains

$$P_{\text{orb}} \rightarrow M_p, \quad T_{\text{eff}} \rightarrow M_p, \quad P_{\text{orb}} \rightarrow R_p, \quad T_{\text{eff}} \rightarrow R_p. \quad (2)$$

Thus R_p retains direct conditional dependence on orbital period and stellar effective temperature, but not on planet mass. Because $F_{\text{inc}} \propto T_{\text{eff}}^4 P_{\text{orb}}^{-4/3}$ (under the fixed- M_* , fixed- R_* scaling) is not an explicit graph variable, the paired $T_{\text{eff}} \rightarrow R_p$ and $P_{\text{orb}} \rightarrow R_p$ edges are the graph-level signature of irradiation-regulated inflation. The arrows into M_p should instead be read as population-level formation, migration, tidal-survival,




Figure 2. Preferred hot-Jupiter causal graph from BOSS with the FFML score. The inflation-relevant result is that R_p has direct parents P_{orb} and T_{eff} , but no direct $M_p \rightarrow R_p$ edge after conditioning.

photoevaporation/Roche-limit, and detection/selection effects (E. B. Ford & F. A. Rasio 2006; R. I. Dawson & J. A. Johnson 2018).

The graph is not a measurement of the analytic exponents in Table 1: an edge gives conditional dependence, not a slope. The missing $M_p \rightarrow R_p$ edge does not eliminate moderate mass scalings, because the graph is fitted to R_p whereas theory describes $\Delta = (R_p - R_0)/R_0$. Across the restricted Jovian mass range, the shallow non-inflated mass-radius relation, age, heavy-element content, entropy, opacity, and measurement scatter can dilute a mass trend in Δ .

The most compatible single mechanisms have both T_{eff} and P_{orb} dependences and no dominant residual mass term. Gold-Soter thermal tides provide the cleanest parent-set match:

$$\Delta_{\text{th,GS}} \propto M_p^{-0.42} P_{\text{orb}}^{-0.60} T_{\text{eff}}^{0.60},$$

so the two selected parents of R_p enter at least as strongly as mass. Kinetic/mechanical heating remains compatible,

$$\Delta_{\text{kin}} \propto M_p^{-0.40} P_{\text{orb}}^{-0.27} T_{\text{eff}}^{0.80},$$

but its weaker period exponent makes the absent mass edge require more dilution by the use of R_p rather than Δ , the narrow mass range, or astrophysical scatter. Ohmic heating, especially near the empirical efficiency peak, is also viable through Eqs. (A12) and (A13), although its stronger mass exponent makes the missing $M_p \rightarrow R_p$ edge less natural. Thus Gold-Soter thermal tides provide the closest single-mechanism match to the recovered parent set, but the graph alone does not uniquely identify the underlying heating physics. This conservative reading is consistent with the shallow-heating inference of S. P. Schmidt et al. (2026), who find that hot-Jupiter cooling rates are suppressed by $\simeq 95$ – 98% and argue that heating near or just below the radiative-convective boundary can enable reinflation with much weaker deep heating. We note that a

flux-controlled radius excess with little residual P_{orb} dependence would favor shallow ohmic dissipation or temperature advection, whereas a residual period term at fixed F_{inc} would favor a thermal-tide contribution.

This parent set disfavors purely period-controlled gravitational tides as the sole population-level explanation. Gold–Soter thermal tides are driven by stellar heating and predict both T_{eff} and period dependencies, whereas dynamical thermal tides and maintained eccentricity/obliquity tides, Eqs. (A17) and (A21), predict strong period dependencies but no leading T_{eff} dependence. Delayed-cooling channels are not excluded, but radiative blanketing and fixed layered convection, Eqs. (A24) and (A25), have weak or absent direct irradiation/period dependencies. Potential-temperature advection, Eq. (A26), has the correct irradiation signature albeit with a steep mass exponent, making it more plausible as a contributor than as the cleanest single explanation.

The sharper discriminator is whether Δ retains residual P_{orb} dependence after specifying F_{inc} . Existing empirical radius–mass–flux relations and heating-efficiency estimates mostly use F_{inc} or T_{eq} , rather than independent local slopes in T_{eff} and P_{orb} conditioned on age, composition, stellar properties, and selection effects (B.-O. Demory & S. Seager 2011; L. M. Weiss et al. 2013; D. P. Thorngren & J. J. Fortney 2018). Using Eq. (A3), irradiation-dominated mechanisms become

$$\Delta_{\text{kin}} \propto M_p^{-0.40} F_{\text{inc}}^{0.20}, \quad \Delta_{\Omega, \beta\Omega=1} \propto M_p^{-0.62} F_{\text{inc}}^{0.20},$$

with no leading residual period term, while Gold–Soter thermal tides become

$$\Delta_{\text{th,GS}} \propto M_p^{-0.42} F_{\text{inc}}^{0.15} P_{\text{orb}}^{-0.40}.$$

A residual P_{orb} dependence at fixed F_{inc} would favor Gold–Soter; a flux-controlled excess would favor kinetic/mechanical or ohmic heating.

Two-source models remain physically attractive. A kinetic/mechanical plus Gold–Soter mixture,

$$\Delta_{\text{mix}} = A_{\text{kin}} \Delta_{\text{kin}} + A_{\text{GS}} \Delta_{\text{th,GS}}, \quad (3)$$

selects the observed parents with moderate mass exponents and interpolates between irradiation-dominated and balanced period–temperature slopes. Locally,

$$\alpha_{\text{eff}} \simeq (1 - f_{\text{GS}}) \alpha_{\text{kin}} + f_{\text{GS}} \alpha_{\text{th,GS}}, \quad (4)$$

with entries ordered as $(M_p, P_{\text{orb}}, T_{\text{eff}})$, so increasing f_{GS} strengthens the period dependence with little change in mass dependence. An ohmic plus Gold–Soter mixture,

$$\Delta_{\text{mix}} = A_{\Omega} \Delta_{\Omega} + A_{\text{GS}} \Delta_{\text{th,GS}}, \quad (5)$$

is plausible if many planets lie near the ohmic efficiency peak, but requires more dilution of the ohmic mass scaling by R_p -rather-than- Δ fitting, composition scatter, or age diversity.

At the present level, a single-source model is more parsimonious and a two-source model more flexible. The robust conclusion is that hot-Jupiter inflation is irradiation-regulated; purely period-controlled gravitational tides are disfavored as the sole population-level channel; and Gold–Soter thermal tides, kinetic/mechanical heating, and ohmic dissipation remain viable. The decisive follow-up is a radius-excess analysis in $(M_p, F_{\text{inc}}, P_{\text{orb}})$, including stellar properties, age, metallicity, composition, and selection effects. Variation between moderately irradiated and ultra-hot planets would point to mixed mechanisms rather than a universal source.

4. SUMMARY & CONCLUSIONS

We applied nonlinear causal discovery to test which observed properties retain direct information about hot-Jupiter radii after conditioning on the others. Using M_p , R_p , P_{orb} , and stellar T_{eff} , we compared the recovered graph with analytic scalings of the form

$$\Delta \propto M_p^{\alpha_M} P_{\text{orb}}^{\alpha_P} T_{\text{eff}}^{\alpha_T},$$

where $\Delta = (R_p - R_0)/R_0$ is the fractional radius excess. The main conclusions are:

- As a control, the super-Earth sample recovers the expected mass–radius connection and leaves stellar T_{eff} isolated, consistent with radii set mainly by interior structure once primordial envelopes are lost.
- For hot Jupiters, the preferred graph selects $P_{\text{orb}} \rightarrow R_p$ and $T_{\text{eff}} \rightarrow R_p$, but no direct $M_p \rightarrow R_p$ edge after conditioning on the other variables.
- Because $F_{\text{inc}} \propto T_{\text{eff}}^4 P_{\text{orb}}^{-4/3}$, the simultaneous dependence of R_p on T_{eff} and P_{orb} is naturally interpreted as evidence for irradiation-regulated inflation.
- The absence of a direct $M_p \rightarrow R_p$ edge does not imply that mass is physically irrelevant. The graph is fit to R_p , whereas the theory describes Δ ; over the restricted Jovian mass range, composition, age, opacity, entropy, and measurement scatter can dilute moderate mass trends.
- Purely period-controlled gravitational tides are disfavored as the sole population-level explanation

because they lack a leading stellar-temperature dependence. At the graph-parent-set level, Gold–Soter thermal tides are the most compatible single-mechanism candidate: their T_{eff} and P_{orb} dependencies both exceed the residual mass exponent. Kinetic/mechanical heating follows, with strong irradiation but weaker period response. Ohmic dissipation remains viable, though its stronger mass scaling makes the missing $M_p \rightarrow R_p$ edge less straightforward.

The present four-variable graph constrains the parent set of R_p , but not the residual-period behavior or local slopes of Δ . At the graph level, Gold–Soter thermal tides are the cleanest single-source match, followed by kinetic/mechanical heating, with ohmic dissipation still viable. A follow-up analysis should infer partial slopes for Δ in $(M_p, F_{\text{inc}}, P_{\text{orb}})$, including stellar properties, age, metallicity, composition, and selection effects. A residual P_{orb} dependence at fixed F_{inc} would favor Gold–Soter tides; a mainly flux-controlled excess would favor kinetic/mechanical heating or ohmic dissipation.

ACKNOWLEDGMENTS

This material is based upon work supported by Tamkeen under the NYU Abu Dhabi Research Institute grant

CASS. This research has made use of NASA’s Astrophysics Data System. This research was carried out on the high-performance computing resources at New York University Abu Dhabi.

Software: `causal-learn` (Y. Zheng et al. 2023), `Matplotlib` (J. D. Hunter 2007), `NetworkX` (A. A. Hagberg et al. 2008), `NumPy` (C. R. Harris et al. 2020), `Pandas` (W. McKinney 2010), `pgmpy` (A. Ankan & A. Panda 2015), `PyGraphviz`, `Python` (G. Van Rossum & F. L. Drake 2009), `SciPy` (P. Virtanen et al. 2020), `seaborn` (M. Waskom 2021)

ORCID IDS

Zehao Jin (金泽灏) 
<https://orcid.org/0009-0000-2506-6645>
 Mohamad Ali-Dib 
<https://orcid.org/0000-0002-6633-376X>
 Mario Pasquato 
<https://orcid.org/0000-0003-3784-5245>
 Yujia Zheng (郑雨嘉) 
<https://orcid.org/0009-0003-5225-6366>
 Benjamin L. Davis 
<https://orcid.org/0000-0002-4306-5950>
 Andrea V. Macciò 
<https://orcid.org/0000-0002-8171-6507>

APPENDIX

A. DETAILED ANALYTIC RADIUS-EXCESS SCALINGS

The purpose of this section is to obtain explicit local power-law predictions for the fractional radius excess,

$$\Delta_i \equiv \frac{R_p - R_0}{R_0}, \quad (\text{A1})$$

where $R_0(M_p, t, Z)$ is the radius expected from a non-inflated cooling track at the same mass, age, and composition. The derived scalings describe the inflation term Δ_i , not the absolute transit radius R_p . Throughout this section T_{eff} denotes the stellar effective temperature. The planet equilibrium temperature T_{eq} is used only as an intermediate irradiation variable; the final radius-excess scalings are written in terms of M_p , P_{orb} , and stellar T_{eff} .

The incident flux and equilibrium temperature are

$$F_{\text{inc}} = \sigma_{\text{SB}} T_{\text{eff}}^4 \left(\frac{R_\star}{a} \right)^2 \propto T_{\text{eff}}^4 R_\star^2 M_\star^{-2/3} P_{\text{orb}}^{-4/3},$$

$$T_{\text{eq}} = \left[\frac{(1 - A_B) F_{\text{inc}}}{4\sigma_{\text{SB}}} \right]^{1/4} \propto T_{\text{eff}} R_\star^{1/2} M_\star^{-1/6} P_{\text{orb}}^{-1/3}. \quad (\text{A2})$$

In the scalings below, we hold M_\star , R_\star , albedo, age, and composition fixed, or equivalently absorb the residual stellar-radius and stellar-mass factors into the normalization. Thus the working conversion is

$$F_{\text{inc}} \propto T_{\text{eff}}^4 P_{\text{orb}}^{-4/3}, \quad T_{\text{eq}} \propto T_{\text{eff}} P_{\text{orb}}^{-1/3}. \quad (\text{A3})$$

We now replace the free response parameters by representative numerical values. For the uninflated giant-planet mass–radius slope we adopt

$$R_0 \propto M_p^{\xi_R}, \quad \xi_R = -0.06, \quad (\text{A4})$$

consistent with the nearly flat mass–radius relation of Jovian planets over the hot-Jupiter mass range (A. Burrows et al. 2003; D. P. Thorngren & J. J. Fortney 2018). For the intrinsic luminosity, we use a fixed-age Kelvin–Helmholtz scaling. With R_0 only weakly dependent on mass, $L_{\text{int}} \sim GM_p^2/(R_0 t)$, and $L_{\text{int}} \sim R_0^2 T_{\text{int}}^4$, giving $T_{\text{int}} \propto M_p^{1/2}$ to the accuracy needed here. Thus we adopt

$$\xi_T \equiv \frac{\partial \ln T_{\text{int}}}{\partial \ln M_p} = 0.50. \quad (\text{A5})$$

For active deep heating, we adopt a radius-response exponent

$$\chi = 0.20, \quad (\text{A6})$$

which is the order-of-magnitude value found by analytic deep-deposition theory and comparison to numerical models, where the radius perturbation scales approximately as the deposited-power parameter to the power 0.15 to 0.20 (S. Ginzburg & R. Sari 2015; D. S. Spiegel & A. Burrows 2013). This value should be read as a local response exponent for mildly to moderately inflated planets, not as a universal structural constant.

For any active heating mechanism with

$$P_i \propto R_0^{s_i} M_p^{u_i} F_{\text{inc}}^{a_i} P_{\text{orb}}^{-b_i}, \quad (\text{A7})$$

the local cooling-time-consistent response is

$$\Delta_i \propto \left(\frac{P_i}{L_{\text{int}}} \right)^\chi \simeq \left(\frac{P_i}{R_0^2 T_{\text{int}}^4} \right)^\chi. \quad (\text{A8})$$

Using Eq. (A3), this gives

$$\Delta_i \propto M_p^{[u_i + (s_i - 2)\xi_R - 4\xi_T]\chi} T_{\text{eff}}^{4a_i\chi} P_{\text{orb}}^{-(b_i + 4a_i/3)\chi} \propto M_p^{0.20[u_i - 0.06(s_i - 2) - 2.00]} T_{\text{eff}}^{0.80a_i} P_{\text{orb}}^{-0.20(b_i + 4a_i/3)}. \quad (\text{A9})$$

The right-hand side is the numerical form used below. The -2.00 term is the adopted T_{int}^4 cooling penalty.

A.1. Active-heating channels

A.1.1. Ohmic dissipation

A wind-driven ohmic estimate gives

$$P_\Omega \propto B^2 R_0^4 M_p^{-1} F_{\text{inc}}^{\beta_\Omega} \propto B^2 R_0^4 M_p^{-1} T_{\text{eff}}^{4\beta_\Omega} P_{\text{orb}}^{-4\beta_\Omega/3}. \quad (\text{A10})$$

Menou-type scalings imply that ohmic power rises faster than the absorbed stellar power on the cool side of the inflation sequence, while magnetic drag produces a turnover along the equivalent equilibrium-temperature coordinate near $T_{\text{eq}} \simeq 1500$ to 2000 K (K. Menou 2012; T. M. Rogers & T. D. Komacek 2014; D. P. Thorngren & J. J. Fortney 2018). In the final variables, this turnover is a turnover in the irradiation combination $T_{\text{eff}} P_{\text{orb}}^{-1/3}$. We adopt a representative pre-turnover local slope

$$\beta_\Omega = 1.25, \quad (\text{A11})$$

which corresponds to $(s, u, a, b) = (4, -1, 1.25, 0)$ in Eq. (A9). The resulting scaling is

$$\Delta_\Omega \propto M_p^{-0.62} P_{\text{orb}}^{-0.33} T_{\text{eff}}^{1.00} \quad (\text{A12})$$

for fixed B . At the empirical efficiency peak, the effective β_Ω is closer to unity, which gives

$$\Delta_{\Omega, \beta_\Omega=1} \propto M_p^{-0.62} P_{\text{orb}}^{-0.27} T_{\text{eff}}^{0.80}. \quad (\text{A13})$$

A.1.2. *Thermal tides*

For the Gold-Soter thermal-tide estimate,

$$P_{\text{th,GS}} \propto R_0^4 P_{\text{orb}}^{-2} T_{\text{eq}}^3 \propto R_0^4 T_{\text{eff}}^3 P_{\text{orb}}^{-3} \propto R_0^4 P_{\text{orb}}^{-2} F_{\text{inc}}^{3/4} \quad (\text{A14})$$

(A. Socrates 2013). Thus $(s, u, a, b) = (4, 0, 3/4, 2)$ and

$$\Delta_{\text{th,GS}} \propto M_p^{-0.42} P_{\text{orb}}^{-0.60} T_{\text{eff}}^{0.60}. \quad (\text{A15})$$

The dynamical thermal-tide estimate instead has

$$P_{\text{th,dyn}} \propto R_0^5 P_{\text{orb}}^{-4} \quad (\text{A16})$$

at fixed Q'_p and a normalized forcing frequency. With $(s, u, a, b) = (5, 0, 0, 4)$,

$$\Delta_{\text{th,dyn}} \propto M_p^{-0.44} P_{\text{orb}}^{-0.80} T_{\text{eff}}^0. \quad (\text{A17})$$

If the forcing-frequency normalization is assumed to vary with the mean motion, the period exponent becomes even simpler.

A.1.3. *Kinetic or mechanical heat burial*

For mechanical heating or heat burial, we use

$$P_{\text{kin}} = \epsilon_{\text{kin}} \pi R_0^2 F_{\text{inc}} \propto \epsilon_{\text{kin}} R_0^2 T_{\text{eff}}^4 P_{\text{orb}}^{-4/3}. \quad (\text{A18})$$

Youdin and Mitchell's mechanical-greenhouse model does not require a universal power-law increase of ϵ_{kin} with irradiation, so we use the constant-efficiency closure $\epsilon_{\text{kin}} = \text{const.}$ (A. N. Youdin & J. L. Mitchell 2010). This gives $(s, u, a, b) = (2, 0, 1, 0)$ and

$$\Delta_{\text{kin}} \propto M_p^{-0.40} P_{\text{orb}}^{-0.27} T_{\text{eff}}^{0.80}. \quad (\text{A19})$$

A.1.4. *Maintained eccentricity or obliquity tides*

For eccentricity tides in a synchronously rotating planet,

$$P_{\text{tide}} \propto R_0^5 \left(\frac{e^2}{Q'_p} \right) P_{\text{orb}}^{-5} \quad (\text{A20})$$

at fixed stellar mass. Obliquity tides have the same role after replacing e^2/Q'_p by the corresponding obliquity-damping factor. With $(s, u, a, b) = (5, 0, 0, 5)$,

$$\Delta_{\text{tide}} \propto M_p^{-0.44} P_{\text{orb}}^{-1.00} T_{\text{eff}}^0 \quad (\text{A21})$$

for fixed e or obliquity and fixed Q'_p .

A.2. *Delayed-cooling and circulation channels*

Delayed-cooling mechanisms do not correspond to a separate additive P_i and therefore are not described by Eq. (A9). For the pure radiative-blanketing or opacity-delayed-cooling limit, analytic irradiated-envelope models give a weak dependence of the interior entropy on irradiation (A. Burrows et al. 2007; S. Ginzburg & R. Sari 2015). At fixed age, the passive irradiated-envelope cooling relation gives approximately

$$T_c \propto M_p^{1/4} R_0^{-1/2} T_{\text{eq}}^{1/4}, \quad (\text{A22})$$

and the small-inflation radius response gives

$$\frac{\Delta R}{R_0} \propto \frac{T_c}{g R_0} \propto \frac{T_c R_0}{M_p}. \quad (\text{A23})$$

506 Using $R_0 \propto M_p^{-0.06}$ and $T_{\text{eq}} \propto T_{\text{eff}} P_{\text{orb}}^{-1/3}$ gives

$$507 \quad \boxed{\Delta_{\text{rad}} \propto M_p^{-0.78} P_{\text{orb}}^{-0.08} T_{\text{eff}}^{0.25}}. \quad (\text{A24})$$

508 Equation (A24) should be interpreted as the passive radiative-blanketing limit. It produces a much weaker period and
509 stellar-effective-temperature dependence than the active heating channels.

510 For layered convection with a fixed compositional-staircase strength, the leading dependence is through surface
511 gravity rather than irradiation. Approximating the imposed entropy perturbation as fixed gives

$$512 \quad \boxed{\Delta_{\text{layer}} \propto M_p^{-1.06} P_{\text{orb}}^0 T_{\text{eff}}^0}. \quad (\text{A25})$$

513 Thus layered convection can change the normalization of R_p substantially, but by itself it does not predict a direct
514 period or stellar-effective-temperature edge unless the layer strength correlates with irradiation.

515 For potential-temperature advection, circulation models show that non-uniform irradiation can drive the deep at-
516 mosphere toward a hotter adiabat than in one-dimensional radiative-convective models (P. Tremblin et al. 2017; F.
517 Sainsbury-Martinez et al. 2019). Those simulations do not provide a unique universal power law, so we close the ana-
518 lytic scaling with the simplest irradiation-set deep-adiabat assumption, $\delta T_{\text{deep}} \propto T_{\text{eq}}$. Using $\Delta R/R_0 \propto \delta T_{\text{deep}}/(gR_0)$
519 gives

$$520 \quad \boxed{\Delta_{\text{adv}} \propto M_p^{-1.06} P_{\text{orb}}^{-0.33} T_{\text{eff}}^{1.00}}. \quad (\text{A26})$$

521 The mass exponent is steep because a fixed deep-temperature perturbation produces a smaller fractional radius response
522 at higher surface gravity.

523 A.3. Empirical observational scalings

524 The empirical relations below are observational fits and should not be identified with any single physical channel.
525 They mix heating physics, age, composition, selection effects, stellar evolution, and correlated system parameters.

526 The numerical scalings in Table 1 provide the theoretical comparison set for the causal-discovery analysis. The
527 strongest direct period dependencies come from maintained gravitational tides and thermal tides. Ohmic dissipation,
528 kinetic heating, and potential-temperature advection predict stronger dependence on T_\star than on P because their leading
529 forcing is irradiation. Radiative-delayed cooling has a much weaker irradiation exponent, while layered convection with
530 fixed layer strength predicts no direct P or T_\star dependencies.

531 B. DETAILED METHODS

532 B.1. Dataset

533 Our dataset was constructed from the **Encyclopaedia of exoplanetary systems** (<https://exoplanet.eu/catalog/>)
534 by applying uniform cuts in mass, radius, and orbital period. Specifically, we retained only planets with cataloged
535 masses in the range $0.5 M_\oplus \leq M_p \leq 2.0 M_\oplus$, orbital periods $P_{\text{orb}} \leq 10$ days, and radii satisfying $0.5 R_\oplus < R_p \leq 3.0 R_\oplus$,
536 for 328 data points in total (Fig. 3). These criteria select short-period, Jupiter-scale exoplanets while excluding objects
537 with very small radii, extremely inflated radii, or masses outside the adopted giant-planet regime. Because the selection
538 was implemented through direct numerical comparisons, planets lacking reported values in any of the filtered quantities
539 were not included in the final sample. For the control sample, we constructed a super-Earth data set from the same
540 catalog using $R_p \leq 1.8 R_\oplus$ and requiring complete measurements of M_p , R_p , P_{orb} , and T_{eff} . This yielded 179 exoplanets
541 (Fig. 4). No additional causal-discovery variables were used in the control analysis.

542 B.2. Causal discovery

543 We inferred the causal structure among planetary mass M_p , planetary radius R_p , orbital period P_{orb} , and stellar
544 effective temperature T_{eff} using Best Order Score Search (BOSS; B. Andrews et al. 2023) with the Fourier Feature
545 Marginal Likelihood (FFML) score (J. D. Ramsey 2026).

546 B.2.1. Best Order Score Search (BOSS)

547 BOSS is a score-based method for learning directed acyclic graphs (DAGs). The original BOSS algorithm was
548 described with a BIC local score; here we use the same search procedure, but replace the local score by FFML to allow
549 nonlinear conditional relationships.

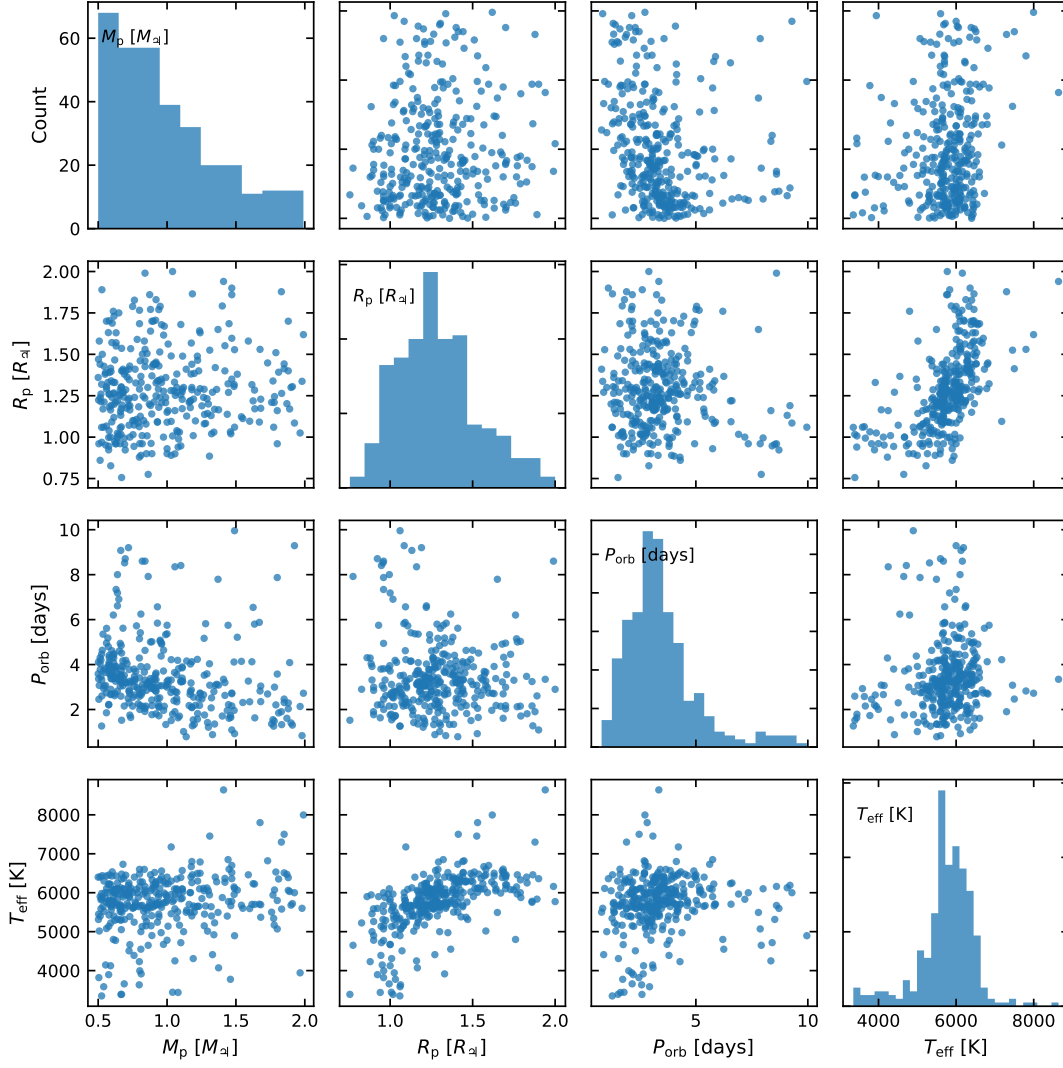


Figure 3. Pair-plot matrix for the filtered hot-Jupiter sample, showing planet mass (M_p), planet radius (R_p), orbital period (P_{orb}), and host-star effective temperature (T_{eff}) for 328 exoplanets.

550 Let X_1, \dots, X_p denote the variables, with $p = 4$ in the main analysis. For a DAG G , let $\text{Pa}_G(j)$ be the parent set
 551 of X_j . The score of G is decomposable:

$$552 \quad S(G) = \sum_{j=1}^p S_{\text{FFML}}(X_j \mid X_{\text{Pa}_G(j)}). \quad (\text{B27})$$

553 Thus each candidate edge is evaluated by asking whether it improves the conditional model of the child after accounting
 554 for its other parents.

555 BOSS searches over variable orderings rather than directly enumerating all DAGs. For an ordering π , define the
 556 prefix of X_j as

$$557 \quad \text{Pre}_\pi(j) = \{k : \pi(k) < \pi(j)\}. \quad (\text{B28})$$

558 Only variables in $\text{Pre}_\pi(j)$ are allowed to be parents of X_j , which guarantees that the resulting graph is acyclic. For a
 559 fixed ordering, BOSS projects the ordering to a DAG by selecting parents for each node using a grow-shrink procedure.
 560 Starting from an empty parent set $A = \emptyset$, the grow step repeatedly adds the prefix variable that gives the largest
 561 positive gain

$$562 \quad \Delta_{\text{add}}(k; A, j) = S_{\text{FFML}}(X_j \mid X_{A \cup \{k\}}) - S_{\text{FFML}}(X_j \mid X_A), \quad (\text{B29})$$



Figure 4. Pair-plot matrix for the filtered super-earths sample, showing planet mass (M_p), planet radius (R_p), orbital period (P_{orb}), and host-star effective temperature (T_{eff}) for 179 exoplanets.

563 until no addition improves the score. The shrink step then checks the selected parents and removes any parent whose
 564 deletion gives a positive gain

$$565 \quad \Delta_{\text{del}}(k; A, j) = S_{\text{FFML}}(X_j | X_{A \setminus \{k\}}) - S_{\text{FFML}}(X_j | X_A). \quad (\text{B30})$$

566 The final parent set is denoted $\widehat{\text{Pa}}_{\pi}(j)$, and the projected DAG has edges $X_k \rightarrow X_j$ for all $k \in \widehat{\text{Pa}}_{\pi}(j)$.

567 BOSS then improves the ordering itself. For each variable, it tries moving that variable to each possible position in
 568 the ordering, recomputes the projected-DAG score, and keeps the move only if it increases the total score in Eq. (B27).
 569 This sweep over variables is repeated until no single-variable move improves the score. Grow-shrink trees cache the
 570 local parent-set scores encountered during this process, so parent sets shared by nearby orderings are not recomputed.
 571 After the best ordering is found, BOSS converts the selected DAG into a partially-directed graph when orientations
 572 are not identifiable from the score and conditional independence structure alone. When FFML is used in DAG mode,
 573 it can also prefer one direction within a Markov-equivalent pair because FFML is not score equivalent.

B.2.2. *Search criteria: Fourier Feature Marginal Likelihood (FFML) score*

The FFML score is a nonlinear marginal-likelihood score. For each local regression, let $Y \in \mathbb{R}^n$ be the child variable observed for n planets, and let $Z \in \mathbb{R}^{n \times d}$ be a candidate parent set. FFML starts from the Gaussian-process model

$$Y = f(Z) + \epsilon, \quad \epsilon \sim \mathcal{N}(0, \sigma^2 I_n), \quad f \sim \mathcal{GP}(0, k), \quad (\text{B31})$$

where k is an RBF kernel. Let K_Z be the kernel matrix over the parent observations and define

$$C_Z = K_Z + \sigma^2 I_n. \quad (\text{B32})$$

After integrating out the unknown function f , the exact kernel marginal log-likelihood, up to constants independent of the parent set, is

$$S_{\text{KML}}(Y | Z) = -\frac{1}{2} Y^\top C_Z^{-1} Y - \frac{1}{2} \log |C_Z|. \quad (\text{B33})$$

The first term rewards fit, while the log-determinant term is an Occam factor that penalizes overly flexible parent sets.

Direct evaluation of Eq. (B33) requires operations on an $n \times n$ matrix. FFML avoids this by approximating the RBF kernel with random Fourier features. For

$$k(z, z') = \exp\left(-\frac{\|z - z'\|^2}{h^2}\right), \quad (\text{B34})$$

draw frequencies $\omega_\ell \sim \mathcal{N}(0, 2h^{-2}I_d)$ and phases $b_\ell \sim \text{Unif}(0, 2\pi)$, and define

$$\phi(z) = \sqrt{\frac{2}{m}} \begin{bmatrix} \cos(\omega_1^\top z + b_1) \\ \vdots \\ \cos(\omega_m^\top z + b_m) \end{bmatrix}. \quad (\text{B35})$$

Stacking $\phi(z_i)$ over all planets gives $\Phi \in \mathbb{R}^{n \times m}$, with

$$K_Z \approx \Phi \Phi^\top. \quad (\text{B36})$$

Define

$$H = \Phi^\top \Phi, \quad r = \Phi^\top Y. \quad (\text{B37})$$

Using the Woodbury identity and the matrix determinant lemma, FFML computes the local score as

$$S_{\text{FFML}}(Y | Z) = -\frac{1}{2\sigma^2} \left[Y^\top Y - r^\top (H + \sigma^2 I_m)^{-1} r \right] - \frac{1}{2} \left[(n - m) \log \sigma^2 + \log |H + \sigma^2 I_m| \right]. \quad (\text{B38})$$

All linear solves and log determinants are now evaluated in the m -dimensional feature space rather than the n -dimensional sample space. For one local FFML score with d parents and m random features, forming the feature matrix Φ costs $\mathcal{O}(ndm)$, forming $H = \Phi^\top \Phi$ costs $\mathcal{O}(nm^2)$, and factorizing $H + \sigma^2 I_m$ costs $\mathcal{O}(m^3)$. Thus the parent-set score cost is $\mathcal{O}(ndm + nm^2 + m^3)$, ignoring lower-order terms. For fixed d and $m \ll n$, this score evaluation is effectively linear in the number of planets. The total BOSS runtime further depends on the number of order moves and candidate parent sets evaluated. When the parent set is empty, FFML reduces to the Gaussian noise model $Y \sim \mathcal{N}(0, \sigma^2 I_n)$.

All variables in our analysis are continuous, so the mixed discrete-parent extension of FFML is not needed. Continuous variables are standardized for kernel evaluation, and the RBF bandwidth is selected from pairwise distances among candidate parent observations. The $\mathcal{O}(ndm + nm^2 + m^3)$ scaling refers to the FFML linear-algebra evaluation for a fixed bandwidth and random-feature basis. A naive all-pairs bandwidth selection for each candidate parent set would add $\mathcal{O}(n^2 d)$ to the corresponding local-score evaluation. The random feature basis is coupled by target variable, so local-score differences such as $S_{\text{FFML}}(Y | Z \cup \{X\}) - S_{\text{FFML}}(Y | Z)$ are evaluated with the same random features and are therefore more stable.

REFERENCES

- Andrews, B., Ramsey, J., Sanchez-Romero, R., Camchong, J., & Kummerfeld, E. 2023, [arXiv e-prints](#), [arXiv:2310.17679](#)
- Ankan, A., & Panda, A. 2015, in Proceedings of the 14th Python in Science Conference (SCIPY 2015), Citeseer

- 614 Arras, P., & Socrates, A. 2009, [arXiv e-prints](#),
615 [arXiv:0901.0735](#)
- 616 Batygin, K., & Stevenson, D. J. 2010, [ApJL](#), **714**, L238
- 617 Bodenheimer, P., Lin, D. N. C., & Mardling, R. A. 2001,
618 [ApJ](#), **548**, 466
- 619 Burrows, A., Hubeny, I., Budaj, J., & Hubbard, W. B.
620 2007, [ApJ](#), **661**, 502
- 621 Burrows, A., Sudarsky, D., & Hubbard, W. B. 2003, [ApJ](#),
622 **594**, 545
- 623 Chabrier, G., & Baraffe, I. 2007, [ApJL](#), **661**, L81
- 624 Charbonneau, D., Brown, T. M., Latham, D. W., & Mayor,
625 M. 2000, [ApJL](#), **529**, L45
- 626 Davis, B. L., Ali-Dib, M., Zheng, Y., et al. 2025, [MNRAS](#),
627 **543**, L34
- 628 Davis, B. L., More, S., Jin, Z., et al. 2026, [ApJ](#), **1000**, 100
- 629 Dawson, R. I., & Johnson, J. A. 2018, [ARA&A](#), **56**, 175–221
- 630 Demory, B.-O., & Seager, S. 2011, [ApJS](#), **197**, 12
- 631 Ford, E. B., & Rasio, F. A. 2006, [ApJL](#), **638**, L45
- 632 Fortney, J. J., Marley, M. S., & Barnes, J. W. 2007, [ApJ](#),
633 **659**, 1661
- 634 Friedman, N. 2004, [Science](#), **303**, 799
- 635 Ginzburg, S., & Sari, R. 2015, [ApJ](#), **803**, 111
- 636 Gold, T., & Soter, S. 1969, [Icarus](#), **11**, 356
- 637 Grunblatt, S. K., Huber, D., Gaidos, E., et al. 2017, [AJ](#),
638 **154**, 254
- 639 Hagberg, A. A., Schult, D. A., & Swart, P. J. 2008, in
640 Proceedings of the 7th Python in Science Conference, ed.
641 G. Varoquaux, T. Vaught, & J. Millman, Pasadena, CA
642 USA, 11 – 15
- 643 Harris, C. R., Millman, K. J., van der Walt, S. J., et al.
644 2020, [Nature](#), **585**, 357
- 645 Hunter, J. D. 2007, [Computing in Science and Engineering](#),
646 **9**, 90
- 647 Jackson, B., Greenberg, R., & Barnes, R. 2008, [ApJ](#), **681**,
648 **1631**
- 649 Jin, Z., Lu, Y., Ting, Y.-S., Zheng, Y., & Buck, T. 2025a,
650 [arXiv e-prints](#), [arXiv:2507.00134](#)
- 651 Jin, Z., Pasquato, M., Davis, B. L., Macciò, A. V., &
652 Hezaveh, Y. 2024, [arXiv e-prints](#), [arXiv:2410.14775](#)
- 653 Jin, Z., Pasquato, M., Davis, B. L., et al. 2025b, [ApJ](#), **979**,
654 **212**
- 655 Komacek, T. D., Thorngren, D. P., Lopez, E. D., &
656 Ginzburg, S. 2020, [ApJ](#), **893**, 36
- 657 Kramm, U., Nettelmann, N., Fortney, J. J., Neuhäuser, R.,
658 & Redmer, R. 2012, [A&A](#), **538**, A146
- 659 Laughlin, G., Crismani, M., & Adams, F. C. 2011, [ApJL](#),
660 **729**, L7
- 661 Leconte, J., & Chabrier, G. 2012, [A&A](#), **540**, A20
- 662 Lopez, E. D., & Fortney, J. J. 2016, [ApJ](#), **818**, 4
- 663 McKinney, W. 2010, in Proceedings of the 9th Python in
664 Science Conference, Vol. 445, Austin, TX, 51–56
- 665 Menou, K. 2012, [ApJ](#), **745**, 138
- 666 Millholland, S., Petigura, E., & Batygin, K. 2020, [ApJ](#), **897**,
667 **7**
- 668 Mucesh, S., Hartley, W. G., Gilligan-Lee, C. M., & Lahav,
669 O. 2024, [arXiv e-prints](#), [arXiv:2412.02439](#)
- 670 Owen, J. E., & Wu, Y. 2013, [ApJ](#), **775**, 105
- 671 Pasquato, M., Jin, Z., Lemos, P., Davis, B. L., & Macciò,
672 A. V. 2023, [arXiv e-prints](#), [arXiv:2311.15160](#)
- 673 Pearl, J. 2009, Causality (Cambridge university press)
- 674 Rahimi, A., & Recht, B. 2007, in Advances in Neural
675 Information Processing Systems 20, ed. J. C. Platt,
676 D. Koller, Y. Singer, & S. T. Roweis (Curran Associates,
677 Inc.), 1177–1184
- 678 Ramsey, J. D. 2026, [arXiv e-prints](#), [arXiv:2605.05743](#)
- 679 Rasmussen, C. E., & Williams, C. K. I. 2006, Gaussian
680 Processes for Machine Learning (Cambridge, MA: MIT
681 Press)
- 682 Rogers, T. M., & Komacek, T. D. 2014, [ApJ](#), **794**, 132
- 683 Runge, J., Bathiany, S., Bollt, E., et al. 2019, Nature
684 communications, **10**, 2553
- 685 Sachs, K., Perez, O., Pe’er, D., Lauffenburger, D. A., &
686 Nolan, G. P. 2005, [Science](#), **308**, 523
- 687 Sainsbury-Martinez, F., Wang, P., Fromang, S., et al. 2019,
688 [A&A](#), **632**, A114
- 689 Sarkis, P., Mordasini, C., Henning, T., Marleau, G. D., &
690 Mollière, P. 2021, [A&A](#), **645**, A79
- 691 Schmidt, S. P., Thorngren, D. P., & Schlaufman, K. C.
692 2026, [ApJ](#), **1001**, 35
- 693 Schneider, J., Dedieu, C., Le Sidaner, P., Savalle, R., &
694 Zolotukhin, I. 2011, [A&A](#), **532**, A79
- 695 Seager, S., Kuchner, M., Hier-Majumder, C. A., & Militzer,
696 B. 2007, [ApJ](#), **669**, 1279
- 697 Sestovic, M., Demory, B.-O., & Queloz, D. 2018, [A&A](#), **616**,
698 **A76**
- 699 Socrates, A. 2013, [arXiv e-prints](#), [arXiv:1304.4121](#)
- 700 Spiegel, D. S., & Burrows, A. 2013, [ApJ](#), **772**, 76
- 701 Spirtes, P., Glymour, C. N., & Scheines, R. 2000,
702 Causation, Prediction, and Search, 2nd edn. (Cambridge,
703 MA: MIT Press)
- 704 Thorngren, D., Gao, P., & Fortney, J. J. 2019, [ApJL](#), **884**,
705 **L6**
- 706 Thorngren, D. P., & Fortney, J. J. 2018, [AJ](#), **155**, 214
- 707 Thorngren, D. P., Fortney, J. J., Lopez, E. D., Berger,
708 T. A., & Huber, D. 2021, [ApJL](#), **909**, L16
- 709 Tremblin, P., Chabrier, G., Mayne, N. J., et al. 2017, [ApJ](#),
710 **841**, 30
- 711 Van Rossum, G., & Drake, F. L. 2009, Python 3 Reference
712 Manual (Scotts Valley, CA: CreateSpace)

- 713 Virtanen, P., Gommers, R., Oliphant, T. E., et al. 2020,
714 [Nature Methods](#), 17, 261
- 715 Waskom, M. 2021, [The Journal of Open Source Software](#), 6,
716 3021
- 717 Weiss, L. M., Marcy, G. W., Rowe, J. F., et al. 2013, [ApJ](#),
718 768, 14
- 719 Youdin, A. N., & Mitchell, J. L. 2010, [ApJ](#), 721, 1113
- 720 Zhang, W., Lin, Q., Ting, Y.-S., et al. 2025, [A&A](#), 703,
721 A276
- 722 Zheng, Y., Huang, B., Chen, W., et al. 2023, [arXiv e-prints](#),
723 [arXiv:2307.16405](#)