























A Carbon-rich Atmosphere on a Windy Pulsar Planet

Michael Zhang^{1,14} , Maya Belezny² , Timothy D. Brandt³ , Roger W. Romani² , Peter Gao⁴ , Hayley Beltz⁵ ,
Matthew Bailes^{6,15} , Matthew C. Nixon^{5,7,14} , Jacob L. Bean¹ , Thaddeus D. Komacek⁸ , Brandon P. Coy⁹ ,
Guangwei Fu¹⁰ , Rafael Luque^{1,16} , Daniel J. Reardon^{6,15} , Emma Carli^{6,15} , Ryan M. Shannon^{6,15} ,
Jonathan J. Fortney¹¹ , Anjali A. A. Piette¹² , M. Coleman Miller⁵ , and Jean-Michel Desert¹³ 

¹ University of Chicago, Department of Astronomy & Astrophysics, Chicago, IL 60637, USA; mzzhang2014@gmail.com

² Stanford University, KIPAC, Stanford, CA 94305, USA

³ Space Telescope Science Institute, Baltimore, MD 21218, USA

⁴ Carnegie Science, Washington, DC 20015 USA

⁵ University of Maryland, Department of Astronomy, College Park, MD 20742, USA

⁶ Swinburne University of Technology, Centre for Astrophysics and Supercomputing, Melbourne, Victoria 3122, Australia

⁷ Arizona State University, School of Earth and Space Exploration, Tempe, AZ 85287, USA

⁸ University of Oxford, Department of Physics (Atmospheric, Oceanic and Planetary Physics), Oxford OX1 2JD, UK

⁹ University of Chicago, Department of Geophysical Sciences, Chicago, IL 60637, USA

¹⁰ Johns Hopkins University, Department of Physics and Astronomy, Baltimore, MD 21218, USA

¹¹ University of California Santa Cruz, Astronomy & Astrophysics Department, Santa Cruz, CA 95064, USA

¹² University of Birmingham, School of Physics & Astronomy, Birmingham, B15 2TT, UK

¹³ University of Amsterdam, Anton Pannekoek Institute for Astronomy, Amsterdam 1098 XH, The Netherlands

Received 2025 September 2; revised 2025 September 28; accepted 2025 October 8; published 2025 December 16

Abstract

A handful of enigmatic Jupiter-mass objects have been discovered orbiting pulsars. One such object, PSR J2322–2650b, uniquely resembles a hot-Jupiter exoplanet, due to its minimum density of 1.8 g cm^{-3} and its $\sim 1900 \text{ K}$ equilibrium temperature. We use JWST to observe PSR J2322–2650b’s emission spectrum across an entire orbit. In stark contrast to every known exoplanet orbiting a main-sequence star, we find an atmosphere rich in molecular carbon (C_3 , C_2) with strong westward winds. Our observations open up new exoplanetary chemical (ultrahigh C/O and C/N ratios of >100 and $>10,000$, respectively) and dynamical regimes (ultrafast rotation with external irradiation) to observational study. The extreme carbon enrichment poses a severe challenge to the current understanding of “black-widow” companions, which were expected to consist of a wider range of elements due to their origins as stripped stellar cores.

Unified Astronomy Thesaurus concepts: [Pulsar planets \(1304\)](#); [Carbon planets \(198\)](#); [Hot Jupiters \(753\)](#); [High energy astrophysics \(739\)](#); [Stellar atmospheres \(1584\)](#)

1. Introduction

In “black-widow” systems, a millisecond pulsar is orbited closely ($P_{\text{orb}} < 1 \text{ day}$) by a low-mass ($< 0.1 M_{\odot}$), often-degenerate companion. They are so named because in a prior low-mass X-ray binary phase, the pulsar was spun up by accreting mass from the companion, while at present the companions are being evaporated by the pulsar. There are ~ 50 known black-widow systems (K. I. I. Koljonen & M. Linares 2025), several of which have estimated companion masses below $10 M_J$. Of the handful with short (\lesssim few days) orbital periods, only one has a minimum density similar to that of gas giants orbiting main-sequence stars: PSR J2322–2650b (R. Spiewak et al. 2017; M. Shamohammadi et al. 2024), with a minimum mass of $0.8 M_J$ and a minimum density of 1.8 g cm^{-3} . The pulsar’s unusually low spindown luminosity of $2 \times 10^{32} \text{ erg s}^{-1}$ —arising from its unusually low magnetic field—would give the companion an equilibrium temperature of 1300 K if isotropically radiated and fully converted to heat.

However, the gamma rays that dominate a millisecond pulsar’s energy output and are responsible for heating the companion are typically beamed toward the equator (e.g., A. Philippov et al. 2020), making an accurate prediction of the companion temperature difficult.

The atmospheres of a few of PSR J2322–2650b’s denser and more massive black-widow cousins have been studied from the ground (e.g., P. Draghis et al. 2019; D. Kandel & R. W. Romani 2020; V. S. Dhillon et al. 2022). D. Kandel & R. W. Romani (2023) summarize 10 sets of previously published photometric phase curves of black-widow companions, with masses ranging from 12 to $67 M_J$ and densities ranging from 20 to 40 g cm^{-3} . All are extremely hot—nine have *nightside* temperatures above 2200 K , while J1311–3430 has no detectable *nightside* emission, but its *dayside* is $>10,000 \text{ K}$. These studies have measured the atmospheric circulation, Roche-lobe fill factors, and inclinations of these extreme objects. They have even provided tentative estimates of pulsar masses and put constraints on the neutron star equation of state. Spectra have been taken of several black-widow companions (e.g., M. H. van Kerkwijk et al. 2011; R. W. Romani et al. 2015, 2016, 2022; M. R. Kennedy et al. 2022; J. A. Simpson et al. 2025) and are generally stellar, indicating roughly solar composition. However, J1311–3430 and J1653–0158 both have 10 – $15 M_J$ companions with H-free surface spectra; these represent a black-widow subclass, the “Tidarrens,” which are likely descendants of ultracompact low-mass X-ray binaries with evolved companions

¹⁴ 51 Pegasi b Fellow.

¹⁵ OzGrav: The ARC Centre of Excellence for Gravitational Wave Discovery.

¹⁶ NHFP Sagan Fellow.






Figure 1. PSR J2322–2650 compared to the black widows from SpiderCat (K. I. I. Koljonen & M. Linares 2025) and to PSR J1719–1438 (M. Bailes et al. 2011), showing it is in a distinct area of parameter space. PSR J1719–1438, though similar in mass, has a much higher mean density of 21 g cm^{-3} ; its bulk composition is likely very different. The temperatures of all companions were very roughly estimated by assuming the spindown luminosity is isotropically radiated and all of the energy hitting the companion is converted to heat. The minimum density was derived from the orbital period via $\rho_{\min} = 3\pi/(0.46^3 GP^2)$, combining Paczyński’s approximation for the Roche lobe ($R_L = 0.46^3 q^{1/3} a$; B. Paczyński 1971) with Kepler’s third law.

(R. W. Romani et al. 2015). Y. Guo et al. (2022) showed with MESA simulations that these ultralight companions—including PSR J1719–1438, J2322–2650, and J1311–3430—could have formed by the Roche-lobe overflow and photoevaporation of helium stars.

PSR J2322–2650b is different from other ultralight pulsar companions, being the only pulsar companion with a mass, a density, and a temperature similar to those of hot Jupiters (Figure 1). The atmosphere of such an object has never been observed. It is, however, very easily observable with JWST: at infrared wavelengths, the pulsar is undetectable, giving us the unprecedented opportunity to obtain exquisite spectra of an externally irradiated planetary mass object. In this Letter, we present JWST spectra of this exotic gamma-ray-heated “exoplanet,” unveiling a bizarre atmosphere that raises more questions than it answers.

2. Observations

On 2024 November 8, we used NIRSpec/PRISM on JWST to observe PSR J2322–2650b’s $0.6\text{--}5.3 \mu\text{m}$ spectroscopic phase curve. These observations started shortly before inferior conjunction (the “nightside”) and continued until shortly after the following inferior conjunction, covering the whole 7.8 hr orbit. Two days later, we used NIRSpec/G235H ($1.7\text{--}3.1 \mu\text{m}$) to take a 2 hr long higher-resolution spectral sequence bracketing superior conjunction (the “dayside”). This data set spans an orbital phase range of 0.25, allowing us to measure the radial velocity of the planet as a function of time. Both the PRISM and G235H observations are taken with the $0.2''$ fixed slit. All the JWST data used in this Letter can be found at MAST: doi:10.17909/shv0-yq03. We develop a custom pipeline to reduce the data, significantly reducing the scatter compared to JWST standard pipeline processing (see Appendix B).

We show in Figure 2 the average PRISM spectrum on the dayside, quadrature, and nightside (defined as phases

0.5 ± 0.085 , 0.25 ± 0.05 or 0.75 ± 0.05 , and 0.00 ± 0.05 or 1.00 ± 0.05 , respectively). The spectra are shown both in μJy and in brightness temperature, calculated with a fiducial radius of $1.1R_J$ and an updated radio pulsar timing parallax distance of 630 pc (see Appendix A). While the nightside spectrum is featureless and consistent with a near-isothermal temperature profile or a thick gray dust/cloud deck, the dayside spectrum has clear absorption features. We can identify the molecules giving rise to these features by comparing the brightness temperature spectrum against $-\ln(\sigma(\lambda))$, where $\sigma(\lambda)$ is the absorption cross section of a candidate molecule. After a complete search of the DACE¹⁷ molecular database and an extensive search of the ExoMol (J. Tennyson & S. N. Yurchenko 2018) molecular database, we conclude that C_3 and C_2 are firmly detected. As shown in Figure 2, the cross section of $^{12}\text{C}_3$ (A. E. Lynas-Gray et al. 2024) rises suddenly redward of $3.014 \mu\text{m}$, matching the sudden flux drop seen in the data. (The isotopologues $^{12}\text{C}^{13}\text{C}^{12}\text{C}$ and $^{12}\text{C}^{12}\text{C}^{13}\text{C}$ have opacity cliffs at slightly different wavelengths and are inconsistent with the data, but their line lists are less reliable because they are not matched to experiments.) C_3 also has an opacity minimum at $3.9 \mu\text{m}$, matching the maximum in the emission spectrum, and an opacity maximum at $5.1 \mu\text{m}$, matching the emission plateau. The sawtooth pattern at $2.45\text{--}2.85 \mu\text{m}$ matches the series of band heads in the opacity of C_2 (S. N. Yurchenko et al. 2018).

C_3 and C_2 are necessary but not sufficient to explain the PRISM spectrum. First, the absorption feature at $3.0\text{--}3.6 \mu\text{m}$ is wider than the C_3 opacity peak, requiring another source of opacity around $3.4 \mu\text{m}$. Similarly, while a C_2 opacity peak around $1.4 \mu\text{m}$ could explain the sudden drop in flux at that wavelength, the absorption feature extends much redder than the opacity peak, suggesting that there is missing opacity around $1.6 \mu\text{m}$. We confirm that isotopologues of C_3 and C_2 cannot explain the missing opacity at either wavelength and tentatively identify both absorption features as arising from the C–H bond. The fundamental stretching mode of this bond is at $3.4 \mu\text{m}$, resulting in the ubiquitous presence of absorption at this wavelength among organic molecules (see Table 1.2 of M. Tammer 2004). The first overtone of the $3.4 \mu\text{m}$ stretch mode is at half the wavelength, or $1.7 \mu\text{m}$. Due to the similarity of absorption features across large classes of organic molecules, the exact molecule(s) giving rise to the C–H absorption is difficult to determine. We consider the identification of the features as C–H absorption to be tentative, because we detect no absorption feature around $2.3 \mu\text{m}$, whereas such a feature is common in hydrocarbons.

While the sawtooth feature in the PRISM spectrum strongly hints at the existence of C_2 , the higher-resolution G235H spectra conclusively demonstrate its existence. Using techniques from high-resolution cross-correlation spectroscopy (I. Snellen 2025), we cross-correlate the smoothed logarithm of C_2 opacities with each of the 151 integrations in the grating observation (see Appendix F). Plotting the cross-correlation function (CCF) with respect to time and blueshift reveals the radial velocity track of the planet. Given an assumed projected orbital velocity K_p and barycenter velocity v_{sys} , we can shift the CCFs into the planetary frame and sum them along the time axis to obtain a C_2 detection significance. Maximizing this significance over the $K_p\text{--}v_{\text{sys}}$ plane, as shown in the left

¹⁷ <https://dace.unige.ch/opacity/>




Figure 2. The observed emission flux and brightness temperature of PSR J2322–2650b. (a) Observed PRISM spectrum of the dayside, the average of the two quadratures, and the nightside at native resolution. The time-averaged G235H spectrum is also plotted, binned in wavelength by a factor of 16. The two shaded wavelength regions are those of the light curves in Figure 5. (b) PRISM spectra expressed in brightness temperatures by assuming $R/D = (1.1 R_J) / (630 \text{ pc})$. We plot HELIOS radiative-convective equilibrium forward models in gray. A linear function of the log of C_3 's absorption cross sections is plotted in green, showing that C_3 absorption explains the sudden dip at 3.014 μm , the recovery at 4 μm , and the slow decline toward the red edge.

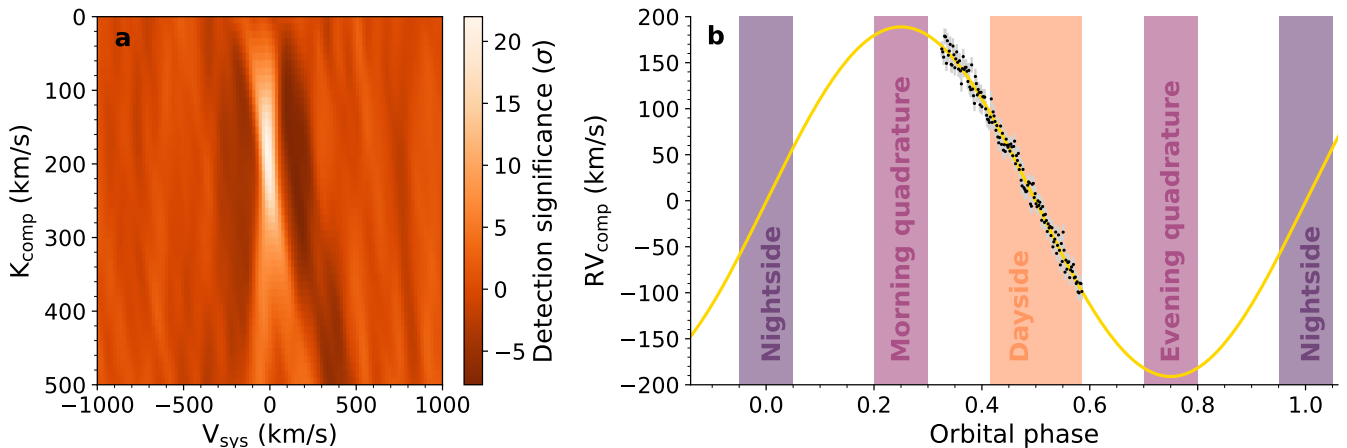


Figure 3. C_2 detection significance and radial velocities obtained from the JWST NIRSpect/G235H data. (a) The detection significance from the cross-correlation of C_2 in the G235H data, as a function of projected orbital velocity and barycentric radial velocity. C_2 is detected at 21σ at $V_{\text{sys}} \approx 0$ and $K_{\text{comp}} = 190 \text{ km s}^{-1}$. (b) In black, the companion radial velocity measured from each integration of the G235H observations using a data-derived template; in yellow, a sinusoidal fit to the data. The vertical offset is arbitrary. In the colored bars, we show the phases we define to be the nightside, quadrature, and dayside, for the purposes of calculating the average spectra in Figure 2.

panel of Figure 3, yields a 21σ detection of C_2 at a barycentric velocity close to 0 and a projected orbital velocity $K_p \sim 200 \text{ km s}^{-1}$.

3. Composition Constraints

In exoplanetary atmospheres, carbon-bearing molecules such as CO, CH_4 , and CO_2 are commonly detected, but

molecular carbon has never been seen. To quantify how depleted the atmosphere is in noncarbon elements, we obtain robust model-independent constraints on molecular abundance ratios by constructing a simulated G235H spectrum from only the cross sections of C_2 and a non- C_2 molecule, adjusting the line contrast until C_2 is detected at the same significance as in the real data, and computing the detection significance of the other molecule. As detailed in Appendix F, we perform this




Figure 4. The relative abundance of C_2 to CO in chemical equilibrium. (a) Abundance ratios for $C/He = 10^{-2}$ and $P = 10$ mbar, for an atmosphere with no other elements. (b) Ratios in an atmosphere with $H/He = 0.01$ as well as solar values of N/He , P/He , and S/He . The solid white curves represent the 3σ constraint from our cross-correlation test ($C_2/CO > 0.17$).

procedure with C_2/CO using the $2.30\text{--}2.47\ \mu\text{m}$ range, where both molecules have substantial opacity, and we find a 3σ lower limit on C_2/CO of 0.17. We repeat the procedure with CN , which has higher opacity than CO over a larger wavelength range, finding $C_2/CN > 32$ at 3.2σ . We additionally search for a wide variety of other molecules in the grating data via cross-correlation—including CS , CH , NS , OCS , H_2O , and CH_4 —but find no $>3\sigma$ detections.

Our inferred abundances suggest an atmosphere rich in carbon and relatively depleted in oxygen and nitrogen. Hydrogen is also heavily depleted, as otherwise the ample carbon atoms would bond with H to form hydrocarbons, resulting in minuscule abundances of molecular carbon, as has been predicted for hot-Jupiter exoplanets with high C/O ratios (N. Madhusudhan 2012). Some $C\text{--}H$ features are still possible, however, as trace amounts of hydrogen can be produced via spallation by the pulsar wind (B. M. S. Hansen 1996). On the other hand, this object is unlikely to be dominated by carbon in bulk, due to its equation of state. To test this, we calculate the mass–radius relationship of a pure carbon object using an interior structure model (M. C. Nixon & N. Madhusudhan 2021; M. C. Nixon et al. 2024) with the equation of state of G. I. Kerley & L. Chhabildas (2001) (see Appendix C) and find that, at zero temperature, the radius of a $1.5\text{--}2.5\ M_J$ object ($0.39\ R_J$) is much smaller than the near-Roche-lobe-filling radius our PRISM phase curve requires (see Section 5). A central temperature of $500,000\ \text{K}$ would be needed to inflate a carbon object to a Jupiter radius, but this is unrealistically hot, given (as we will later argue) its likely past as a white dwarf with efficient conductive heat transport. If, however, we assume an object that is mostly helium, with 1% C , then the radius increases to $0.92\ R_J$ for a $1.5\ M_J$ object with a photosphere temperature of $1500\ \text{K}$, using the helium equation of state from G. Chabrier & F. Debras (2021) and an additive volume law. We thus conclude that the planet’s bulk composition is likely helium-dominated.

3.1. Equilibrium Chemistry

We can test whether a helium-dominated atmosphere enriched in carbon can reproduce our observations by

performing equilibrium chemistry calculations with FastChem (J. W. Stock et al. 2018, 2022). Figure 4 plots the ratio of C_2/CO as a function of temperature and atomic C/O , either assuming $C/He = 0.01$ and no other elements or assuming $H/He = 0.01$ as well as solar N/He , P/He , and S/He . In either case, CO is highly favored over C_2 until $C/O \gg 1$ and far more favored at low temperatures. Even at $2300\ \text{K}$, the hottest point in our ICARUS model (see below), our constraint of $C_2/CO > 0.17$ requires $C/O > 100$. At the average dayside brightness temperature of $2000\ \text{K}$, $C_2/CO > 0.17$ would require $C/O > 2000$ (the pure-He-and-C scenario) or $C/O > 20,000$ (the many-elements scenario). We generate multidimensional abundance grids to explore a wide range of metallicities, C/O and C/N ratios, temperatures, and pressures, finding that our observational constraints on C_2/CO and C_2/CN imply $C/O > 100$ and $C/N > 10,000$ over a wide range of conditions (see Appendix H).

In our high- C/O calculations with only C , O , and He , C_3 is the most abundant molecule, consistent with our observations; at $C/O \sim 1$, CO is the most abundant molecule and C_3 has negligible (parts per billion or less) abundances. In our calculations including hydrogen and other elements, C_3H and C_2H are comparably abundant to C_3 and C_2 over a wide range of temperatures and C/O ratios. The $C\text{--}H$ bonds on these molecules may explain the opacity we see at 3.4 and $1.6\ \mu\text{m}$, but since there are no published line lists for these molecules, we are unable to include them in our modeling. Colder carbon-rich compositions ($T < 1000\ \text{K}$, $C/O \gg 1$) are predicted to have abundant quantities of larger carbon molecules, such as C_4 and C_5 , which also have no published line lists. To compare these results to data, we generate forward models using HELIOS (M. Malik et al. 2017, 2019; see Appendix D), assuming 1D radiative-convective equilibrium in a bottom-heated helium-dominated atmosphere with an extended graphite dust/cloud layer and \sim part per thousand abundances of C_3 , C_2 , and C_2H_4 (a stand-in for a generic molecule with $C\text{--}H$ bonds). These are shown in Figure 2, which demonstrates that this simple prescription reproduces the main spectral features at all phases, with the featureless nightside explained by invoking more dust absorption.

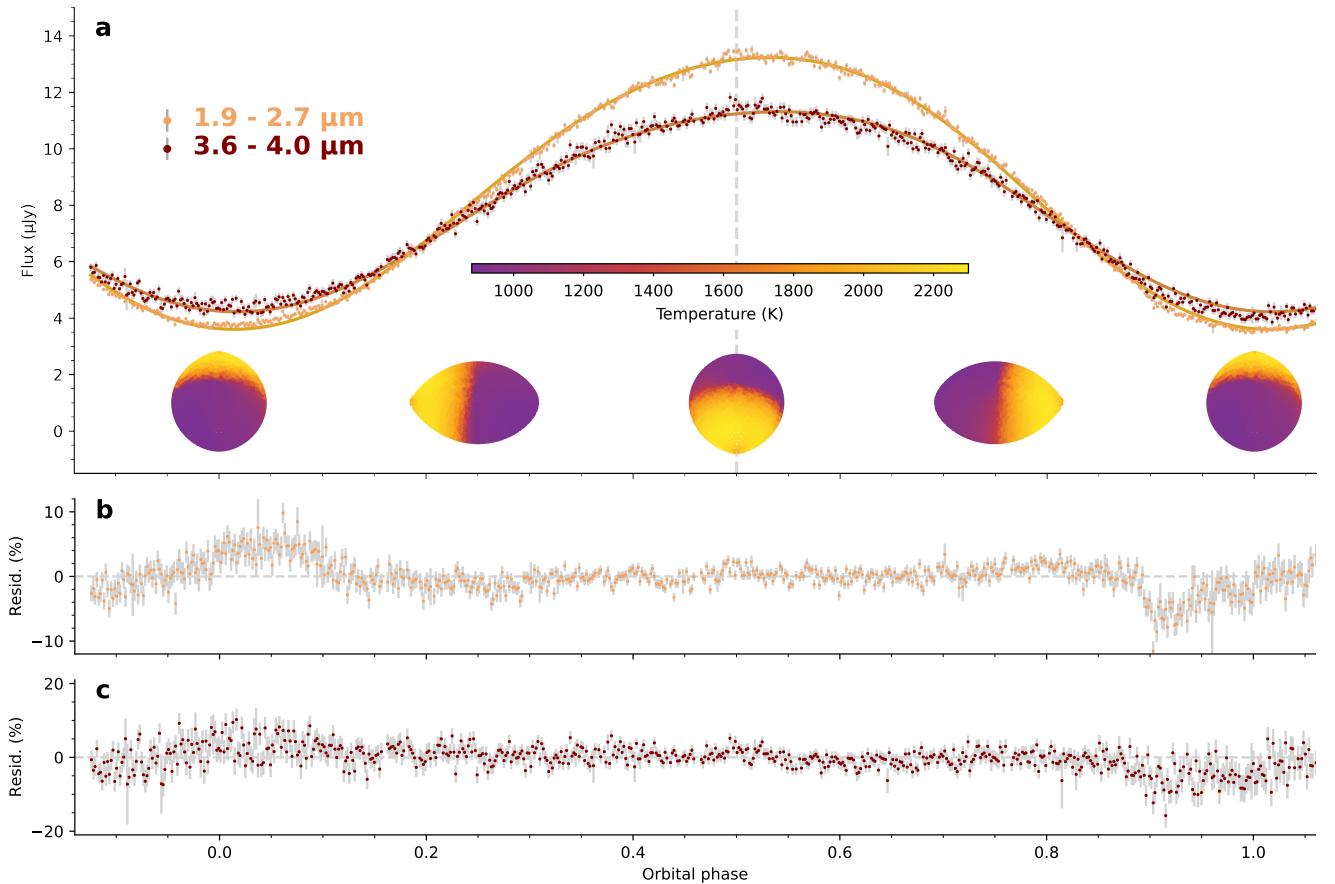


Figure 5. The companion light curve in two continuum bands interpreted with a pulsar direct heating model. (a) The models are shown by the curves and the data are shown by the points. The model does not capture the strong orbit-to-orbit variations at minimum ($\phi \sim 0$ and $\phi \sim 1$) but provides a good description of the flux from the heated face and constrains the binary parameters (see Section 5). Schematics of the Earth view of the heated companion’s surface temperature are shown below the light curve. (b) and (c) The residuals of the models compared to their corresponding light curves.

4. Radial Velocities

Refining the value of the projected orbital velocity K_p allows us to pin down the inclination of the companion orbit and the masses of the companion and pulsar, as well as assess the validity of our fiducial companion radius. We use the G235H spectra to monitor the radial velocity variations (D. K. Sing et al. 2024). We generate a template spectrum by shifting every spectrum into the planetary frame using our preliminary K_p value ($\sim 200 \text{ km s}^{-1}$) and taking the timewise average. We then fit this template to each individual spectrum to find the radial velocity shift as a function of time, as shown in Figure 3 (right). Finally, we fit the sinusoid $v(\phi) = v_0 + K_p \sin(2\pi\phi)$ to the velocities. We obtain $K_p = 190 \pm 2 \text{ km s}^{-1}$. This value represents the center-of-light orbital velocity K_{CoL} , which is lower than the planet’s center-of-mass orbital velocity K_{CoM} , because the bright, irradiated side of the planet is closer to the pulsar. The ratio $K_{\text{cor}} = K_{\text{CoM}}/K_{\text{CoL}}$ depends on the heating pattern, temperature dependence of the lines dominating the radial velocity, Roche-lobe fill factor, and inclination. For our system with its mass ratio $q \approx 1000$, the maximum possible value of K_{cor} is $1/(1 - (3q)^{-1/3}) = 1.07$, which reduces to 1.04 for a more realistic $T_0 \cos^{1/4}(\theta)$ temperature distribution (M. H. van Kerkwijk et al. 2011).

For a binary system in a circular orbit where the primary is much more massive than the secondary, $K_{\text{CoM}} = K_{\text{CoL}} K_{\text{cor}} = \left(\frac{2\pi GM_*}{P}\right)^{1/3} \sin i$. Assuming the pulsar mass M_p can range from

the canonical $1.4M_\odot$ to an unusually massive $2.4M_\odot$ (see Figure 7 of J. M. Lattimer 2012), we infer an inclination range of $34.7\text{--}28.4^\circ$ (for $K_{\text{cor}} = 1.04$). This low i , combined with the fact that pulsar gamma rays are beamed toward the spin equator, explains a number of system features: the lack of a Fermi gamma-ray detection (S. Abdollahi et al. 2022), the high companion temperature compared to the pulsar spindown power, and the lack of radio eclipses from a companion outflow. This range of inclinations implies a planet mass of $1.4\text{--}2.4M_J$, a Roche-lobe radius $R_L = 0.462a(m_p/m_*)^{1/3} = 0.462\left(\frac{GM_p P^2}{4\pi^2}\right)^{1/3}$ of $0.99\text{--}1.18R_J$, and density of 1.8 g cm^{-3} . All plausible pulsar masses therefore imply a roughly Jupiter-sized and Jupiter-mass companion on a $\sim 30^\circ$ inclination orbit.

5. PRISM Phase Curve

Next, we model the PRISM phase curve. The C_3 absorption depth varies dramatically with orbital phase, complicating modeling of the heated companion surface. However, at wavelengths of relatively low opacity, a basic thermal model can be used to constrain the binary properties. Figure 5 shows the light curves in two wavelength ranges between the strongest absorption bands. The flux varies by \sim three times from day to night, a modest amplitude compared to other black-widow companions, which can show modulations of 30–100 times (P. Draghis et al. 2019). The light curves exhibit nightside variability—the second minimum is 10% lower than the first,

possibly due to flares (D. Kandel & R. W. Romani 2020). We model the light curves in the two opacity windows by assuming Planckian emergent spectra and direct pulsar heating, using a version of the ICARUS heated binary light-curve code (R. P. Breton et al. 2012; see Section 5). All fits require a near-filled Roche lobe. Fitting the full light curves, we find $i \sim 31^\circ$, with photosphere temperatures ranging from 900 K at the coldest points on the nightside to 2300 K at the hottest points on the dayside. We infer a distance consistent with the radio parallax measurement and a westward wind. Given the closeness of the inferred inclination to the values we obtained previously by assuming reasonable pulsar masses, it is not surprising that we now infer a pulsar mass ($\sim 2.0M_\odot$ with $0.1M_\odot$ statistical error and $0.5M_\odot$ systematic error) consistent with astrophysical expectations (see Appendix 5). The inferred isotropic heating luminosity of $9 \times 10^{32} \text{ erg s}^{-1}$ is $5/I_{45}$ times the pulsar’s spindown luminosity, where $10^{45} I_{45} \text{ g cm}^2$ is the pulsar’s moment of inertia. The apparent over-unity heating efficiency is an artifact of the pulsar’s gamma-ray emission being beamed along the spin equator (and, hence, toward the companion), possibly in combination with a large moment of inertia arising from a high pulsar mass.

The flux maximum in the phase curve occurs 12° after phase 0.5, indicating a westward offset that is suggestive of strong winds blowing opposite to the planet’s rotation direction. A westward thermal phase offset is seen in some black-widow companions (D. Kandel & R. W. Romani 2020) and is rare among hot Jupiters orbiting main-sequence stars (T. J. Bell et al. 2021), having only been robustly observed for CoRoT-2b (L. Dang et al. 2018). All of these objects have rotation periods equal to their orbital periods, due to tidal locking, but PSR J2322–2650b’s exceptionally short rotation period of 0.32 days places it in a different dynamical regime from almost all hot Jupiters. E. K. H. Lee et al. (2020) and X. Tan & A. P. Showman (2020) demonstrated with 3D general circulation models (GCMs) that, at very short rotation periods of $\lesssim 10$ hr, the equatorial eastward jet ubiquitous on slower-rotating hot Jupiters narrows, while off-equatorial westward zonal winds become increasingly distinctive, leading to a westward phase offset. Because we do not expect the dynamics on PSR J2322–2650b to be identical to X. Tan & A. P. Showman (2020), due to the very different form of irradiation, we run a planet-specific atmospheric dynamics model with MITgcm (see Appendix E). Our temperature maps are qualitatively similar to theirs; in particular, we also obtain a westward phase offset. Our data are therefore observational evidence of the dynamical regime that X. Tan & A. P. Showman (2020) predicted—a regime rarely probed by hot Jupiters orbiting main-sequence stars.

Better line lists and improvements in the modeling of planetary atmospheres will allow more accurate computation of inclination and pulsar mass through modeling of the spectroscopic phase curve. This could be particularly interesting for PSR J2322–2650, because the pulsar’s low intrinsic spindown \dot{P} implies an exceptionally low neutron star dipole field $2.5 \times 10^7 I_{45} \text{ G}$. This is relevant, since the reduced magnetic field of millisecond pulsars is often attributed to accretion during the spinup process (G. S. Bisnovaty-Kogan & B. V. Komberg 1974; R. E. Taam & E. P. J. van den Heuvel 1986; R. W. Romani 1990; D. Mukherjee 2017); with such a low field, the pulsar may have accreted a large amount of matter, achieving a larger moment of inertia and a high

mass. Thus, precise measurement of the masses in this system could help constrain M_{TOV} , the Tolman–Oppenheimer–Volkoff mass, which sets the maximum mass for a slowly rotating neutron star and constrains the dense matter equation of state.

6. Conclusion

Our findings pose a challenge to the current understanding of black-widow formation, in which a pulsar strips the outer layers of its stellar companion by a combination of Roche-lobe overflow and photoevaporation (O. G. Benvenuto et al. 2012; Y. Guo et al. 2022). This mechanism can produce a helium-dominated Jupiter-mass object if stripping occurs before core helium burning, during the main sequence or red giant branch. In this picture, these black widows are the descendants of ultracompact X-ray binaries with He-star donors, and the core C/N/O ratios can be sensitive to the evolutionary state of the donor at the start of Roche-lobe overflow (see Figure 4 of G. Nelemans et al. 2010). How this process produces a C/O ratio greater than 100 or a C/N ratio greater than 10,000 is difficult to explain. Although G. Nelemans et al. (2010) predict extreme N depletion in one scenario (see their Figure 4, bottom middle), O becomes significantly enhanced relative to C in this scenario. In the other two scenarios (corresponding to different initial helium-star masses), C is enhanced relative to O by a factor of ~ 15 , but it is only enhanced relative to N by a factor of \sim several. Although extreme photospheric ratios are possible via gravitational settling for the much-higher-gravity white dwarfs (the immediate progenitors of black-widow companions in the standard formation scenario), gravitational settling is not expected to differentiate elements once the companion becomes a planetary mass object with a molecular atmosphere.

Other mechanisms of carbon enrichment pose their own challenges. For example, the helium-dominated, hydrogen-poor, and carbon-enriched composition of PSR J2322–2650b is reminiscent of R Coronae Borealis (RCB) stars and dustless hydrogen-deficient carbon (dLHdC) stars, which are believed to result from mergers of He and CO white dwarfs (R. F. Webbink 1984; H. Saio & C. S. Jeffery 2002). However, the rarity of RCB/dLHdC stars and the necessity of invoking three stars makes them unlikely progenitors, and the measured C/O and C/N ratios of these stars (12–81 and 4–130, respectively; A. Mehla et al. 2025) are still lower than the very high minimum values we obtain. The triple-alpha process, which fuses helium into carbon, creates asymptotic giant branch “carbon stars” with C/O ratios of up to \sim several (C. Abia et al. 2008) but not the $C/O > 100$ inferred from our data. Carbon stars have a dusty outflow rich in carbon grains (M. Di Criscienzo et al. 2013), providing a further source of carbon enrichment, but further modeling is required to explain how this carbon ultimately ends up in a Jupiter-mass object. More spectroscopic observations of ultralow-mass black-widow companions (“Tidarrens”) are needed to determine whether PSR J2322–2650b’s composition is unusual or representative of the class. We particularly encourage observations of PSR J1719–1438, which has a similar mass but a far higher minimum density of 21 g cm^{-3} , leading M. Bailes et al. (2011) to suggest it is an ultralow-mass carbon white dwarf.

Data Access

A Jupyter notebook is available to reproduce the figures in this work in Zenodo (M. Zhang 2025) at doi:[10.5281/zenodo.17400581](https://doi.org/10.5281/zenodo.17400581). The repository also includes the spectra plotted in Figure 2(a) in ASCII format.

Acknowledgments

The authors thank Gerald Kerley for providing the carbon equation of state used for the interior structure models. T.D.K. acknowledges Jayne Birkby, Luke Parker, and Stephen Smartt for insightful discussions. M.Z. thanks Qiao Xue, Keeyoon Sung, Sergey Yurchenko, John Tonry, and Jeremy Goodman for insightful feedback. M.B. thanks Jaikhomba Singha and Andrea Possenti for comments on the pulsar timing section and Sarah Buchner for observing. P.G. acknowledges Shazrene Mohamed and Mikako Matsuura for enlightening discussions.

These observations are associated with JWST GO #5263 and the work is funded by the associated grant. The MeerKAT telescope is operated by the South African Radio Astronomy Observatory, which is a facility of the National Research Foundation, an agency of the Department of Science and Innovation. M.Z. acknowledges support from the 51 Pegasi b Fellowship funded by the Heising–Simons Foundation. Part of this work was undertaken for the Australian Research Council Centre of Excellence for Gravitational Wave Discovery (project number CE230100016). Support for this work was provided by NASA through the NASA Hubble Fellowship grant HST-HF2-51559.001-A awarded by the Space Telescope Science Institute, which is operated by the Association of Universities for Research in Astronomy, Inc., for NASA, under contract NAS5-26555.

Author Contributions

M.Z. proposed the observations, led the project, identified C₃ and C₂, and led the writing of the Letter. M.B. fit the continuum light curves with ICARUS to obtain temperature maps and constrain binary parameters. T.D.B. performed the data analysis using primarily custom software. R.W.R. gave advice on the pulsar physics, proposed a plausible formation mechanism, and helped with light-curve fitting. P.G. searched opacity databases to identify spectral features and ran atmospheric forward models to fit the spectrum. H.B. ran, postprocessed, and interpreted GCMs with a bespoke heating scheme. M.B. is the PI of the MeerTime Large Survey Project, oversees the pulsar portal/pipeline, and created the standard profile, arrival times, and initial model fits. M.N. performed the interior structure calculations. J.L.B. helped conceive the project and gave advice through all stages. T.D.K. helped with GCM development and interpretation and writing the manuscript. B.P.C. suggested the identification of C–H spectral features. G.F. performed an independent data reduction. R.L. helped with the figures and with data interpretation. D.J.R. conducted pulsar timing analysis to determine system parameters. E.C. ran and maintained the pipeline producing the MeerKAT pulsar timing data. R.M.S. is the PI of the MeerKAT PTA timing project and conducted pulsar timing analysis to determine system parameters. J.J.F. advised on the interior structure and atmosphere modeling. A.A.A.P. performed a blackbody retrieval and gave comments on the manuscript. M.C.M. edited the Letter and contributed to understanding the pulsar heating. J-M.D. contributed to the proposal.

Facilities: JWST (NIRSpec), MeerKAT.

Software: astropy (Astropy Collaboration et al. 2013, 2018, 2022), JWST Calibration Pipeline (H. Bushouse et al. 2025), HELIOS (M. Malik et al. 2017, 2019), HELIOS-K (S. L. Grimm et al. 2021), FastChem (J. W. Stock et al. 2022).

Appendix A Analysis of MeerKAT Radio Data

The distance to the pulsar is important for the interpretation of this system in at least two ways: it determines the luminosity of the companion, and it allows us to derive the intrinsic spindown rate \dot{P}_1 , due to the Shklovskii correction for pseudo-acceleration (I. S. Shklovskii 1970). The pulsar’s spindown luminosity and inferred magnetic field strength are both related to \dot{P}_1 . The PSR J2322–2650 discovery paper (R. Spiewak et al. 2017) reported a low timing parallax distance of 230_{-50}^{+90} pc, much smaller than their dispersion measure (DM) model distance of 760 pc. M. Shamohammadi et al. (2024), using the more accurate MeerKAT radio telescope data set, obtained a timing parallax of 1.3 ± 0.2 mas, from which was derived a distance of 770_{-100}^{+130} pc. Here, we improve the distance estimate by extending the MeerKAT data set.

The 3.46 ms pulsar J2322–2650 is observed regularly as part of the MeerKAT Pulsar Timing Array (MPTA; M. T. Miles et al. 2023), using the PTUSE pulsar instrument (M. Bailes et al. 2020). Observations span from 2019 July to 2025 April, and we obtain arrival times at 136 different epochs in 16 coherently dispersed frequency subbands, spanning 775.0 MHz centered at 1283.582 MHz. A standard pulse profile is created by aligning all data with the optimal DM and with the phase-connected timing solution, removing the dispersion and averaging into 16 subbands, before wavelet smoothing with the PSRCHIVE (W. van Straten et al. 2012) application PSRSMOOTH. Topocentric arrival times are measured with a standard Fourier-domain pulse profile cross-correlation technique, with uncertainties estimated with a Monte Carlo simulation. Arrival times with a signal-to-noise ratio of cross-correlation below 10 are removed, and we perform a least-squares fit using TEMPO2 (G. B. Hobbs et al. 2006) for the pulsar spin, binary, and astrometric parameters on the 2625 remaining arrival times. The reduced χ^2 is 1.0, and the weighted rms and median residuals are $2.2 \mu\text{s}$ and $3.4 \mu\text{s}$, respectively, indicating a good model fit.

We infer posterior probability distributions for the timing model parameters together with noise parameters using the




Figure 6. Posterior probability distributions for key pulsar parameters from the timing of PSR J2322–2650 with the MeerKAT radio telescope.

Table 1
Measured Timing Parameters of PSR J2322–2650

Parameter	Value
Pulsar Parameters	
R.A., α (J2000)	23:22:34.638364(4)
decl., δ (J2000)	−26:50:58.39619(8)
Proper motion in R.A., $\mu_\alpha \cos \delta$	−2.26(3) mas yr ^{−1}
Proper motion in decl., μ_δ	−8.31(5) mas yr ^{−1}
Spin period (P)	3.463099179266747(4) ms
Period derivative (\dot{P})	$5.8586(16) \times 10^{-22}$ s s ^{−1}
Epoch (MJD)	59693.0
DM	6.14449(4) pc cm ^{−3}
Parallax	1.59(17) mas
Semimajor axis ^a	0.00278444 (s)
Orbital period	0.3229640004(7) days
Eccentricity, e	<0.00012 (95%)
MJD ascending node	59692.9699957(14)
Derived Parameters	
Distance, D	630_{-60}^{+75} pc
Transverse velocity	26 ± 3 km s ^{−1}
Intrinsic \dot{P}	$1.8(4) \times 10^{-22}$

Note.

^a Projected semimajor axis of the pulsar (not the companion), in light-seconds.

Bayesian pulsar timing package inference ENTERPRISE (J. A. Ellis et al. 2020). We use a noise model that is standard in pulsar timing array analyses (including for the MPTA; M. T. Miles et al. 2025), which describes white noise, DM variations, and achromatic red noise in timing residuals. However, no significant red-noise processes were detected. Corner plots (D. Foreman-Mackey 2016) of the 1D and 2D marginal posterior probability distributions for the main timing model parameters of interest are shown in Figure 6. The measured timing model parameters are provided in Table 1, including the pulsar parallax, proper motion, period, and period derivative. We do not detect orbital eccentricity and place a 95% upper limit of $e < 0.00012$. The pulsar distance derived from the parallax is $D = 630_{-60}^{+75}$ pc and the transverse velocity with respect to the Sun is just 26 km s^{-1} . Correcting for the peculiar motion of the Sun means that the velocity of the pulsar is small in the local standard of rest and indistinguishable from random motions in the Galaxy. The pulsar is experiencing Shklovskii pseudo-acceleration due to its transverse velocity and distance, and this implies that the intrinsic period derivative is $\dot{P}_i = \dot{P}_{\text{obs}} - PD\mu^2/c = 1.8(4) \times 10^{-22}$, where μ is the pulsar proper motion and c is the speed of light. Such a low period derivative makes PSR J2322–2650 one of the pulsars with the weakest magnetic field strengths known ($B^2 \propto P\dot{P}$), which may have helped the planet survive irradiation from the pulsar.

Appendix B Analysis of JWST Data

We extract 1D spectra from the uncalibrated JWST data using a mixture of the JWST pipeline and our own software and methods. The first stage of the pipeline consists of fitting a count rate to each pixel’s measured sequence of accumulated counts. While we follow the general approach of the JWST pipeline, we use our own software for every step of this stage. Some of the changes in our approach have since been implemented in the official pipeline.

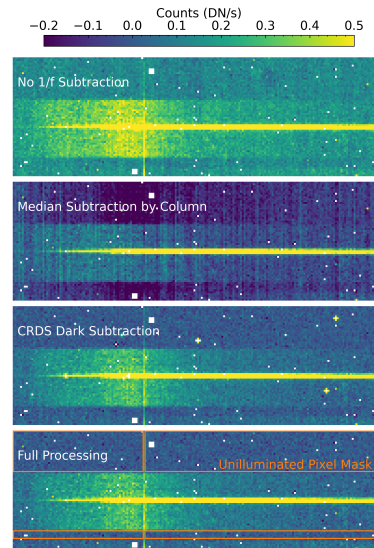


Figure 7. Comparison of different data reduction processes for the PRISM data, averaged over 10 integrations. The region shown encompasses just under half of the full PRISM spectrum.

Correlated $1/f$ noise (S. H. Moseley et al. 2010) is typically corrected by the JWST pipeline using non-light-sensitive reference pixels. Unfortunately, our data are in subarray mode, where most of these reference pixels are skipped rather than being read out. Our treatment of $1/f$ noise differs somewhat from that of the JWST pipeline. We describe that approach first, as we use it for all of our data processing, including the construction of dark reference files.

The $1/f$ noise is correlated in time. The pixels in the NIRSPEC subarrays are read out in columns, with each pixel in a column being read out sequentially, before moving on to the next column. As a result, temporally correlated noise appears as a spatial correlation, with the value of the noise being similar among all pixels in a column. There are several approaches that can be used to mitigate this $1/f$ noise. One is to compute and subtract the median value column by column from each read, while another is to determine a set of pixels that receive no signal (or nearly no signal) and to use these to construct a noise model. Figure 7 shows both of these approaches, together with the result of doing no correction, in the top three panels. In each case, we show the count rate averaged over the same 10 integrations with the same linear color scale.

The top panel of Figure 7 shows the result with no $1/f$ correction. The background regions have a visibly nonzero count rate, due to the low-frequency component of the $1/f$ that is shared by all reads; it would have a different value if a different set of integrations were averaged together. The second panel shows the result of a median subtraction column by column. It shows large single-column deviations from systematically different pixels along with biases from the fact that science data are not excluded.

Our $1/f$ noise correction, shown in the third and fourth panels of Figure 7, is based on the NSClean algorithm (B. J. Rauscher 2024) as implemented in the JWST pipeline (H. Bushouse et al. 2025). Our changes to the algorithm consist of tuning the frequencies to be corrected and designing a custom mask for the pixels to be used. For PRISM data, we use four rows (numbers 7–10 of 64, beginning from 1)




Figure 8. Empirical profile of the PRISM data used to extract spectra using the optimal extraction algorithm. We use the 3 central pixels along the trace, which itself is nearly parallel to a detector row, to extract the spectrum. The pixel values correspond to the JWST pipeline’s `xldints` files; the corresponding wavelengths are given on the top axis. Wavelengths below $0.8 \mu\text{m}$ are not used in the analysis.

between shutters, in addition to the top 22 rows. These rows (with the exception of one bad column) are shown in the bottom panel of Figure 7. The frequencies to be corrected are specified by four numbers. Frequencies below the lowest number or above the highest number (in Hz) are not corrected; frequencies between the second and third numbers are fully fit. Frequencies between the first and second numbers, and between the third and fourth numbers, are apodized—i.e., fit with a weight that smoothly varies between 0 and 1. The default values for these four frequencies are 1061, 1211, 49,943, and 49,957 Hz. The latter number is just below 50,000 Hz; odd–even noise appears at a frequency of half the pixel rate of 100 kHz and is corrected. We adopt values of 200, 600, 49,943, and 49,957 for our correction frequencies—i.e., we fit for lower-frequency noise than the pipeline defaults. We are able to do this because we have more pixels available for fitting in our custom mask than in the default mask.

In addition to the $1/f$ correction, we also revisit the JWST pipeline step of the dark subtraction. The dark signal and bad pixels available at the JWST Calibration Reference Data System (CRDS)¹⁸ have evolved since those data were taken and do not provide a perfect match to our data. Further, the dark signal for most pixels is extremely small, so that it is measured by the dark reference files only with low signal-to-noise ratios. As a result, we do not perform a dark subtraction for most pixels. Any dark that we would construct or use from CRDS would have a realization of read noise in it, even if this read noise is averaged over many integrations. Our mean spectra are themselves averaged over nearly 700 integrations, so the dark could easily make a significant contribution to the total read noise. To avoid adding unnecessary read noise, we use darks only to correct the neighbors of hot pixels.

We define hot pixels for our purposes as those that accumulate more than 250 counts in 30 reads. We do not aim to recover the value of the hot pixel but simply to remove the impact it makes on its neighbors due to interpixel capacitance. To get this right pixel by pixel for our PRISM data, we construct a dark image for each read, from 151 integrations of grating data with the same subarray configuration. These data do have a faint visible trace. We fit this trace with a quadratic in pixel coordinates. Near the trace, the grating-derived dark would be an inappropriate choice, as it is not truly dark. We therefore substitute the value of the CRDS dark for pixels within 3 pixels of the center of the grating trace.

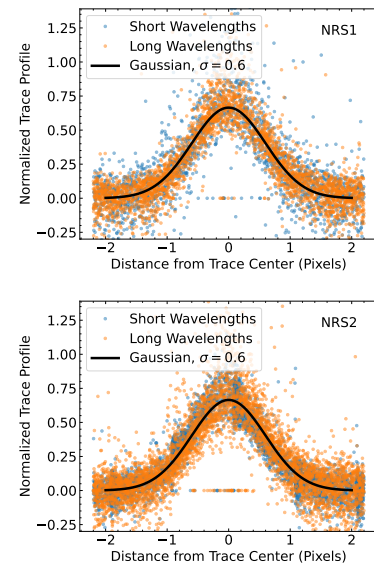


Figure 9. Adopted empirical profiles of the G235 grating traces for detectors NRS1 (top) and NRS2 (bottom) shown with points for the measured pixel values after normalizing in the cross-dispersion direction and fitting the trace position. The different colored points show the shorter- and longer-wavelength ranges of each detector.

Finally, we use our hybrid dark to correct hot pixels and their immediate neighbors, before masking the hot pixels themselves. For an isolated hot pixel, we scale the dark to the science frame and subtract it in a 3×3 patch around the hot pixel. The hot pixel itself now has zero flux; we mask it for all further analysis. The immediate neighbors of the hot pixel now have had the effects of interpixel capacitance removed. If a hot pixel is not isolated, we need to be careful not to doubly remove the interpixel capacitance. We use the same algorithm described above, but for pixels that border more than one hot pixel, we split their values across multiple 3×3 patches. The bottom panel of Figure 7 shows the impact of the hot-pixel processing, with neighbors of hot pixels appearing to be consistent with the large-scale background.

The steps above create a calibrated set of individual reads, but we still must fit for a count rate for each pixel in each integration. For this, we use the likelihood-based ramp-fitting and jump-detection algorithms described in T. D. Brandt (2024a) and T. D. Brandt (2024b). These were integrated into the JWST pipeline after these data were reduced. So-called snowballs (M. Regan 2024) are clusters of pixels showing a jump in counts after an energetic cosmic-ray hit. Some of these jumps are at levels that are difficult to detect in an individual pixel. To remove snowballs, we search for 7×7 pixel clusters where at least 80% of the pixels show evidence of a jump. We then compute the mean radius of this cluster of jumps, expand it by a factor of 2.5, and apply a mask to these reads for these pixels. Finally, to mitigate persistence from strong cosmic-ray hits, we exclude the remainder of the ramp for all pixels that had jumps registered on at least three consecutive reads.

Our suppression of correlated noise is good enough that the uncertainties in the ramps are now overestimated. We quantify this by the distribution of the ratio of the ramp slope fit divided by its estimated uncertainty. Because most pixels are either unilluminated or nearly so, this distribution should closely approximate a unit Gaussian. Instead, we find that the distribution closely approximates a Gaussian with a standard

¹⁸ <https://jwst-crds.stsci.edu>

deviation of 0.8. We therefore reduce the estimated uncertainties on the ramp slopes by 20%.

Our processing of the grating data matches our processing of the PRISM data as described above and shown in Figure 7. For the background pixels for a $1/f$ correction, we empirically determine the location of the trace by fitting a quadratic in pixel coordinates and exclude pixels within 4 pixels of the trace.

The next step of our data reduction is spectral extraction. We use the JWST pipeline for the initial step of flux calibration, continuing with our own routines subsequent to the generation of the `calints` files. We proceed with optimal extraction using an empirical trace profile. For the PRISM data, the trace is nearly horizontal, lining up with one of the detector rows. We limit the consideration of the trace to 3 pixels: the central pixel and its two nearest neighbors. We first smooth the 2D spectrum, using a running median in the dispersion direction; we then normalize each column in the cross-dispersion direction. We next fit a quadratic to the (normalized) pixel value as a function of position in the dispersion direction. Finally, we renormalize this empirical trace profile at each position in the dispersion direction.

Figure 8 shows the empirical trace profiles that we obtain from the JWST pipeline’s `calints` files using our custom `rateints` files as input. These consist of scaled 2D calibrated cutouts from the ramp slope files that we obtain as described earlier in this section. We use the trace profiles in Figure 8 to perform an optimal extraction (K. Horne 1986) but without weighting by the estimated uncertainties in each pixel. The estimated uncertainties are correlated with the measurement errors themselves (fewer recorded photons lead to a lower estimated photon noise), so weighting the profile by the estimated uncertainties would introduce a bias. The code that we use for spectral extraction, although not the code for constructing the trace profile, has also been added to the JWST pipeline since this reduction was done. We use the JWST pipeline’s wavelength calibration.

Our approach for grating spectral extraction differs for two reasons: the signal-to-noise ratio is lower, and the trace shows significant curvature relative to the detector axes. We adopt a multistage process. We begin by averaging all of our grating exposures. We then fit the trace profile using a quadratic in pixel coordinates. Next, we normalize the average grating exposure column by column. Finally, we plot each pixel value as a function of its distance from the trace. A Gaussian with a standard deviation of 0.6 pixels provides a visually good fit. Figure 9 shows this empirical trace profile together with the measured values as a function of distance from the trace center. The trace profile appears to be consistent between the NRS1 and NRS2 detectors and does not appear to vary significantly with wavelength: the short-wavelength and long-wavelength halves of the data are consistent when plotted against one another. A higher-fidelity reconstruction of the trace profile for the grating data could modestly improve our results. Finally, we adopt the JWST pipeline’s wavelength solution.

Appendix C Internal Structure Modeling

We generate interior structure models (M. C. Nixon & N. Madhusudhan 2021; M. C. Nixon et al. 2024) of PSR J2322–2650b to explore its possible bulk composition.

We consider three illustrative compositions for the planet: a pure carbon interior, a pure helium interior, and an interior of 99% helium and 1% carbon by mass. We use temperature-dependent equations of state for helium (G. Chabrier & F. Debras 2021) and carbon (G. I. Kerley & L. Chhabildas 2001). For the mixed carbon and helium interior, we assume an additive volume law to obtain the mixed equation of state. We assume an isothermal–adiabatic temperature with a photospheric pressure of 1 mbar and a radiative–convective boundary of 10 bar. The isothermal region of the atmosphere is fixed to 1500 K, with an adiabatic temperature profile used at deeper pressures.

Within the mass range of $1.5\text{--}2.5M_J$, we find that a pure carbon object would have a radius of $0.393\text{--}0.394 R_J$. Adjusting the temperature profile so that the radius becomes $1 R_J$ requires a photospheric temperature of 5300 K, which yields a central temperature of 4.6×10^5 K. This is significantly hotter than would be expected for PSR J2322–2650b. The pure helium and mixed helium + carbon models yield radii much closer to $1 R_J$ using our nominal temperature profile. The pure helium models range from 0.837 to $0.947 R_J$, and the mixed helium + carbon models range from 0.821 to $0.925 R_J$, within the mass range of $1.5\text{--}2.5M_J$.

Given the apparent depletion of hydrogen from the planet’s atmosphere, and the fact that model planets with significant amounts of elements heavier than helium result in radii $\ll 1R_J$, it seems likely from this initial analysis that the planet consists primarily of helium. A more complete analysis, using a wider range of compositions and accounting for the partly differentiated interior as well as the nonspherical shape, would be required to confirm this and refine constraints on the bulk composition. This detailed analysis is left for a future study.

Appendix D HELIOS Models

HELIOS is an open-source 1D radiative transfer code developed for planetary atmospheres (M. Malik et al. 2017). Given planetary parameters and an atmospheric composition, it finds the radiative–convective equilibrium and outputs the resulting emission spectrum as well as the temperature–pressure profile. HELIOS self-consistently supports external irradiation by main-sequence stars, but we model the pulsar planet as an internally heated brown dwarf, because most of the pulsar’s radiation is in \sim GeV gamma rays, which penetrate well below the infrared photosphere. In a helium atmosphere, the electromagnetic shower induced by the gamma-ray collision deposits energy at a mass column density of $\Sigma \sim 2000\text{g cm}^{-2}$ (Equation (7) of D. Kandel & R. W. Romani 2020), corresponding to a pressure of $P = \Sigma g = 10$ bar for the subpulsar point. (This calculation neglects tidal forces from the pulsar, which reduce the effective g and therefore the characteristic P .) We therefore turn off external irradiation in HELIOS by setting the stellar temperature to 0.

HELIOS requires opacity tables of all significant absorbers. We generate correlated-k tables for C_2 and C_2H_4 using the `ktable.py` script included with HELIOS, starting from the high-resolution opacities in the DACE database.¹⁹ C_2H_4 is a stand-in for a generic molecule with carbon–hydrogen bonds. C_3 is not in the DACE database, so we first generate opacities in DACE format from the line list (A. E. Lynas-Gray et al. 2024) using HELIOS-K (S. L. Grimm et al. 2021), before running `ktable.py` to produce the correlated-k table.

¹⁹ <https://dace.unige.ch/opacityDatabase/>

Table 2
Gas Volume Mixing Ratios, Cloud Parameters, and Internal Temperatures
Adopted in the HELIOS Models

	Day	Quadrature	Night
VMR_{He}	0.994	0.997	0.9994
VMR_{C_3}	3×10^{-3}	1.5×10^{-3}	1×10^{-4}
VMR_{C_2}	1×10^{-3}	5×10^{-4}	1×10^{-4}
$\text{VMR}_{\text{C}_2\text{H}_4}$	2×10^{-3}	1×10^{-3}	2×10^{-4}
$r_{\text{dust}} (\mu\text{m})$	0.03	0.04	0.1
$T_{\text{int}} (\text{K})$	2250	2000	1675

For all our HELIOS models, we adopt a radius of $1 R_J$ and a gravity of 4000 cm s^{-2} . The atmosphere is assumed to be helium-dominated, with the gas abundances specified in Table 2. In our prescription, non-helium molecular abundances are lower on the nightside than on the dayside, because larger molecules with unknown opacities (such as C_4 and C_5) are more abundant under chemical equilibrium, and because more carbon should condense into clouds. All models have a graphite “cloud” (dust) extending upward from 0.1 bar, with a cloud bottom mixing ratio of graphite particles of 1×10^{-14} and a scale height of 1.5 times the gas scale height. The only difference in cloud properties between day and night is in the particle radius, for which we prescribe $0.03 \mu\text{m}$ on the dayside, $0.04 \mu\text{m}$ at quadrature, and $0.1 \mu\text{m}$ at night. These radii were tuned to mute the spectral features, which we find to be much too strong in a clear atmosphere. We also tune the one crucial parameter that controls the temperature structure of the atmosphere— T_{int} , which quantifies the internal heat flux—so that the average brightness temperature of the resulting spectrum matches the observed brightness temperature. A T/P profile more isothermal than predicted can also mute spectral features, reducing the need for a cloud/dust layer. Given the large uncertainties in the heating mechanism, abundances, and opacity sources, we do not consider HELIOS’ T/P profile prediction to be robust.

These HELIOS models show that a cloudy helium-dominated atmosphere with trace amounts of C_3 , C_2 , and a hydrocarbon can reproduce the main features of the spectrum at all phases. They are not intended to be a fully accurate model of the atmosphere. Many parameters—such as the gas abundances, cloud particle mixing ratio, and cloud scale height—are simply assumed. In addition, many potentially important physical processes are not modeled, such as the temperature inhomogeneity across the planetary disk, the nonspherical nature of the planet, and the nonconstant effective gravity across the planet. Finally, while we use C_2H_4 as a stand-in for a molecule with C–H bonds, we do not expect it to survive at 2000 K temperatures. More likely candidates are C_2H and C_3H , which equilibrium chemistry calculations predict to be more abundant. We will be able to include their opacities once high-temperature line lists are available.

Since HELIOS is a 1D plane-parallel model, it does not account for dayside inhomogeneity. The irradiation, temperature, effective gravity, and viewing angle all vary significantly across the disk. We will explore the impact of the plane-parallel assumption in future work.

Appendix E GCMs

To model the atmosphere of the pulsar companion, we use the 3D GCM MITgcm (A. Adcroft et al. 2004) in the ADAM framework, which solves the primitive equations of meteorology on a cubed-sphere grid and has previously been used to model hot gaseous planets extensively (A. P. Showman et al. 2009; N. K. Lewis et al. 2014; X. Tan & T. D. Komacek 2019; V. Parmentier et al. 2021; M. E. Steinrueck et al. 2021; A. Roth et al. 2024; X. Tan et al. 2024). These equations are coupled with a double-gray radiative transfer scheme (see T. D. Komacek et al. 2017 for details) routine, with the thermal absorption coefficient calculated as a power law in pressure (T. D. Komacek et al. 2017). Our irradiation scheme (described below) deposits a set total amount of heat in the atmosphere centered upon a prescribed pressure. As such, there is no shortwave heating in the double-gray radiative transfer scheme. We set the internal temperature, T_{int} , to be 500 K, such that the net thermal flux emerging from the bottom of the domain is σT_{int}^4 .

To our knowledge, no GCM of a pulsar planet has been published previously (however, GCMs of neutron stars have been explored; see J. Nätilä et al. 2024). As discussed in the previous subsection, the companion is primarily heated by the pulsar’s GeV gamma rays, which deposit their energy at a typical pressure of ~ 10 bar at the subpulsar point (and $\sim 10 \cos(\theta)$ bar at an angular distance θ from the subpulsar point). We model this heating by artificially adding in a heat source at a characteristic pressure of $P_0 \cos \theta$, where $P_0 = 10$ bar and θ is the angular distance from the substellar point. The flux received from the pulsar is

$$F = \frac{L_H}{4\pi a^2} \cos \theta, \quad (\text{E1})$$

where L_H is the pulsar heating luminosity. We adopt $L_H = 9.1 \times 10^{32} \text{ erg s}^{-1}$, from the ICARUS light-curve model fit. We distribute this flux as a Gaussian in $\ln(P)$ space with a standard deviation of 1 (roughly corresponding to one scale height):

$$\frac{dF}{d\ln P} = \frac{L_H}{4\pi a^2} \cos \theta \exp\left(-\frac{1}{2} \ln^2 \left[\frac{P}{P_0 \cos \theta} \right]\right) / \sqrt{2\pi}. \quad (\text{E2})$$

To convert $\frac{dF}{d\ln P}$ into a volumetric heating rate, we multiply by $|\frac{d\ln P}{dz}|$. In hydrostatic equilibrium, $P = P_a \exp(-z/H)$ over infinitesimally small dz , so $|\frac{d\ln P}{dz}| = 1/H$. We obtain a volumetric heating rate Γ of

$$\Gamma = \frac{L_H}{4\pi a^2} \cos \theta \frac{1}{\sqrt{2\pi} H} \exp\left(-\frac{1}{2} \ln^2 \left[\frac{P}{P_0 \cos \theta} \right]\right), \quad (\text{E3})$$

where H is the *local* scale height. Dividing Γ by ρ gives the heating rate per unit mass, which we put into the GCM following A. P. Showman et al. (2009) as a simple prescription for gamma-ray heating.

As the physical properties of the companion are not well constrained, we take the mass and radius of the companion to be $1.5 M_{\text{Jup}}$ and $1 R_{\text{Jup}}$, respectively. We use the published rotation period of 0.322 days (R. Spiewak et al. 2017). We use a




Figure 10. Temperature structure from 3D GCMs. Here, we show maps of the time-averaged temperature and streamlines at two isobars (left) and equatorial temperature–pressure profiles (right) for two different heat deposition cases. The subpulsar point is indicated with a star. The isobaric projections correspond to roughly 20 mbar (the top row of the maps)—shown by the blue substellar marker and dashed line in the T – P profiles—and 2 bar (the bottom row of the maps)—corresponding to the purple dashed line and marker from our GCMs with a bespoke heating scheme given heat depositions centered at 3 bar and 0.3 bar. The combination of a rapid rotation rate and deep heating results in the circulation pattern. Specifically, we find westward-shifted midlatitude hot regions at the 2 bar pressure level for both cases of heat deposition. However, at 20 mbar, the model with more shallow heating shows a different temperature structure that is more symmetric and unlikely to produce a westward offset. In the more deeply heated 3 bar case, the hot region being centered at negative longitudes implies a flux peak after phase 0.5, as observed.

dynamical time step of 10 s for each simulation, spanning from roughly 100 – 10^{-6} bars in pressure over 70 vertical levels. We assume a mean molecular weight of 4 amu and set the heat capacity to that of a monoatomic gas, reflecting the likely helium-dominated composition of the atmosphere. Each model was run for at least 1500 planetary days to ensure sufficient spinup time.

We plot the isobaric temperature and streamlines for two pressure levels as well as equatorial temperature–pressure profiles in Figure 10 for heating centered at 3 bar and 0.3 bar. The subpulsar point is indicated by the star, whose color shows the corresponding pressure level on the temperature–pressure profiles (purple for roughly 2 bar and blue for 20 mbar). At 2 bar, both cases show midlatitude hot anticyclones that are centered west of the substellar point. The resulting westward offset is a result of the rapid rotation of the planet; see E. K. H. Lee et al. (2020), X. Tan & A. P. Showman (2020), and R. Zhan et al. (2024) for previous studies of planets in this regime. Because the pulsar planet is viewed at low inclination, the midlatitudes contribute a disproportionate percentage of the total flux, further increasing the westward offset in the data. Notably, at 20 mbar, the more shallowly heated 0.3 bar model diverges from this circulation pattern and appears unlikely to produce a westward offset. This suggests that this heating may be too shallow to explain the JWST phase curve of PSR J2322–2650b, and instead deep heating is required to match the observations.

We note that our GCMs display stronger heat redistribution than in our observed phase curves (see Figure 5), particularly the 0.3 bar model. Future work will improve the realism of the GCMs by exploring a larger model parameter space, including further variation of the heat deposition, examining the influence of the infrared opacity, and self-consistently postprocessing the model to produce spectroscopic phase curves to gauge the magnitude of the westward offset.

Appendix F Cross-correlation

Cross-correlation is used extensively to detect molecules in high-resolution ($R \gtrsim 20,000$) transmission and emission

spectra of exoplanets (I. Snellen 2025). JWST spectral resolutions are much lower, but cross-correlation has nevertheless revealed CO in the NIRSpec G395H transmission spectrum of WASP-39b (E. Esparza-Borges et al. 2023).

To check for the presence of a molecule in the atmosphere, we first make a template from $-\ln(\sigma_\lambda)$, where σ_λ is the absorption cross section of the molecule at a fiducial temperature of 2000 K and 10 mbar (a rough approximation of the planet’s infrared photospheric conditions). The precise pressure is unimportant, because thermal broadening dominates over pressure broadening at $P < 1$ bar, for the default pressure broadening assumed by ExoMol of $0.07 \text{ cm}^{-1} \left(\frac{300 \text{ K}}{2000 \text{ K}} \right)^{0.5} = 0.03 \text{ cm}^{-1}$. We generate the C_2 cross sections with ExoCross (S. N. Yurchenko et al. 2018) at $R = 10^6$, take their natural log, broaden the result with a Gaussian filter to the maximum instrumental resolution ($R = 4200$), interpolate onto a wavelength grid with spacing of $\lambda/16,000$, and then detrend using a uniform filter with a width of 20 (equivalent to $\lambda/800$). This broadened and detrended $-\ln(\sigma_\lambda)$ is the cross-correlation template.

We process the G235H data into a residuals grid using a methodology similar to that used in HRCCS. We first compute the median spectrum across the whole observation and fit a second-order polynomial to the region blueward of $3 \mu\text{m}$ (to avoid the C_3 cliff at $3.014 \mu\text{m}$) to estimate the continuum. We divide every spectrum by this polynomial, then divide it again by the median of the continuum-corrected spectrum to account for planetary flux variations over time. Next, we subtract the column median from every column (corresponding to one wavelength). This subtraction step removes all spectral or pseudo-spectral features due to bad pixels, the thermal background, or flatfielding errors, so long as the features do not vary with time. The planet’s spectral features are mostly protected from this subtraction, because they shift over time due to orbital motion. Finally, to match the processing done on the template, we detrend each row of the residuals grid with a uniform filter of width 20 and interpolate it onto the same wavelength grid.

Figure 11 shows the cross-correlation value as a function of time and blueshift. C_2 is strongly detected, and the orbital motion of the planet is clearly visible. Using exactly the same methodology, CO is not detected. CO’s absorption cross section is many orders of magnitude smaller than C_2 ’s throughout most of the G235H bandpass, but is only ~ 10 times smaller at 2.30 – $2.47 \mu\text{m}$. To improve sensitivity to CO, we tried only using this wavelength range but still obtained no detection. We do, however, detect C_2 at 6.1σ .

To derive a robust model-independent upper limit on the CO-to- C_2 abundance ratio, we perform an analysis based solely on the cross sections of CO and C_2 . We model the brightness temperature as $T_b(\lambda) = T_0 - r \ln \sigma(\lambda)$, where $\sigma(\lambda) = f\sigma_{\text{C}_2}(\lambda) + (1 - f)\sigma_{\text{CO}}(\lambda)$ is the absorption cross section at 2000 K and 10 mbar. T_0 is computed so that the median brightness temperature over 2.30 – $2.47 \mu\text{m}$ equals 2350 K (in accordance with observations). We convert T_b to F_ν with Planck’s law, bin the $R = 10^6 F_\nu$ down to an instrumental resolution of $R = 3100$ with a Gaussian filter, interpolate to a wavelength grid with spacing of $\lambda/6000$ (equal to the G235H wavelength spacing at $2.39 \mu\text{m}$), and add random Gaussian noise to each data point with a standard deviation equal to 8% of the mean F_ν (equal to the error on the real observations). The cross-correlation value at $V_{\text{sys}} = 0$ of the detr-




Figure 11. Cross-correlation value between the residuals grid and the C_2 template (left) or the CO template (right), as a function of time/phase and blueshift. The negative “shadow” on the left is from self-subtraction during the columnwise subtraction step.

ended simulated observations with the detrended $\ln \sigma_{C_2}(\lambda)$ ($\ln \sigma_{CO}(\lambda)$), divided by the standard deviation of the CCFs in the $1250 < |V_{\text{sys}}| < 5000 \text{ km s}^{-1}$ regions, is the detection significance of C_2 (CO). For each f , we first tune r so that the detection significance of C_2 reaches 6.0σ . We then compute the detection significance of CO with the same r . This exercise reveals that at $f=1.0$, $r=50 \text{ K}$ is required to detect C_2 at 6.0σ , at which point CO is (unsurprisingly) “detected” at 0.1σ . At $f=0.25$, r must be 50 K , and CO is detected at 2.4σ . The detection significance of CO increases to 3.0σ at $f=0.15$ ($r=75 \text{ K}$) and to 3.5σ at $f=0.1$ ($r=80 \text{ K}$). These results are highly insensitive to both r and the injected noise, provided that the combination of the two results in a 6σ C_2 detection. To prove this point, we double our injected noise from 8% to 16% and find that at $f=0.25$, $r=140 \text{ K}$ achieves a 6.1σ C_2 detection and a 2.4σ CO detection—identical to our fiducial results. We also halve our injected noise from 8% to 4% and find that at $f=0.25$, $r=35 \text{ K}$ results in an identical 6.1σ C_2 detection and 2.4σ CO detection.

Appendix G PRISM Phase-curve Model

We construct 3D models of the heated companion surface in the Roche geometry, assuming that the pulsar heating is dominated by gamma rays, as is true for other black widows. Then, using the ICARUS binary light-curve code (R. P. Breton et al. 2012), with improvements including surface winds and more realistic gravity and limb darkening (D. Kandel & R. W. Romani 2020), we integrate over the companion surface at each orbital phase, computing the emergent spectrum and light curve. Normally, this employs a grid of stellar atmospheres, but at present we do not adequately understand the temperature/pressure variations of the very strong molecular carbon bands to compute reliable emergent spectra for the range of temperatures on the companion surface. Nevertheless, the PRISM phase curves show clear thermal modulation, with near-featureless Planckian spectra at minimum. Therefore, we isolate (as most continuum-dominated) two bands away from the strongest broad absorption features and fit these to light-curve models, assuming Planck emission at the local effective temperature.

The light-curve model fits for the companion’s base (nighttime) temperature, the isotropic equivalent pulsar heating power, the system distance, and the orbital inclination. We find that all fits converge to a high Roche-lobe fill factor and so fix the ratio of the nose radius to the L1 point distance to $f_1 = 0.99$. The phase offset requires a significant surface wind, where the wind strength is characterized by the ratio of radiation to advection timescales $\epsilon = \tau_{\text{rad}}/\tau_{\text{adv}} \approx -0.12$; the negative sign indicates westward flow predominates at the latitudes that dominate the observed emission. While we use a standard analytic model for the limb-darkening coefficients based on solar values (H. Neckel 2005), we find that the final results are sensitive to this assumption. Lacking detailed atmospheric emission models, we parameterize this sensitivity with a simple rescaling of the limb-darkening amplitude.

This simple model does well at explaining the bulk of the light curve, indicates a low inclination, and gives a photometric distance consistent with the parallax constraint. The Earth line-of-sight views shown below the curve show how even the night-phase spectral energy distribution (SED) will be dominated by dayside fluxes from moderate latitude, viewed close to the limb. However, the model does not capture the nighttime variability of the planet from orbit to orbit—the nightside is hotter the first time we observe it than the second, and by a larger amount at shorter wavelengths. This may be from an eruptive event, as seen in other black-widow systems (e.g., R. W. Romani et al. 2015), although such eruptions are associated with very strong pulsar heating and/or magnetic activity on the companion, neither of which apply here. Alternatively, variable high-latitude winds at the terminator might cause fluctuations in emissivity or limb darkening. This variability at orbital minimum induces a large systematic that dominates our present inclination uncertainty and hence pulsar mass uncertainty, exacerbated by our poor knowledge of the true limb darkening. We report parameters for fits to the full light curves, and for fits excluding the first or second minimum, in Table 3. Note that systematic errors dominate. The fit distance is at the large side of the parallax-allowed range; imposing this range as a prior makes little difference to the best-fit parameters. The parameters are particularly sensitive to the limb-darkening enhancement C_{LD} ; because the system is viewed at small inclination i , this allows a large

Table 3
Light-curve Fit Parameters

	All Phases	−0.12 to 0.88	0.07–1.06
T_N (K)	$904 \pm 15 \pm 40$	833	915
$L_H(10^{32}\text{erg s}^{-1})$	$9.1 \pm 0.2 \pm 2$	11.4	7.7
i (deg)	$30.9 \pm 0.5 \pm 2.6$	28.1	33.1
d (pc)	$750 \pm 10 \pm 45$	800	710
ϵ	-0.123 ± 0.003	−0.133	−0.124
C_{LD}	$2.1 \pm 0.2 \pm 0.8$	2.9	1.4
χ^2/DOF	2.22	1.74	2.00
$M_{\text{PSR}} (M_{\odot})$	$1.97 \pm 0.1 \pm 0.5$	2.56	1.65

Note. The first errors are $\pm 1\sigma$ fit errors, inflated by χ^2/DoF , while the second errors represent systematics from the range of the three fits.

range of inclinations and hence of system masses. The PRISM data quality is such that, if reliable emergent spectra and limb darkening could be computed, we could fit to the full spectral range and substantially constrain the allowed parameters. Until better atmosphere models are available, SEDs covering several orbits would be required to understand the minimum variability and establish the true thermal minimum and precise orbital inclination.

At first sight, the large inferred isotropic heating power L_H is puzzling. However, the pulsar gamma-ray emission—and to a lesser extent the particle flux—is concentrated toward the spin equator, especially for low magnetic inclination. Also, for large masses and the preferred stiff equation of state (J. M. Lattimer & B. F. Schutz 2005), the pulsar moment of inertia is large, with $I_{45} \approx 1.2(M/M_{\odot})^{1.5}$, increasing the inferred pulsar spindown luminosity and magnetic field. For example, with $M_{\text{PSR}} = 2M_{\odot}$, we infer $\dot{E} \approx 6 \times 10^{32}\text{erg s}^{-1}$ and so the heating power should be beamed by >1.5 times toward the equator. Thus, the strong companion heating implies strong beaming, possibly in combination with a high pulsar mass.

Perhaps the most interesting system parameter is the pulsar mass, M_{PSR} , since a precise measurement should help us understand the accreted increment and the associated decrease in the pulsar dipole field. If large, the mass might also contribute to our understanding of the dense matter equation of state. At present, our fits are not constraining. However, as noted above, with improved modeling of the molecular carbon-dominated atmosphere, the JWST data can allow us to address these important questions.

Appendix H Equilibrium Chemistry Grid

To derive conservative constraints on the elemental C/O and C/N ratios from our C_2/CO and C_2/CN lower limits, themselves derived from cross-correlation of the G235H spectra, we generate two equilibrium chemistry abundance grids with FastChem. The grids cover $\log_{10}(\text{C}/\text{He})$ of -5 to 1 in steps of 0.05 , temperature of $600\text{--}4000\text{ K}$ in steps of 100 K , and pressure of 10^{-9} bar to 10^3 bar in steps of 1 dex. The first grid covers $\log_{10}(\text{C}/\text{O})$ of -1 to 5 in 60 steps, while the second grid covers $\log_{10}(\text{C}/\text{N})$ of 0 to 7 in 60 steps. For reference, the solar photospheric values are $\log_{10}(\text{C}/\text{O}) = -0.22$ and $\log_{10}(\text{C}/\text{N}) = 0.63$. The first grid only includes the elements He, C, and O, while the second grid only includes He, C, and N.

In the first grid, $C_2/\text{CO} > 0.17$ is violated at all grid points unless $\text{C}/\text{O} > 2$. Restricting the temperature to $<2300\text{ K}$ strengthens the upper limit to $\text{C}/\text{O} > 11$. Further imposing the restrictions $\text{C}/\text{He} > 0.0035$ (the solar value) and $P > 1$ mbar raises the upper limit to $\text{C}/\text{O} > 100$.

In the second grid, $C_2/\text{CN} > 32$ is violated at all grid points unless $\text{C}/\text{N} > 300$. Restricting the temperature to $<2300\text{ K}$ raises the limit to $\text{C}/\text{N} > 2000$. Further imposing the restrictions $\text{C}/\text{He} > 0.0082$ (the solar value) and $P > 1$ mbar raises the upper limit to $\text{C}/\text{N} > 10,000$.

ORCID iDs

Michael Zhang  <https://orcid.org/0000-0002-0659-1783>
 Maya Beleznyay  <https://orcid.org/0000-0003-1285-8170>
 Timothy D. Brandt  <https://orcid.org/0000-0003-2630-8073>
 Roger W. Romani  <https://orcid.org/0000-0001-6711-3286>
 Peter Gao  <https://orcid.org/0000-0002-8518-9601>
 Hayley Beltz  <https://orcid.org/0000-0002-6980-052X>
 Matthew Bailes  <https://orcid.org/0000-0003-3294-3081>
 Matthew C. Nixon  <https://orcid.org/0000-0001-8236-5553>
 Jacob L. Bean  <https://orcid.org/0000-0003-4733-6532>
 Thaddeus D. Komacek  <https://orcid.org/0000-0002-9258-5311>
 Brandon P. Coy  <https://orcid.org/0000-0002-0508-857X>
 Guangwei Fu  <https://orcid.org/0000-0002-3263-2251>
 Rafael Luque  <https://orcid.org/0000-0002-4671-2957>
 Daniel J. Reardon  <https://orcid.org/0000-0002-2035-4688>
 Emma Carli  <https://orcid.org/0000-0003-3265-2866>
 Ryan M. Shannon  <https://orcid.org/0000-0002-7285-6348>
 Jonathan J. Fortney  <https://orcid.org/0000-0002-9843-4354>
 Anjali A. A. Piette  <https://orcid.org/0000-0002-4487-5533>
 M. Coleman Miller  <https://orcid.org/0000-0002-2666-728X>
 Jean-Michel Desert  <https://orcid.org/0000-0002-0875-8401>

References

- Abia, C., de Laverny, P., & Wahlin, R. 2008, *A&A*, **481**, 161
 Abdollahi, S., Acero, F., Baldini, L., et al. 2022, *ApJS*, **260**, 53
 Adcroft, A., Campin, J.-M., Hill, C., & Marshall, J. 2004, *MWRv*, **132**, 2845
 Astropy Collaboration, Price-Whelan, A. M., Lim, P. L., et al. 2022, *ApJ*, **935**, 167
 Astropy Collaboration, Price-Whelan, A. M., Sipőcz, B. M., et al. 2018, *AJ*, **156**, 123
 Astropy Collaboration, Robitaille, T. P., Tollerud, E. J., et al. 2013, *A&A*, **558**, A33
 Bailes, M., Bates, S. D., Bhalariao, V., et al. 2011, *Sci*, **333**, 1717
 Bailes, M., Jameson, A., Abbate, F., et al. 2020, *PASA*, **37**, e028
 Bell, T. J., Dang, L., Cowan, N. B., et al. 2021, *MNRAS*, **504**, 3316
 Benvenuto, O. G., De Vito, M. A., & Horvath, J. E. 2012, *ApJL*, **753**, L33
 Bisnovatyi-Kogan, G. S., & Komberg, B. V. 1974, *SvA*, **18**, 217
 Brandt, T. D. 2024a, *PASP*, **136**, 045004
 Brandt, T. D. 2024b, *PASP*, **136**, 045005
 Breton, R. P., Rappaport, S. A., van Kerkwijk, M. H., & Carter, J. A. 2012, *ApJ*, **748**, 115
 Bushouse, H., Eisenhamer, J., Dencheva, N., et al., 2025 JWST Calibration Pipeline v1.18.0, Zenodo, doi:10.5281/zenodo.6984365
 Chabrier, G., & Debras, F. 2021, *ApJ*, **917**, 4
 Dang, L., Cowan, N. B., Schwartz, J. C., et al. 2018, *NatAs*, **2**, 220
 Dhillon, V. S., Kennedy, M. R., Breton, R. P., et al. 2022, *MNRAS*, **516**, 2792
 Di Criscienzo, M., Dell’Aglia, F., Ventura, P., et al. 2013, *MNRAS*, **433**, 313
 Draghis, P., Romani, R. W., Filippenko, A. V., et al. 2019, *ApJ*, **883**, 108
 Ellis, J. A., Vallisneri, M., Taylor, S. R., & Baker, P. T., 2020 ENTERPRISE: Enhanced Numerical Toolbox Enabling a Robust Pulsar Inference Suite, v3.0.0, Zenodo, doi:10.5281/zenodo.4059815
 Esparza-Borges, E., López-Morales, M., Adams Redai, J. I., et al. 2023, *ApJL*, **955**, L19
 Foreman-Mackey, D. 2016, *JOSS*, **1**, 24

- Grimm, S. L., Malik, M., Kitzmann, D., et al. 2021, *ApJS*, **253**, 30
- Guo, Y., Wang, B., & Han, Z. 2022, *MNRAS*, **515**, 2725
- Hansen, B. M. S. 1996, PhD thesis, California Institute of Technology
- Hobbs, G. B., Edwards, R. T., & Manchester, R. N. 2006, *MNRAS*, **369**, 655
- Horne, K. 1986, *PASP*, **98**, 609
- Kandel, D., & Romani, R. W. 2020, *ApJ*, **892**, 101
- Kandel, D., & Romani, R. W. 2023, *ApJ*, **942**, 6
- Kennedy, M. R., Breton, R. P., Clark, C. J., et al. 2022, *MNRAS*, **512**, 3001
- Kerley, G. I., & Chhabildas, L. 2001, Multicomponent-multiphase Equation of State for Carbon SAND2001-2619, Sandia National Laboratories <https://www.osti.gov/servlets/purl/787608>
- Koljonen, K. I. I., & Linares, M. 2025, arXiv:2505.11691
- Komacek, T. D., Showman, A. P., & Tan, X. 2017, *ApJ*, **835**, 198
- Lattimer, J. M. 2012, *ARNPS*, **62**, 485
- Lattimer, J. M., & Schutz, B. F. 2005, *ApJ*, **629**, 979
- Lee, E. K. H., Casewell, S. L., Chubb, K. L., et al. 2020, *MNRAS*, **496**, 4674
- Lewis, N. K., Showman, A. P., Fortney, J. J., Knutson, H. A., & Marley, M. S. 2014, *ApJ*, **795**, 150
- Lynas-Gray, A. E., Polyansky, O. L., Tennyson, J., Yurchenko, S. N., & Zobov, N. F. 2024, *MNRAS*, **535**, 1439
- Madhusudhan, N. 2012, *ApJ*, **758**, 36
- Malik, M., Grosheintz, L., Mendonça, J. M., et al. 2017, *AJ*, **153**, 56
- Malik, M., Kitzmann, D., Mendonça, J. M., et al. 2019, *AJ*, **157**, 170
- Mehla, A., Kasliwal, M. M., Karambelkar, V., et al. 2025, *PASP*, **137**, 044201
- Miles, M. T., Shannon, R. M., Bailes, M., et al. 2023, *MNRAS*, **519**, 3976
- Miles, M. T., Shannon, R. M., Reardon, D. J., et al. 2025, *MNRAS*, **536**, 1467
- Moseley, S. H., Arendt, R. G., Fixsen, D. J., et al. 2010, *Proc. SPIE*, **7742**, 77421B
- Mukherjee, D. 2017, *JApA*, **38**, 48
- Nättilä, J., Cho, J. Y. K., Skinner, J. W., Most, E. R., & Ripperda, B. 2024, *ApJ*, **971**, 37
- Neckel, H. 2005, *SoPh*, **229**, 13
- Nelemans, G., Yungelson, L. R., van der Sluys, M. V., & Tout, C. A. 2010, *MNRAS*, **401**, 1347
- Nixon, M. C., & Madhusudhan, N. 2021, *MNRAS*, **505**, 3414
- Nixon, M. C., Piette, A. A. A., Kempton, E. M. R., et al. 2024, *ApJL*, **970**, L28
- Paczyński, B. 1971, *ARA&A*, **9**, 183
- Parmentier, V., Showman, A. P., & Fortney, J. J. 2021, *MNRAS*, **501**, 78
- Philippov, A., Timokhin, A., & Spitkovsky, A. 2020, *PhRvL*, **124**, 245101
- Rauscher, B. J. 2024, *PASP*, **136**, 015001
- Regan, M. 2024, Detection and Flagging of Showers and Snowballs in JWST Technical Report JWST-STScI-008545, Space Telescope Science Institute
- Romani, R. W. 1990, *Natur*, **347**, 741
- Romani, R. W., Filippenko, A. V., & Cenko, S. B. 2015, *ApJ*, **804**, 115
- Romani, R. W., Graham, M. L., Filippenko, A. V., & Zheng, W. 2016, *ApJ*, **833**, 138
- Romani, R. W., Kandel, D., Filippenko, A. V., Brink, T. G., & Zheng, W. 2022, *ApJL*, **934**, L17
- Roth, A., Parmentier, V., & Hammond, M. 2024, *MNRAS*, **531**, 1056
- Saio, H., & Jeffery, C. S. 2002, *MNRAS*, **333**, 121
- Shamohammadi, M., Bailes, M., Flynn, C., et al. 2024, *MNRAS*, **530**, 287
- Shklovskii, I. S. 1970, *SvA*, **13**, 562
- Showman, A. P., Fortney, J. J., Lian, Y., et al. 2009, *ApJ*, **699**, 564
- Simpson, J. A., Linares, M., Casares, J., et al. 2025, *MNRAS*, **536**, 2169
- Sing, D. K., Evans-Soma, T. M., Rustamkulov, Z., et al. 2024, *AJ*, **168**, 231
- Snellen, I. 2025, *ARA&A*, **63**, 83
- Spiewak, R., Bailes, M., Barr, E. D., et al. 2017, *MNRAS*, **475**, 469
- Steinrueck, M. E., Showman, A. P., Lavvas, P., et al. 2021, *MNRAS*, **504**, 2783
- Stock, J. W., Kitzmann, D., & Patzer, A. B. C. 2022, *MNRAS*, **517**, 4070
- Stock, J. W., Kitzmann, D., Patzer, A. B. C., & Sedlmayr, E. 2018, *MNRAS*, **479**, 865
- Taam, R. E., & van den Heuvel, E. P. J. 1986, *ApJ*, **305**, 235
- Tammer, M. 2004, *Colloid Polym Sci*, **283**, 235
- Tan, X., & Komacek, T. D. 2019, *ApJ*, **886**, 26
- Tan, X., Komacek, T. D., Batalha, N. E., et al. 2024, *MNRAS*, **528**, 1016
- Tan, X., & Showman, A. P. 2020, *ApJ*, **902**, 27
- Tennyson, J., & Yurchenko, S. N. 2018, *Atoms*, **6**, 26
- van Kerkwijk, M. H., Breton, R. P., & Kulkarni, S. R. 2011, *ApJ*, **728**, 95
- van Straten, W., Demorest, P., & Osłowski, S. 2012, *AR&T*, **9**, 237
- Webbink, R. F. 1984, *ApJ*, **277**, 355
- Yurchenko, S. N., Al-Refaie, A. F., & Tennyson, J. 2018, *A&A*, **614**, A131
- Yurchenko, S. N., Szabó, I., Pyatenko, E., & Tennyson, J. 2018, *MNRAS*, **480**, 3397
- Zhan, R., Koll, D. D. B., & Ding, F. 2024, *ApJ*, **971**, 125
- Zhang, M. 2025, Plots for the Paper “A Carbon-rich Atmosphere on a Windy Pulsar Planet,” v1.0, Zenodo, doi:10.5281/ZENODO.17400581