

PROBING NEUTRON STAR CRUST PROPERTIES VIA ACCRETION-INDUCED HEATING  
AND COOLING

By

Justin Grace

A DISSERTATION

Submitted to  
Michigan State University  
in partial fulfillment of the requirements  
for the degree of

Astrophysics and Astronomy—Doctor of Philosophy  
Computational Mathematics, Science and Engineering—Dual Major

2024

## ABSTRACT

Accreting neutron stars in low-mass X-ray binary (LMXB) systems provide an avenue for studying the thermal properties of the crust and core. When material falls onto a neutron star's surface, the matter in its outer layers is compressed and nuclear reactions are induced. In quasi-persistent transient systems, accretion outbursts last years to decades followed by periods of quiescence that can last years to decades. During an outburst, the accretion-induced reactions heat the crust out of thermal equilibrium with the core. When accretion ends, the core cools back into thermal equilibrium with the core which we can observe as a decrease in surface temperature: the cooling curve. The shape of the cooling curve depends on the thermal properties of the crust.

In this work, I model the thermal evolution of the crust using dStar. By fitting dStar models to observed cooling curves, we can estimate the properties of the crust such as the core temperature ( $T_c$ ), the impurity of the composition ( $Q_{\text{imp}}$ ), and the accretion-induced heating in the shallow layers ( $Q_{\text{sh}}$ ) and depths of the crust ( $Q_{\text{in}}$ ).

Prior to this work, the cooling curves of several LMXBs have been independently analyzed. It is unknown to what extent different neutron stars share crust properties, such as composition and accretion-induced heating. In this thesis, I perform joint fits of five LMXBs, where I simultaneously fit the cooling curves of all sources with the value of  $Q_{\text{imp}}$ ,  $Q_{\text{sh}}$ , or  $Q_{\text{in}}$  being shared across all sources. I compare the goodness-of-fit of the joint fits to independent fits for each source. I find that jointly fitting  $Q_{\text{imp}}$  or  $Q_{\text{sh}}$  has little impact on the quality of the fits, suggesting that the sources either do share the same value for these parameters or that the data does not sufficiently constrain them. When jointly fitting  $Q_{\text{sh}}$ , the quality of the fits significantly decreases, suggesting that different sources must have different values of  $Q_{\text{sh}}$ .

The predicted composition, accretion-induced heating rates, and location of heat release depend on the behavior of free neutrons in the deep crust. Crust models that allow free neutrons to diffuse throughout the inner crust predict that the crust is more pure and the inner heating rate is lower than models that do not. Additionally, some nuclear models predict the existence of non-spherical “pasta” phases of nuclear matter at the bottom of the crust. A pasta layer acts as a thermally

insulating layer in thermal evolution models and its presence affects the rate of cooling. I compare models that allow neutron diffusion to those that do not, and models with pasta layers to those that do not by estimating the Bayesian evidence with nested modeling. I find that the cooling curves of the five LMXBs to which I fit the models do not consistently favor the same models. This is inconsistent with expectations because the the existence of pasta and diffusion of free neutrons depend solely on the local density; in particular, they are not expected to depend on the accretion history and therefore should not differ among neutron stars.

Copyright by  
JUSTIN GRACE  
2024

## **ACKNOWLEDGEMENTS**

This work was enabled by the National Science Foundation under grant PHY-1430152 (JINA Center for the Evolution of the Elements), and supported by the NSF under grant AST-1812838 and by NASA under grant 80NSSC20K0503. We thank L. Ootes for providing the tabulated accretion rates for KS 1731-260 and Rahul Jain for tabulated crust compositions.

## TABLE OF CONTENTS

LIST OF TABLES . . . . .	vii
LIST OF FIGURES . . . . .	viii
LIST OF ABBREVIATIONS . . . . .	x
CHAPTER 1 INTRODUCTION . . . . .	1
1.1 History . . . . .	2
1.2 Neutron Star Structure . . . . .	4
1.3 Investigating the Crust Properties of LMXBs in Quiescence . . . . .	9
CHAPTER 2 COOLING QUIESCENT NEUTRON STARS . . . . .	11
2.1 The Composition of the Accreted Crust . . . . .	12
2.2 Observations of Accreting Neutron Stars . . . . .	18
2.3 Modelling the Thermal Evolution of the Crust . . . . .	19
2.4 Fitting Model Parameters to Observational Data . . . . .	23
2.5 Accounting for Variable Accretion History . . . . .	27
CHAPTER 3 JOINTLY FITTING CRUSTAL PROPERTIES OF MULTIPLE NEU- TRON STARS . . . . .	37
3.1 Independently Fitting Neutron Star Models to Observations . . . . .	38
3.2 Testing Whether All Sources Share Crust Properties . . . . .	59
3.3 Conclusion . . . . .	64
CHAPTER 4 THE EFFECT OF NEUTRON DIFFUSION ON ACCRETION-INDUCED CRUSTAL HEATING . . . . .	67
4.1 Neutron Diffusion . . . . .	67
4.2 Incorporating Neutron Diffusion Into Cooling Models . . . . .	68
4.3 Model Selection . . . . .	73
4.4 Comparison of Models . . . . .	74
4.5 Conclusion . . . . .	78
CHAPTER 5 CONCLUSIONS . . . . .	79
5.1 Jointly Fitting Crustal Properties of Multiple Neutron Stars . . . . .	79
5.2 The Effect of Neutron Diffusion on Accretion-Induced Crustal Heating . . . . .	80
5.3 Future Work . . . . .	81
BIBLIOGRAPHY . . . . .	82
APPENDIX DETAILED RESULTS FOR MODELS WITH AND WITHOUT NEUTRON DIFFUSION . . . . .	94

## LIST OF TABLES

Table 3.1	Sources and fitting parameters . . . . .	39
Table 3.2	Bounds for independent fit priors . . . . .	40
Table 3.3	Best-fit parameters for all independent models . . . . .	58
Table 3.4	Best-fit parameters for independent fits and all joint fits . . . . .	62
Table 4.1	Approximate heating rates and zones for reaction network calculations . . . . .	71
Table 4.2	Bayes factor threshold for model selection. . . . .	74
Table 4.3	Evidence for non-diffusive and nHD models. . . . .	75

## LIST OF FIGURES

Figure 1.1	Density-temperature plane for dense matter . . . . .	2
Figure 1.2	Schematic of a neutron star crust . . . . .	6
Figure 1.3	Melting and Fermi temperatures in the crust and envelope. . . . .	8
Figure 2.1	Observed accretion rates of KS 1731–260 during its 12.5 year outburst . . . . .	29
Figure 2.2	Time-resolved versus average outburst accretion rates . . . . .	31
Figure 2.3	Different partitioning schemes for the accretion history . . . . .	33
Figure 2.4	First-point errors for different partitioning schemes . . . . .	35
Figure 2.5	Growth rates for exponential partitioning schemes . . . . .	36
Figure 3.1	Posterior distribution and cooling curve for KS 1731–260 . . . . .	44
Figure 3.2	Posterior distribution and cooling curve for MXB 1659–29 . . . . .	46
Figure 3.3	Posterior distribution and cooling curve for XTE J1701–462 . . . . .	48
Figure 3.4	Posterior distribution and cooling curve for EXO 0748–676 . . . . .	49
Figure 3.5	Posterior distribution and cooling curve for IGR J17480–2446 . . . . .	51
Figure 3.6	Posterior distribution and cooling curve for MAXI J0556–332 . . . . .	52
Figure 3.7	Posterior distribution and cooling curve for Aql X-1 . . . . .	55
Figure 3.8	Posterior distribution and cooling curve for Swift 174805.3–244637 . . . . .	56
Figure 3.9	Cooling curves for joint fits . . . . .	63
Figure 4.1	Composition and impurity profiles for reaction networks . . . . .	70
Figure 4.2	Heating profiles for reaction networks . . . . .	72
Figure 4.3	Cooling curves for non-diffusive and nHD models . . . . .	76
Figure A.1	Non-diffusive and nHD models for KS 1731–260 without a pasta layer . . . . .	96
Figure A.2	Same as Fig. A.1, but with a pasta layer . . . . .	97
Figure A.3	Non-diffusive and nHD models for MXB 1659–29 without a pasta layer . . . . .	98
Figure A.4	Same as Fig. A.3, but with a pasta layer . . . . .	99



Figure A.5	Non-diffusive and nHD models for XTE J1701–462 without a pasta layer . . . .	100
Figure A.6	Same as Fig. A.5, but with a pasta layer . . . . .	101
Figure A.7	Non-diffusive and nHD models for IGR J17480–2446 without a pasta layer . . .	102
Figure A.8	Same as Fig. A.7, but with a pasta layer . . . . .	103
Figure A.9	Non-diffusive and nHD models for Swift 174805.3–244637 without a pasta layer	104
Figure A.10	Same as Fig. A.9, but with a pasta layer . . . . .	105

## LIST OF ABBREVIATIONS

<b>bcc</b>	Body-centered cubic
<b>CLDM</b>	Compressible liquid-drop model
<b>dof</b>	Degrees of freedom
<b>EOS</b>	Equation of state
<b>LMXB</b>	Low-mass X-ray binary
<b>MCMC</b>	Markov chain Monte Carlo
<b>MCP</b>	Multi-component plasma
<b>nHD</b>	Neutron hydrostatic and diffusive equilibrium
<b>NS</b>	Neutron star
<b>OCP</b>	One-component plasma
<b>PRE</b>	Photospheric radius expansion
<b>TOV</b>	Tolman-Oppenheimer-Volkoff
<b>WS</b>	Wigner-Seitz

# CHAPTER 1

## INTRODUCTION

Neutron stars, which result from core-collapse supernovae, have played an important role in the development of astronomy and nuclear physics and are essential in advancing our understanding of stellar evolution and nuclear physics. They allow us to study matter in extreme conditions, in particular at high densities.

Neutron stars are observed in all bands across the electromagnetic spectrum. They were the first non-solar sources of X-rays detected and are the brightest point sources of X-ray and  $\gamma$ -ray emission. In recent years, binary mergers of neutron stars have produced detectable gravitational waves. They also sit in an important place in our understanding of stellar evolution. As one possible remnant of core-collapse supernova, they mark an end stage in the evolution of massive stars, and during binary mergers contribute to heavy element nucleosynthesis.

With masses and radii of about  $1.4 M_{\odot}$  and 10 km, respectively, neutron stars contain the densest matter in the universe. Their surface gravity exceeds  $10^{11}$  times that of the Earth, and their gravity is significantly non-Newtonian. The matter in their cores may reach densities 10–20 times that of a heavy atomic nucleus. Because of their extreme densities and gravity, neutron stars lie at the intersection of many branches of physics including nuclear physics and general relativity. In addition, neutron stars are host to a wealth of plasma processes and (magneto) hydrodynamics. Their extreme central densities make neutron stars especially appealing for testing theoretical models of cold, dense matter. The conditions in their cores are so far from terrestrial conditions that experimental data provides limited information that can describe the matter there (see Figure 1.1 for a schematic of the phase space explored by heavy ion collisions compared to neutron stars). The interiors of neutron stars are at high baryon chemical potential (i.e., super-nuclear density) and low temperature. As such, they are complementary to the conditions probed by heavy-ion collisions, which are at low chemical potential and high temperature. Under these conditions, other forms of matter (kaons, hyperons) could appear or the matter could be described in terms of quarks.

We don't need to look only at the core to find interesting nuclear physics processes. In the crust




Figure 1.1 Schematic of the chemical potential-temperature phase space.

and envelope, the layers outside of the core, the matter is still in an environment of strong gravity and its density spans over 14 orders of magnitude. In particular, neutron stars in binary systems that accrete matter from a companion exhibit many interesting nuclear processes in their outer layers. The continual accretion of matter induces nuclear reactions throughout the outer layers of the neutron star on timescales ranging from seconds to centuries.

Additionally, the electromagnetic emission that we observe from a neutron star originates at the surface. So in order to properly interpret the observational data necessary for probing the physics of the interior, we need to understand the properties of the outer layers. One way of studying the properties of the crust is to observe the thermal evolution of a neutron star after a period of active accretion.

## 1.1 History

Our understanding of neutron stars has been closely tied to the development of nuclear physics. Baade & Zwicky (1934) first theoretically predicted the existence of neutron stars just two years after the discovery of the neutron (Chadwick, 1932). For the next three decades, neutron stars remained firmly in the realm of theoretical curiosities. Tolman (1939) and Oppenheimer & Volkoff (1939) independently derived the equation for hydrostatic equilibrium of a spherically symmetric star in the framework of general relativity, known as Tolman-Oppenheimer-Volkoff (TOV) equation. Given an equation of state (EOS), which describes the relationship between pressure and density  $P(\rho)$ , one can integrate the TOV equation to obtain the mass-radius relationship  $M(R)$  of neutron stars.

Oppenheimer & Volkoff (1939) further showed that the TOV equation implied a maximum stable mass for a given EOS. Assuming neutron star matter was composed of non-interacting degenerate neutrons, they calculated a maximum mass of  $0.7 M_{\odot}$ . Cameron (1959) calculated a maximum mass of about  $2 M_{\odot}$ , with an EOS derived using the nuclear potential of Skyrme (1958). Pulsars in neutron star-white dwarf binaries have been measured to have masses up to  $\approx 2 M_{\odot}$  (Demorest et al., 2010; Antoniadis et al., 2013); see Özel & Freire (2016) for a review. These measurements provide evidence for a repulsive strong force at supranuclear densities.

The first widely accepted direct detection of a neutron star came in 1967, when Jocelyn Bell, a student of Antony Hewish at the University of Cambridge, discovered a rapidly pulsating radio source using the recently constructed Interplanetary Scintillation Array (Hewish et al., 1968)<sup>1</sup>. Three other similar pulsating sources were discovered in the following months. Hewish et al. proposed that they were either oscillating white dwarfs or neutron stars. Pacini (1967) and Gold (1968) independently proposed that these pulsating radio sources, or *pulsars*, were rotating magnetized neutron stars. Because of the extreme regularity of their pulsation periods, pulsars have been instrumental for measuring neutron star masses and testing general relativity. Notably, Hulse & Taylor (1975)<sup>2</sup> measured the orbital period of a pulsar in a binary system with another neutron star by an observed systematic variation in its pulsation period. After several years of observation, the orbital period was observed to decay, consistent with the predictions of general relativity (Taylor & Weisberg, 1982, 1989).

While the pulsar discovery of Hewish et al. was the first widely accepted direct detection of a neutron star, the first detection actually occurred several years earlier by the rocket-borne X-ray detectors of Giacconi et al. (1962)<sup>3</sup>. Giacconi et al. intended to detect fluorescent X-rays produced on the lunar surface by solar X-rays but incidentally detected X-rays coming from the direction of

---

<sup>1</sup>Hewish was awarded a share of the Nobel Prize in Physics in 1974 “for his decisive role in the discovery of pulsars.” Controversially, Bell was omitted from the award, despite her contribution to the construction of the radio telescope and analysis of the data that led to the discovery.

<sup>2</sup>Hulse and Taylor were awarded the Nobel Prize in Physics in 1993 “for the discovery of a new type of pulsar, a discovery that has opened up new possibilities for the study of gravitation.”

<sup>3</sup>Giacconi was awarded a share of the 2002 Nobel Prize in Physics “for pioneering contributions to astrophysics, which have led to the discovery of cosmic X-ray sources.”

the Scorpius constellation. As this source was the first X-ray source detected in the constellation Scorpius, it was given the designation Scorpius X-1 or Sco X-1. Follow-up observations identified the optical counterpart to the source (Sandage et al., 1966). Using the optical data, Shklovsky (1967) correctly argued that it was an accreting neutron star, though this conclusion was not widely accepted at the time. Sco X-1 was the first observed low-mass X-ray binary (LMXB), a binary system in which a lower mass donor star transfer mass onto a higher mass black hole or neutron star. Because of the compactness of the accretor, the infalling matter releases substantial gravitational potential energy which is emitted in the form of X-rays

Detailed X-ray and gamma ray observations came from space-based observatories, beginning with the launch of NASA's *Uhuru* satellite in 1970. In the following decades, more advanced X-ray observatories, including *Einstein*, *RXTE*, *XMM-Newton*, and *Chandra* expanded the catalog of known LMXBs and provided detailed observations of their behavior. These observations have uncovered phenomena such as X-ray bursts, thermonuclear explosions on the neutron star surface (see Strohmayer & Bildsten, 2004, for a detailed review of X-ray burst observations). Many LMXBs are transient, in which the source alternates between states of *outburst*, during which the X-ray luminosity is much higher, indicating a high accretion rate, and *quiescence*, during which accretion stops or is low enough to be unobservable (for a review, see King, 2004). Quiescence can last years to decades while outbursts typically last weeks to months. This varying accretion rate is probably caused by a thermo-viscous instability in the accretion disk, similar to dwarf novae in accreting white dwarf systems. Some transients, however, are *quasi-persistent*, in which outbursts can last years to decades. The origin of these quasi-persistent transients, however, is much less clear (see, e.g., King & Ritter, 1998; Dubus et al., 1999, 2001; Ritter et al., 2000). As we shall discuss later, this time-varying accretion induces nuclear reactions in the ashes of X-ray bursts; the response of the neutron star provides us with a probe of the interior physics.

## 1.2 Neutron Star Structure

Here I will summarize the structure of a neutron star. For a detailed description, see Haensel et al. (2007).

A neutron star with a mass  $M = 1.4 M_{\odot}$  and radius  $R = 10$  km will have a mean density exceeding the density of a heavy nucleus<sup>4</sup>,  $\rho_0 = 2.8 \times 10^{14} \text{ g cm}^{-3}$ . To satisfy hydrostatic equilibrium and mechanical stability, the pressure and density are greatest in the center of the neutron star and decrease outward. Therefore, the outer layers are composed of matter at sub-nuclear densities where the EOS is well-known from terrestrial experiments. The high-density interior is composed of uniform nuclear matter, but where  $\rho \lesssim 0.5\rho_0$ , the matter breaks into nuclei. The thickness of the layer composed of nuclei is roughly a pressure scale height,  $\sim P/(\rho g)$ , where the gravity  $g$  depends on  $M$  and  $R$ , and  $P$  and  $\rho$  are set by the nuclear equation of state at this transition from uniform matter to nuclei. The thickness of this outer layer is thus approximately 1 km and contains a mass of  $\sim 0.01 M_{\odot}$ .

This outer layer, which includes the *crust* and *envelope*, is composed of nuclei, electrons, and, in the deepest layers, free neutrons. Measured surface temperatures of X-ray-emitting neutron stars have temperatures of  $\sim 10^6$  K. Assuming the opacity is set by Thomson scattering, the pressure at the photosphere (where the optical depth  $\tau = 2/3$ ) is  $\approx g\tau \sim 10^{14} \text{ dyn cm}^{-2}$  and the density is  $\rho \sim 1 \text{ g cm}^{-3}$ . Let us now explore the neutron star, beginning at the photosphere.

### 1.2.1 The Envelope and Crust

Starting from the photosphere, the cm-thick layer where observable emission originates, and moving inwards toward the core of the neutron star, we encounter several distinct layers (Figure 1.2). The envelope is the regime in which the electrons are non-degenerate and the plasma is fluid ( $\rho \lesssim 10^5 \text{ g cm}^{-3}$ ). Because the electrons are non-degenerate, the opacity is high and the temperature has a significant gradient.

Beneath the photosphere, the matter is a fully ionized plasma. To describe the properties of the different regions of the neutron star, we need to begin by describing the properties of the plasma. For simplicity, let us consider a one-component plasma (OCP) with ions with charge number  $Z$ , mass number  $A$ , and number density  $n_N$ . Due to charge neutrality, the electron number density is  $n_e = Zn_N$ . The ratio of the Coulomb energy to the thermal energy of the ions is quantified by the

---

<sup>4</sup> $\rho_0$  is called the *nuclear saturation density*. Technically, it is defined as the density that minimizes the energy density of bulk nuclear matter. Heavy nuclei have densities near  $\rho_0$ .

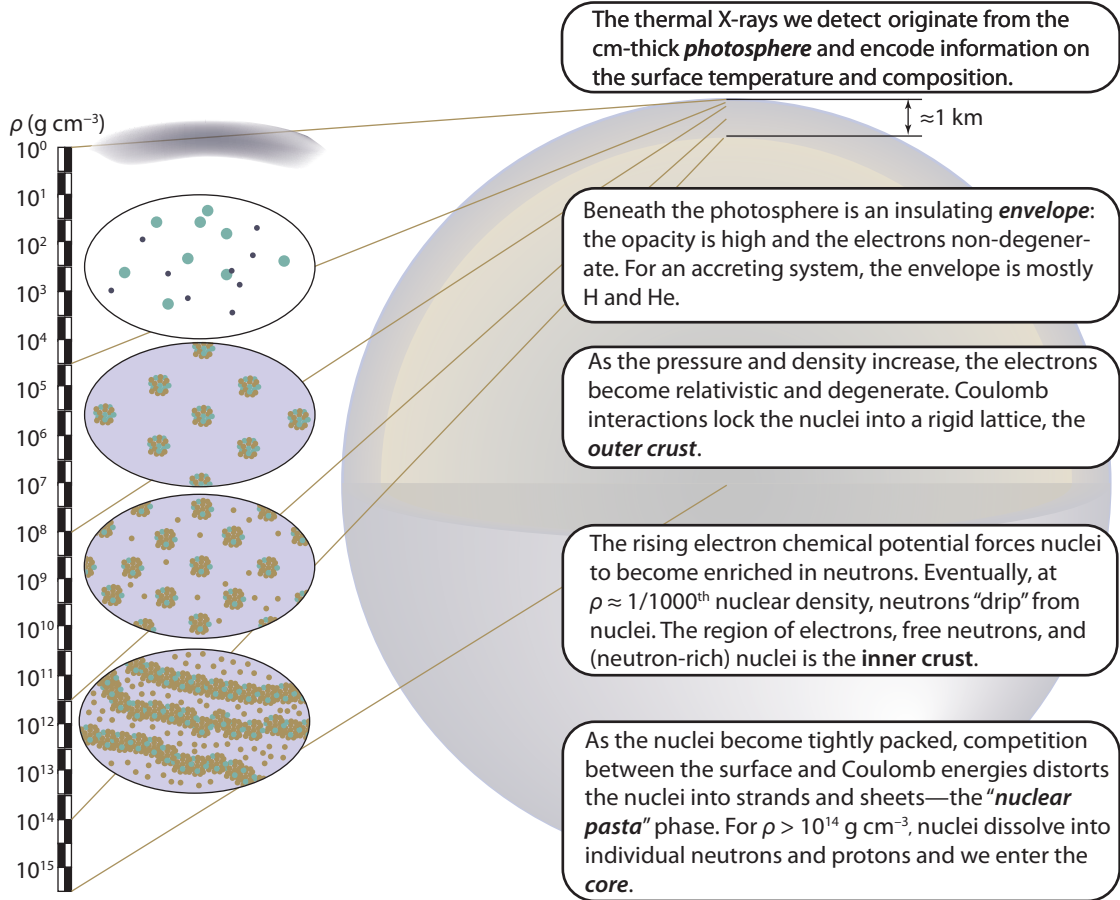


Figure 1.2 The crust of a neutron star. Starting at the photosphere, the density is roughly that of water,  $\approx 1 \text{ g cm}^{-3}$ , but rapidly increases by  $10^{14}$  over the  $\approx \text{km}$  descent to the crust-core boundary.

*plasma coupling parameter,*

$$\Gamma = \frac{Z^2 e^2}{a k_B T}, \quad (1.1)$$

where  $e$  is the elementary charge,  $k_B$  is the Boltzmann constant,  $a = [3/(4\pi n_N)]^{1/3}$  is the ion sphere radius, and  $T$  is the plasma temperature. When  $\Gamma \gtrsim 1$ , ion correlations are important, and the ions form a non-ideal Coulomb liquid. When  $\Gamma > 175$ , the plasma forms a crystal lattice (Hughto et al., 2011). Re-arranging Equation (1.1), we can define the melting temperature  $T_m$ ,

$$T_m = \frac{Z^2 e^2}{175 k_B} \left( \frac{4\pi}{3} \right)^{1/3} \left( \frac{\rho}{A m_u} \right)^{1/3} = 1.3 \times 10^6 \text{ K} \left( \frac{Z^2}{A^{1/3}} \right) \left( \frac{\rho}{10^9 \text{ g cm}^{-3}} \right)^{1/3}. \quad (1.2)$$

Because  $T_m$  increases with density, it is easier for the plasma to crystallize at greater depth in the neutron star.



As we move deeper into the neutron star and the density increases, the electrons become more degenerate. The level of degeneracy is characterized by the Fermi temperature,

$$T_{F,e} = \frac{E_{F,e} - m_e c^2}{k_B}, \quad (1.3)$$

where  $E_{F,e} = m_e c^2 (1 + x_r^2)^{1/2}$  is the electron Fermi energy and  $x_r = p_{F,e}/m_e c = \hbar(3\pi^2 n_e)^{1/3}/m_e c$  is a dimensionless relativity parameter. Where  $T \ll T_{F,e}$ , the electrons are strongly degenerate. Degenerate electrons have high thermal conductivity (see, e.g., Shternin & Yakovlev, 2006), so the temperature becomes nearly isothermal in the deep interior.

Figure 1.3 shows  $T_m$  and  $T_{F,e}$  for a pure  $^{56}\text{Fe}$  plasma along with the temperature-density profile of a neutron star atmosphere with an effective surface temperature  $T_{\text{eff}} = 10^6$  K and consisting of a thin  $^4\text{He}$  layer (column density  $\int \rho dr = 10^4 \text{ g cm}^{-2}$ ) resting on an  $^{56}\text{Fe}$  substrate. Note how the temperature becomes nearly constant where  $\rho \gtrsim 10^5 \text{ g cm}^{-3}$  and the electrons become degenerate. The intermediate regime where the electrons are degenerate but the ions are non-crystallized is sometimes referred to as the *ocean*.

The crust encompasses the regime where the ions form a crystal lattice ( $\rho \gtrsim 10^8 \text{ K}$ ). The free electrons are highly degenerate and relativistic. As we increase the depth, the density increases which increases  $E_{F,e}$ . With higher  $E_{F,e}$ , more neutron-rich compositions are produced via electron capture reactions. Eventually, at a density of about  $4 \times 10^{11} \text{ g cm}^{-3}$ , the composition becomes so neutron-rich that further electron captures induce neutron emissions. The appearance of free neutrons marks the location of neutron-drip and the transition from the outer to the inner crust.

The inner crust is composed of a lattice of neutron-rich nuclei, free, degenerate electrons, and free, degenerate neutrons. As the density increases, the distinction between nuclei and free neutrons becomes ambiguous so groupings of neutrons and protons at the lattice points are sometimes referred to as “nuclear clusters” rather than nuclei (see, e.g., Chamel & Haensel, 2008). The predicted composition and structure of the nuclear clusters strongly depends on theoretical nuclear models.

At the bottom of the crust, where  $\rho \approx \frac{1}{3}\rho_0$  to  $\rho \approx \frac{1}{2}\rho_0$ , the competition between the Coulomb energy and nuclear surface energy of the clusters leads to Coulomb frustration and spherical clusters




Figure 1.3 Fermi temperature  $T_{F,e}$  and melting temperature  $T_m$  as a function of density for a OCP of  $^{56}\text{Fe}$ . The dashed line marks the temperature profile of a neutron star atmosphere comprising a thin  $^4\text{He}$  layer superincumbent on a layer of  $^{56}\text{Fe}$  with  $T_{\text{eff}} = 10^6$  K. Where  $T < T_m$ , the plasma crystallizes. Where  $T \ll T_{F,e}$ , the electrons are strongly degenerate.

become unstable to deformation. Using the compressible liquid-drop model of Mackie & Baym (1977), Ravenhall et al. (1983) found that as the density increased, the most stable shape of the clusters sequentially progressed from spheres to cylinders, slabs, tubes, and then bubbles (see also Hashimoto et al., 1984; Oyamatsu et al., 1984). These different phases have been referred to as “pasta” phases for their similarities to different pastas (rods resemble spaghetti, slabs resemble lasagna, etc.). Not all models, however, predict the existence of pasta phases in the crust (see, e.g., Cheng et al., 1997; Douchin & Haensel, 2000; Maruyama et al., 2005; Oyamatsu & Iida, 2007). The region where the pasta phases may exist (sometimes called the *mantle*) contains the majority of the mass of the crust.

As the density further increases at greater depth, the clusters occupy a greater fraction of the volume until the matter is composed of uniform, nuclear matter. This marks the transition from the crust to the core and occurs at  $\rho \approx 0.5\rho_0$  (Pethick et al., 1995).

### 1.2.2 The Core

At the bottom of the crust, the nuclear clusters blend together forming uniform nuclear matter. This marks the transition from the crust to the core. With a radius of  $\sim 10$  km, the core comprises about 75% of the total volume of the neutron star and over 99% of the mass. The outermost core, where  $0.5\rho_0 \lesssim \rho \lesssim 2.0\rho_0$ , is composed of uniform neutrons, protons, electrons, and—where  $E_{F,e} > m_\mu c^2$ —muons. All of these particles are highly degenerate, and the neutrons and protons may be superfluid. Where  $\rho \gtrsim 2.0\rho_0$ , the composition and EOS are unknown and model-dependent. Some hypothesized compositions include hyperons, a pion condensate, or deconfined quark matter. The central density of the neutron star depends on its total mass (more massive neutron stars have greater central densities), but may be up to  $\sim 20\rho_0$ .

### 1.3 Investigating the Crust Properties of LMXBs in Quiescence

In the absence of accretion, a neutron star’s crust is in thermal equilibrium<sup>5</sup> with the core. As described in Chapter 2, accretion of matter onto the neutron star’s surface compresses the matter beneath it and induces nuclear reactions in the envelope and in the crust. These reactions heat the crust out of thermal equilibrium with the core. When the system enters quiescence, the hot crust cools back into thermal equilibrium with the core. The observed temperature from surface emission following an outburst is called a *cooling curve*.

In the following chapters, I will describe the work I have done investigating the properties of the crust of accreting neutron stars.

The accreting neutron star KS 1731–260, which had been persistently accreting since its discovery in 1989 (Sunyaev, 1989), went into quiescence in 2001 (Wijnands et al., 2001). When KS 1731–260 faded into quiescence, its cool quiescent surface effective temperature  $T_{\text{eff}}$  suggested

---

<sup>5</sup>Strictly speaking, since the core gradually cools over time, the system is never truly in thermal equilibrium. However, the cooling timescale of the core is significantly greater than the cooling timescale for the crust, so to good approximation, we can describe the crust as in thermal equilibrium with the core when discussing its thermal properties.

that the source had a long quiescent interval (Wijnands et al., 2001). Rutledge et al. (2002) pointed out that the long accretion outburst would substantially heat the crust out of thermal equilibrium with the core, and that monitoring observations of the quiescent lightcurve could measure the crust cooling timescale and the amount of heat stored in the crust. Ushomirsky & Rutledge (2001) computed the thermal relaxation timescale for quasi-persistent transients and showed that observations of crust thermal relaxation would convey information about the thermal conductivity of the crust. Fits to the lightcurve of KS 1731–260 and another quasi-persistent transient, MXB 1659–29, found that the thermal conductivity is high (Shternin et al., 2007; Brown & Cumming, 2009).

The shape of the cooling curve is determined by the reactions that are induced by accretion and the thermal properties of the crust. In Ch. 2, I will describe `dStar`, the code I use to model the thermal evolution of accreting neutron stars. I will also explain the methodology by which I fit `dStar` models to observed cooling curves. By modeling the thermal evolution of the crust and comparing the predicted surface emission to observed cooling curves, we can test our theories that describe cold, dense matter. While we cannot directly observe the matter in the core or crust of a neutron star, accreting neutron stars provide a window into the crust. In Chapter 3, I fit `dStar` models to eight of the nine neutron stars for which cooling curve data has been obtained. I test whether we can accurately model these neutron stars using the same crust properties or whether different neutron stars require different properties. In Chapter 4, I compare crustal heating models in which free neutrons in the inner crust diffuse freely to models in which the free neutrons are not allowed to diffuse. I test whether observed cooling curves favor one model or the other.

## CHAPTER 2

### COOLING QUIESCENT NEUTRON STARS

When a neutron star is formed, its temperature is high ( $> 10^{11}$  K) and its interior is in thermodynamic equilibrium with respect to all interactions (Burrows & Lattimer, 1986). During the fiery birth, matter in the primordial crust is in nuclear statistical equilibrium. As the crust thermally relaxes over 10–100 yr (Yakovlev & Pethick, 2004), the composition is annealed and reaches an absolute ground state of lowest energy: the *catalyzed* state (see, e.g., Baym et al., 1971; Baym et al., 1971; Negele & Vautherin, 1973). This state is obtained at any given baryon density by minimizing the energy density under the conditions of  $\beta$ -equilibrium and charge neutrality.

In an accreting system, matter from the outer layers of the neutron star’s companion falls onto the neutron star’s surface and becomes part of the envelope. The accreted matter is often assumed to have roughly solar composition with H and He fractions of  $X_{\text{H}} \approx 0.75$  and  $X_{\text{He}} \approx 0.25$ , respectively. The newly accreted material compresses the matter in the envelope and the crust. The increase in pressure brings the matter out of nuclear equilibrium and induces nuclear reactions.

In the envelope, the increased pressure induces thermonuclear reactions. Over a wide range of observed accretion rates, these reactions are subject to a thin-shell thermal instability (Hansen & Van Horn, 1975), which leads to a rapid increase in X-ray emission, observed as a Type I X-ray burst. These bursts can last tens to hundreds of seconds and recur over hours to days. If enough carbon is produced during repeated Type I X-ray bursts, unstable  $^{12}\text{C}$  burning can eventually occur and produce a *superburst*. Superbursts produce about 1000 times the energy of Type I X-ray bursts, last for hours to days, and have recurrence times of years.

Continual accretion compresses these ashes of light element burning. Where the ions form a crystalline lattice ( $\Gamma > 175$ ), the tunneling through the Coulomb barrier is set by the zero-point motion of nuclei about their lattice sites, rather than the thermal motions that would occur in a rarefied plasma. Instead, where electron captures have reduced the nuclear charge  $Z$  sufficiently, a *pycnonuclear*<sup>1</sup> reaction occurs instead. The accretion-induced nuclear reactions in the crust are

---

<sup>1</sup>From the Greek *pycnos* meaning “dense”

restricted primarily to electron captures, neutron emissions, and pycnonuclear reactions. As a result, the crust cannot reach the catalyzed state, and the composition of an accreting neutron star's crust will deviate significantly from that of an isolated neutron star. Over the course of long-term accretion, the whole crust is replaced by the accreted matter. For a fiducial mass accretion rate  $10^{17} \text{ g s}^{-1} \approx 10^{-9} M_{\odot} \text{ yr}^{-1}$ , the total replacement of the crust takes  $\lesssim 10^6 \text{ yr}$ .

## 2.1 The Composition of the Accreted Crust

The composition of the accreted crust was first calculated by Bisnovatyi-Kogan & Chechetkin (1979) and Sato (1979). Haensel & Zdunik (1990) showed that the pressure-induced nuclear reactions during accretion deposit about 1.4 MeV per accreted nucleon deep in the crust (see also Haensel & Zdunik, 2003, 2008; Lau et al., 2018, for more recent calculations with similar estimates). In this section I will describe the crust composition calculations of Haensel & Zdunik (1990).

Let us follow an element of matter as it is buried deeper in the crust and eventually merges into the core during accretion. The composition at the top of the crust will be the ashes of the nuclear processes in the envelope above. Unstable He burning and superbursts primarily produce iron-peak elements. For simplicity I will use the OCP approximation, so we will begin with pure  $^{56}\text{Fe}$  at the top of the outer crust.

With  $Z = 26$ ,  $T = 10^8 \text{ K}$ , and  $\rho > 10^8 \text{ g cm}^{-3}$ ,  $k_{\text{B}}T/E_{\text{F},e} \ll 1$ . Thermal contributions to thermodynamic potentials are negligible, and the equation of state becomes barotropic,  $P = P(\rho)$ . I therefore use the  $T \rightarrow 0 \text{ K}$  limit for computing the mechanical structure of the crust, while including thermal effects in determining transport coefficients. The density profile will be discontinuous where there are changes in the composition. The pressure, however, is always continuous and monotonically increases with depth to satisfy hydrostatic equilibrium, which makes  $P$  a useful coordinate for the depth of a layer.

We model the crust as a lattice of a single species of atomic nucleus embedded in a degenerate electron gas, along with a free neutron gas in the inner crust. To account for the Coulomb interaction between electrons and nuclei, we will use the Wigner-Seitz (WS) approximation. We decompose the crust into spherical cells centered on the nuclei at the lattice points. The radius of a cell is

defined as  $R_{\text{cell}} = (4\pi n_{\text{N}}/3)^{-1/3}$  so that the volume of a WS cell is equal to  $1/n_{\text{N}}$ . Each WS cell is electrically neutral, so the density of electrons is  $n_e = Zn_{\text{N}}$ . We can assume point-like nuclei since the lattice spacing is much larger than the size of nuclei everywhere except for near the bottom of the crust.

### 2.1.1 Outer Crust

The energy density in a layer in the outer crust consists of the energy contributions from the ions, the electrons, and the Coulomb interactions of the lattice:

$$\epsilon_{\text{tot}} = n_{\text{N}}E_{\text{N}}\{A, Z\} + \epsilon_e + \epsilon_{\text{L}}. \quad (2.1)$$

Here  $n_{\text{N}}$  is the number density of nuclei,  $E_{\text{N}}\{A, Z\}$  is the energy of a nucleus, including the rest mass, with  $A$  total nucleons and  $Z$  protons,  $\epsilon_e$  is the energy density of electrons including the electron rest mass, and  $\epsilon_{\text{L}}$  is the lattice energy density due to electron-electron, electron-ion, and ion-ion Coulomb interactions.  $E_{\text{N}}\{A, Z\}$  has been experimentally measured for over 3000 nuclides (see Wang et al., 2021, for the most up-to-date Atomic Mass Evaluation). For nuclides whose masses have not been measured experimentally,  $E_{\text{N}}\{A, Z\}$  must be calculated using a theoretical model, for example the finite-range droplet model (Möller et al., 2016).

In the conditions of the crust, electron-charge-screening effects are negligible so we can treat the electrons as an ideal Fermi gas with uniform density uniform (Watanabe & Iida, 2003, see, e.g.,).

In the WS approximation, the lattice energy density is the Coulomb energy of one WS cell times the density of nuclei,



$$\epsilon_{\text{L}} = -C \left( \frac{4\pi}{3} \right)^{1/3} Z^{2/3} e^2 n_e^{4/3}. \quad (2.2)$$

The factor  $C$  is determined by the lattice configuration. If we ignore the interactions between WS cells, and only consider the Coulomb energy contribution from each cell,  $C = 9/10$ . For a body-centered cubic (bcc) lattice, which results in the lowest lattice energy,  $C$  decreases to 0.89593 (see, e.g., Chamel & Haensel, 2008).


For a given pressure,  $P$ , the equilibrium composition is determined by minimizing the Gibbs free energy per nucleon of a WS cell,  $g(P) = (\epsilon_{\text{tot}} + P)/n_{\text{b}}$ , where  $n_{\text{b}}$  is the density of baryons. In the outer crust, we are restricted to electron capture reactions, so we generally can't reach the absolute minimum of  $g(P)$ .

The total pressure is made of contributions from the electrons and the lattice,  $P = P_e + P_L$ . The lattice pressure  $P_L = \frac{1}{3}\epsilon_L$  and the electron pressure  $P_e = \mu_e n_e - \epsilon_e$ , where  $\mu_e$  is the electron chemical potential. Since the electrons are degenerate,  $\mu_e = E_{\text{F},e}$ .

As we move to greater depths,  $\mu_e$  and  $\epsilon_e$  increase because  $n_e$  increases. As a result, electron capture reactions will occur when they result in lower  $g(P)$ . Due to nuclear pairing effects, the most stable nuclides tend to have even  $N$  and  $Z$  (even-even nuclides) and the electron captures proceed in pairs. So for our initial composition of  $^{56}\text{Fe}$ , we have



More generally, for some composition with  $A$  nucleons and  $Z$  protons, we have



The first capture (Equation (2.3)) will proceed once  $\mu_e > E\{A, Z - 1\} - E\{A, Z\}$  in a quasi-equilibrium manner with negligible energy release, producing an odd-odd nucleus which is strongly unstable. The second capture (Equation (2.4)) will proceed in a non-equilibrium manner, releasing some energy  $Q_j$  where  $j$  is a label for the non-equilibrium reaction. Because electron capture reactions convert a proton to a neutron,  $A$  is constant through the outer crust.

Electron capture reactions increase the neutron richness of the composition until it reaches neutron drip, where the separation energy of neutrons is negative. At this point, electron captures



will trigger neutron emissions:

$$(A, Z) + e^- \rightarrow (A, Z - 1) + \nu_e \quad (2.5)$$

$$(A, Z - 1) + e^- \rightarrow (A - k, Z - 2) + k n + \nu_e + Q_j. \quad (2.6)$$

Here  $k$  is the number of neutrons emitted. Neutron drip occurs at at the neutron drip density,  $\rho_{\text{ND}} \approx 4 \times 10^{11} \text{ g cm}^{-3}$ . This marks the transition to the inner crust, where a fraction of the neutrons form a neutron gas outside the nuclei.

### 2.1.2 Inner Crust

The conditions of the inner crust are not producible in laboratory experiments due to the presence of the neutron gas. As a result, we cannot use experimental data to determine the energy contribution of the free neutrons nor can we use nuclear mass measurements to determine the energy contribution from nuclei. In fact, we expect that in the deepest layers of the inner crust the nuclei will deform into non-spherical shapes, with the nuclei in adjacent lattice points merging into rod-like or sheet-like phases. At this point, distinct “nuclei” can no longer be defined; rather, the system is divided into two regimes: nuclear clusters of high baryon density that include both protons and neutrons, and a lower density neutron gas outside the clusters. For the sake of clarity and consistency, when discussing a system where free neutrons are present (as is the case throughout the inner crust), I will use the term “nuclear cluster” instead of “nucleus” to refer to the higher density regions.

To determine the energy density of the nuclear clusters and the free neutrons, we must rely on theoretical models. These models may include: microscopic models of the nucleon-nucleon interactions such as the Hartree-Fock approximation; semi-classical models such as the Thomas-Fermi approximation; and classical models such as the compressible liquid drop model (CLDM).

The general process by which we will determine the composition will be similar to how we determined the composition in the outer crust in that we will minimize  $g(P)$  in a WS cell under the condition of charge neutrality. Compositional changes are restricted to electron captures, neutron emissions, or pycnonuclear reactions.

To demonstrate how we can calculate the nuclear energy contributions in a WS cell I will use the CLDM (Mackie & Baym 1977; see Steiner 2012 for a more recent treatment), because it is intuitive and illustrates some of the significant physical properties of the system as it approaches  $\rho_0$ . We treat the nuclear clusters and liquid drops of nuclear matter. In a WS cell centered on a nuclear cluster, the densities of protons and neutrons inside the cluster are constant ( $n_{\text{pi}}$  and  $n_{\text{ni}}$  respectively). We will initially assume spherical clusters with proton radius  $r_p$  such that the number of protons in the cluster is  $r_p = (4\pi n_{\text{pi}}/3Z)^{-1/3}$ . The neutron density outside the clusters  $n_{\text{ni}}$  is constant and less than the density inside the clusters,  $n_{\text{no}} < n_{\text{ni}}$ . We assume there are no protons outside the clusters. The energy density of a WS cell can be expressed as

$$\epsilon_{\text{cell}} = w\epsilon_{\text{bulk,i}} + (1 - w)\epsilon_{\text{bulk,o}} + E_{\text{surf}}/V_c + E_{\text{Coul}}/V_c + \epsilon_e. \quad (2.7)$$

Here  $\epsilon_{\text{bulk,i}}$  and  $\epsilon_{\text{bulk,o}}$  are the energy densities of bulk nuclear matter inside and outside the nuclear cluster respectively.  $E_{\text{surf}}$  is the surface energy of the cluster.  $E_{\text{Coul}}$  is the Coulomb energy of the cell.  $V_c$  is the volume of the cell and  $w = (r_p/R_{\text{cell}})^3$  is the fraction of the cell's volume that is occupied by the nuclear cluster.

The bulk energy density of nuclear matter  $\epsilon_{\text{bulk}}(n_p, n_n)$  depends on the densities of protons and neutrons and can be calculated using a microscopic model such as a Skyrme model (Skyrme, 1958). The bulk nuclear energy terms are then  $\epsilon_{\text{bulk,i}} = \epsilon_{\text{bulk}}(n_{\text{pi}}, n_{\text{ni}})$  and  $\epsilon_{\text{bulk,o}} = \epsilon_{\text{bulk}}(0, n_{\text{ni}})$ .



For the neutron-rich nuclear clusters we expect in the crust, excess neutrons will form a “neutron skin” on the surface of the cluster. The surface term  $E_{\text{surf}}$  accounts for the energy contribution of this skin. Similar to  $\epsilon_{\text{bulk}}(n_p, n_n)$ ,  $E_{\text{surf}}$  can be calculated using a microscopic nuclear model (see, e.g., Douchin et al., 2000). For consistency,  $\epsilon_{\text{bulk}}(n_p, n_n)$  and  $E_{\text{surf}}$  should be calculated using the same microscopic model.

We calculate the Coulomb energy density slightly differently than we did for  $\epsilon_L$  in the outer crust. First, we must now include the Coulomb contribution of protons in the nuclear cluster (the nuclear energy  $E_N\{A, Z\}$  includes the Coulomb contribution of the protons within the nucleus so we did not have to account for it in Equation (2.2)). Second, at the high densities of the inner crust, we must account for the finite size of the nuclear clusters.

Given Equation (2.7), we can determine the composition in a manner similar to the way in which we did for the outer crust. We begin at the top of the inner crust with the composition that resulted from neutron drip. As we increase the pressure (increase depth), neutron emissions and electron captures occur when they result in a decrease in  $g(P)$ . If both neutron emission and electron capture are energetically favorable, then we assume neutron emission occurs first because it is more rapid. When electron captures occur, they generally trigger a neutron emission.

Electron captures systematically reduce the number of protons within the nuclear clusters which reduces the Coulomb barrier between adjacent clusters in the lattice. This, combined with the decreased spacing between lattice points with increased depth, increases the likelihood of pycnonuclear fusion reactions in which nuclear clusters in adjacent lattice points fuse via quantum mechanical tunneling through the Coulomb barrier.

Pycnonuclear fusion reaction rates are highly uncertain and as a result, there is high uncertainty in where they begin to occur in the crust. The pycnonuclear fusion timescale depends sensitively on the number of protons in the nuclear clusters. Pycnonuclear fusion begins to occur when this timescale is less than the accretion timescale,  $y/\dot{M}$ , which describes the time it takes for a layer in the crust to be replaced via accretion, where  $y = \int \rho dr \approx P/g$  is the column density. Because the pycnonuclear fusion timescale is so sensitive to  $Z$ , it decreases significantly when electron captures on the nuclear clusters occur which can induce pycnonuclear fusion. So following the reactions in Equations (2.5) and (2.6), we have a pycnonuclear fusion reaction possibly followed by neutron emission and other non-equilibrium processes:



The reaction in Equation (2.10) represents any nonequilibrium processes that follow the neutron emissions after pycnonuclear fusion. The total energy release resulting from pycnonuclear fusion then is  $Q_j = Q_{j,1} + Q_{j,2} + Q_{j,3}$ . Pycnonuclear reactions account for the majority of the heat release

in the inner crust during accretion.

## 2.2 Observations of Accreting Neutron Stars

After KS 1731–260 was first observed to go into quiescence following an outburst lasting 12 years, Rutledge et al. (2002) proposed that observations of the quiescent cooling neutron star transients could provide insight into the properties of the neutron star’s crust. Motivated by this, Cackett et al. (2006) began monitoring the evolution of the quiescent flux from these quasi-persistent transients. A few years later, detailed thermal evolution models were fit to the quiescent luminosities of KS 1731–260 (Shternin et al., 2007; Brown & Cumming, 2009) and MXB 1659–29 (Brown & Cumming, 2009). Since then, several other LMXBs cooling in quiescence have been observed and described with thermal evolution models (e.g. Degenaar & Wijnands, 2011; Degenaar et al., 2014; Ootes et al., 2016; Parikh et al., 2017a, 2019).

I will summarize here the observational signatures during an accretion outburst and the observed cooling during quiescence. For more detailed descriptions, see the review article by Wijnands et al. (2017). During an accretion outburst, emission from the accretion disk dominates and we are unable to observe the surface of the neutron star. When a nucleon falls onto the neutron star’s surface, it releases  $\eta m_{\text{u}} c^2 \approx GMm_{\text{u}}/R \approx 200 \text{ MeV}$  of gravitational potential energy, most of which is emitted in the form of X-rays. As a result, we can estimate the accretion rate from the X-ray flux via

$$\dot{M}(t) = \frac{\xi 4\pi d^2}{\eta c^2} F(t), \quad (2.11)$$

where  $d$  is the distance to the source,  $\xi$  is a bolometric correction factor, and  $F(t)$  is the observed X-ray flux.

During outburst, sources typically exhibit luminosities of  $L > 10^{36-39} \text{ erg s}^{-1}$  (Wijnands et al., 2017). The X-ray fluxes of many outbursts have been measured by all-sky monitoring telescopes such as the *All-Sky Monitor (ASM)* on the *Rossi X-ray Timing Explorer (RXTE)* and the *Monitor of All-Sky X-ray Image (MAXI)* aboard the International Space Station. By regularly observing X-ray sources via all-sky surveys, these instruments can obtain nearly daily measurements of outburst X-ray fluxes. When an LMXB’s X-ray flux drops below the detectable threshold of the

monitoring telescopes, the source is considered quiescent. Follow-up observations with more sensitive telescopes such as *Chandra* or *XMM-Newton* are necessary to study the quiescent X-ray emission in detail.

Quiescent X-ray spectra can be described by a combination of a soft component (dominating below 1 keV) and a hard component (dominating above 3 keV). The soft component likely originates from the surface of the neutron star and could be due to the cooling emission from a hot neutron star that has been heated due to accretion. It is typically modeled using a neutron star atmosphere model. The hard component is typically modeled using a simple power-law. The source of the hard component is not well understood.

It is usually assumed that when the quiescent spectra are dominated by the soft component that we are observing surface emission a hot neutron star. When the power-law component contributes significantly (more than several tens of percent of the observed quiescent luminosity), then the inferred surface temperatures and quiescent luminosities should be regarded as upper limits.

### **2.3 Modelling the Thermal Evolution of the Crust**

We model the thermal evolution of the neutron star interior using *dStar* (Brown, 2015), which uses a method-of-lines algorithm (Schiesser, 1991) to solve the general relativistic heat transport equations under spherical symmetry in the crust. We use the HELM EOS (Timmes & Swesty, 2000) to compute the electron equation of state; for the strongly coupled ions, we use the results of Chabrier & Potekhin (1998) for the liquid state with corrections from Potekhin & Chabrier (2000) and the treatment of Baiko et al. (2001) for the bcc lattice energy. Anharmonic corrections are handled following Potekhin & Chabrier (2010). The electron-ion EOS routines used in *dStar* are based on those available from the Ioffe Institute<sup>2</sup>. We use a compressible liquid-drop treatment (Mackie & Baym, 1977) for the free neutrons. For the thermal conductivity, we use the treatment of Potekhin et al. (1997) for electron-electron scattering, and that of Baiko et al. (1998) for electron-ion scattering, with impurity corrections from Itoh & Kohyama (1993).

To reduce computational time, *dStar* does not calculate a reaction network and therefore does

---

<sup>2</sup><http://www.ioffe.ru/astro/EIP/>

not directly calculate the heat generated during accretion. Instead, the heating rate and the location of the heat deposition are parameters of the model. For example, we can implement the heating calculated by Haensel & Zdunik (1990) with an inner crustal heating rate of 1.4 MeV per accreted nucleon, deposited over a density range of  $9.1 \times 10^{11} \text{ g cm}^{-3} < \rho < 1.1 \times 10^{13} \text{ g cm}^{-3}$ .

The model parameters we fit to quiescent cooling data are the neutron star core temperature, the impurity of the crust composition, and the accretion-induced heating rates in the shallow layers of the crust (“shallow heating”) and deep in the inner crust (“inner heating”).

The core temperature ( $T_c$ ) sets the crust’s baseline temperature. Computationally, it sets the boundary condition for the crust temperature at the crust-core interface. Constraining its value also provides insight into the structure of the core. In the absence of accretion, a neutron star’s core gradually cools due to neutrino emission. The neutrino emissivity depends on the composition of the core and what neutrino processes dominate (see Yakovlev et al., 2001, for a thorough overview of neutrino emission from neutron stars). The most efficient neutrino cooling process is the direct Urca reactions,

$$p + e \rightarrow n + \nu_e, \quad n \rightarrow p + e + \bar{\nu}_e. \quad (2.12)$$

These reactions maintain  $\beta$ -equilibrium, but each reaction emits a neutrino that carries away energy. In order to satisfy momentum conservation, the direct Urca process can only occur in matter where the proton fraction exceeds 1/9 (see, e.g. Yakovlev et al., 2001). If the proton fraction is not large enough, a spectator nucleon is needed to satisfy momentum conservation. This is the modified Urca process,

$$p + n + e \rightarrow n + n + \nu_e, \quad n + n \rightarrow p + n + e + \bar{\nu}_e \quad (2.13)$$

$$p + p + e \rightarrow n + p + \nu_e, \quad n + p \rightarrow p + p + e + \bar{\nu}_e. \quad (2.14)$$

Because of the spectator particle, the integration over phase space for the modified Urca process has an additional factor of  $(k_B T / E_F)^2$  and proceeds much more slowly than the direct Urca process. For typical conditions in a neutron star core, the direct Urca process is more efficient than the modified Urca process by roughly 6 orders of magnitude (see, e.g., Potekhin et al., 2015).

The rate at which the core heats or cools is also set by its specific heat. Cumming et al. (2017) placed a lower limit of the heat capacities of KS 1731–260, MXB 1659–29, and XTE J1701–462 based on their core temperatures after observed accretion outbursts and ruled out cores dominated by a quark color-flavor-locked phase. Brown et al. (2018) found that the inferred low core temperature of MXB 1659–29 suggests that a rapid neutrino emissivity, such as the direct Urca process, is operating in the core. This is not the case for KS 1731–260, which appears to cool primarily via modified Urca processes.

Electrical and thermal conductivity in the inner crust are primarily determined by impurity scattering, and are inversely proportional to the impurity of the composition,  $Q_{\text{imp}} \equiv n_{\text{ion}}^{-1} \sum_i n_i (Z_i - \langle Z \rangle)^2$  (Itoh & Kohyama, 1993). For dStar models,  $Q_{\text{imp}}$  therefore functions as a proxy for thermal conductivity in the inner crust. The impurity at the top of the crust can vary widely depending on the composition resulting from thermonuclear reactions in the envelope. Lau et al. (2018) found a range of  $Q_{\text{imp}} \approx 4$  for superburst ashes up to  $Q_{\text{imp}} \approx 80$  for extreme X-ray burst ashes, which result from systems with high accretion rate and low metallicity (Schatz et al., 2001). However, in the inner crust, they find that the impurity decreases for all compositions, with the most extreme case yielding  $Q_{\text{imp}} \approx 10$ . By constraining  $Q_{\text{imp}}$ , we can test the validity of crust composition calculations and, in principle, may be able to determine the nuclear processes in the envelope that set the composition of the crust.

The depth and amount of accretion-induced heating in the crust depend on the particular reactions that occur throughout the crust. Modelling the thermal evolution using a detailed heating profile by including heating zones for every exothermic nuclear reaction that occurs is computationally expensive. To increase the computational efficiency, I divide the heating in the crust into three broad regimes: outer heating, inner heating ( $Q_{\text{in}}$ ), and shallow heating ( $Q_{\text{sh}}$ ). These parameters are all implemented as heating rates per accreted nucleon such that a specified amount of heat is deposited uniformly within a specified pressure range for each accreted nucleon. For example, the outer heating, which parameterizes the heat produced from electron capture reactions in the outer crust, has a heating rate of 0.3 MeV per nucleon which is deposited over a region where the pressure

is  $26.86 \leq \log[P/(\text{dyn cm}^{-2})] \leq 30.13$  (Haensel & Zdunik, 1990). In this density regime, the reactions are well-constrained by experimental data, so I use fixed values for the outer heating rate and depth.

The inner heating parameterizes the heat deposited due to reactions in the inner crust, which primarily comes from pycnonuclear reactions (described in Section 2.1.2). This heating is estimated to be  $1.5\text{--}2.0 \text{ MeV u}^{-1}$  deposited over a range of pressures  $30.42 \leq \log[P/(\text{dyn cm}^{-2})] \leq 31.20$  (Haensel & Zdunik, 1990, 2003, 2008; Lau et al., 2018). In this regime, the degenerate, free neutrons drive the composition to extremely neutron-rich nuclides. We have limited data this far from  $\beta$ -stability, so we depend on theoretical models to calculate the inner heating rate. For the calculations in this chapter, rather than fixing inner heating parameters to values from theoretical calculations (e.g. Haensel & Zdunik, 1990, 2003, 2008) or reaction network calculations (e.g. Lau et al., 2018), I leave the heating rate as a free parameter. I do this to determine how well the observational data constrains the heating rate in the inner crust. If the heating rate is well-constrained, we may be able to select between different theoretical models.

Accurately fitting models to the early observed temperatures of several sources requires an additional heat source at shallow depths in the crust (see, e.g., Brown & Cumming, 2009; Degenaar et al., 2011). The heat required for most sources is up to 4 MeV per accreted nucleon and occurs at depths where the density is  $\rho < 10^{10} \text{ g cm}^{-3}$  (see, e.g., Brown & Cumming, 2009; Merritt et al., 2016; Parikh et al., 2019; Degenaar et al., 2014). This density is much less than the density required for pycnonuclear reactions (Haensel & Zdunik, 2008; Horowitz et al., 2008) and the energy release is much greater than what is produced by electron capture reactions. Current nuclear physics models are unable to explain the shallow heating and it is currently unknown what other physical mechanism(s) may be responsible. Inogamov & Sunyaev (2010) proposed that turbulent braking induced by accretion would significantly heat the envelope, but the temperatures they predicted are significantly higher than is observed (see also Inogamov & Sunyaev, 1999). Medin & Cumming (2015) found that including compositionally driven convection in the neutron star ocean lessens but does not entirely eliminate the need for an extra source of shallow heat (see also Medin &



Cumming, 2011; Medin & Cumming, 2014).

The source of shallow heating is currently one of the biggest open questions in the study of neutron star LMXBs. Understanding shallow heating is important for interpreting observed cooling curves which may provide deeper insight into the physics of dense nuclear matter. It may also be an important component in understanding superbursts. The apparent ignition properties of observed superbursts require significantly higher crust temperatures than can be explained by nuclear heating models without additional shallow heating (Altamirano et al., 2012).

## 2.4 Fitting Model Parameters to Observational Data

Estimating the surface effective temperature,  $T_{\text{eff}}$ , of an LMXB in quiescence requires fitting the soft X-ray emission to a spectral model—which requires a model of the neutron star’s atmosphere and assumptions about the neutron star’s mass, radius, and distance. When I fit my dStar models to observations, I do not estimate  $T_{\text{eff}}$  from the X-ray data; rather, I compare my dStar estimates of  $T_{\text{eff}}$  to the estimates of  $T_{\text{eff}}$  based on X-ray data reported in the literature. The masses and radii of the observed neutron stars are unknown and the distance estimates often have large uncertainties. Because the reported estimates of  $T_{\text{eff}}$  depend on assumptions for these parameters as well as the choice of atmosphere model, our dStar model fits are subject to systematic uncertainties.

To fit our models to the temperature data, I use `emcee` (Foreman-Mackey et al., 2013), a Python implementation of the Markov Chain Monte Carlo (MCMC) ensemble sampler proposed by Goodman & Weare (2010). MCMC methods are a class of algorithms for sampling from a probability distribution. They are useful for estimating model parameters, particularly in high-dimensional parameter spaces and when the shape of the probability distribution is not known a priori. For a detailed description and discussion of MCMC methods see a review such as Mackay (2003), Hogg et al. (2010), or Sharma (2017). Here I will give a high-level, conceptual description of the MCMC methods I use in my analysis.

Our goal with our model is to determine the model parameters that yield results that most closely fit the observed data. However, since the model is not a perfect description of the physical system we are studying and because the observational data has some uncertainty to it, there is a non-zero

probability for a large range of parameter sets to yield the data. We therefore want to know the probability distribution of parameter sets given the observed data. Let us consider a model with  $M$  parameters, described by the parameter set  $\Theta = \{\theta_1, \theta_2, \dots, \theta_M\}$ . We can calculate the probability distribution of  $\Theta$  given a set of data  $\mathbf{D}$  using Bayes' Theorem,

$$P(\Theta|\mathbf{D}) = \frac{P(\mathbf{D}|\Theta)P(\Theta)}{P(\mathbf{D})} \quad \text{or} \quad \mathcal{P}(\Theta) = \frac{\mathcal{L}(\theta)\pi(\Theta)}{\mathcal{Z}}. \quad (2.15)$$

The *likelihood*  $\mathcal{L}(\theta) \equiv P(\mathbf{D}|\Theta)$  is the probability that our model will yield  $\mathbf{D}$  given  $\Theta$ . The *prior*  $\pi(\Theta) \equiv P(\Theta)$  describes our prior knowledge of what we expect the parameters may be. The *evidence*  $\mathcal{Z} \equiv P(\mathbf{D}) = \int \mathcal{L}(\theta)\pi(\Theta)d\Theta$  is a normalizing factor that quantifies the overall probability of the model. Since we are updating our knowledge of  $\Theta$  given a new set of data  $\mathbf{D}$ ,  $\mathcal{P}(\Theta) \equiv P(\Theta|\mathbf{D})$  is called the *posterior*<sup>3</sup>.

In practice, we typically cannot determine  $\mathcal{L}(\theta)$  analytically and so must use numerical methods to approximate the posterior distribution  $\mathcal{P}(\Theta)$ . This could be done using grid-based methods, in which we create a grid in parameter space, calculate the likelihood at each grid point, and use numerical integration techniques to calculate  $\mathcal{Z}$ . Since the posterior must be normalized, however, large regions of the parameter space must yield near-zero probability. Therefore, a significant amount of calculation time will be spent gaining little to no useful information. Unless we know about the shape of the posterior distribution *a priori*, we can't know where in the parameter space we should place more grid points.

This computational inefficiency becomes a significant issue when we have a large number of parameters due to the *curse of dimensionality*. The volume of the parameter space increases exponentially with the number of dimensions; as a result, the number of grid points we need to sample the space also increases exponentially. This will become computationally prohibitive.

MCMC methods are designed to generate samples proportionally to  $\mathcal{P}(\Theta)$  so that we can more optimally estimate the posterior distribution. This is done by creating a chain of  $N$  sets of parameter values  $\{\Theta_1 \rightarrow \Theta_2 \rightarrow \dots \rightarrow \Theta_N\}$  such that the number of elements of the chain that are within a

---

<sup>3</sup>When applying Bayes' Theorem to model-fitting as I am,  $\mathcal{P}(\Theta)$  is a distribution over parameter space. I will interchangeably refer to this as the "posterior", "posterior distribution", or "posterior probability distribution".

region of parameter space is approximately proportional to  $\mathcal{P}(\Theta)$  in that region. More precisely, if  $m(\Theta_i)$  is the number of chain elements within a region  $\delta_{\Theta_i}$  centered on  $\Theta_i$ , then the normalized density of samples in that region is  $\rho(\Theta_i) = m(\Theta_i)/N$  and

$$\int_{\Theta \in \delta_{\Theta_i}} \rho(\Theta) d\Theta \approx \int_{\Theta \in \delta_{\Theta_i}} \mathcal{P}(\Theta) d\Theta. \quad (2.16)$$

For a finite chain length  $N$ , the sample density will only approximate the posterior distribution; in the limit as  $N \rightarrow \infty$ , however,  $\rho(\Theta) \rightarrow \mathcal{P}(\Theta)$  throughout the parameter space. So with a sufficiently large chain, the sample density can tell us the model parameters that best fit the data and how well-constrained those parameters are by the data. We can also use the sample density to calculate expectation values of any parameter-dependent functions.

The most well-known MCMC algorithm is the Metropolis-Hastings algorithm, first developed by Metropolis et al. (1953) for the case of symmetrical proposal distributions and later generalized by Hastings (1970). The Metropolis-Hastings algorithm is a random walk algorithm where the chain of samples is generated by following a “walker” as it steps through the parameter space. We initiate the walker at some point  $\Theta_1$  and have it take steps in parameter space to generate the elements of the chain. We will sample from the function  $f(\Theta) = \mathcal{L}(\theta)\pi(\Theta)$  since it is proportional to  $\mathcal{P}(\Theta)$  (see Equation (2.15)). For any element in the chain  $\Theta_i$ , the walker takes steps in parameter space according to the following steps.

1. Propose a new set of parameters  $\Theta'_{i+1} = \Theta_i + \delta\Theta$  where the step  $\delta\Theta$  is selected randomly from a proposal distribution (for example, via a Gaussian centered on  $\Theta_i$ ).
2. Calculate the acceptance ratio  $\alpha = f(\Theta'_{i+1})/f(\Theta_i)$ .
3. Accept or reject the step:
  - a) If  $\alpha \geq 1$ , accept the step.
  - b) If  $\alpha < 1$ , accept the step with a probability of  $\alpha$  and reject the step with a probability of  $1 - \alpha$ .
4. If the step is accepted,  $\Theta_{i+1} = \Theta'_{i+1}$ ; if rejected,  $\Theta_{i+1} = \Theta_i$ .

This process can occur with multiple walkers, each following the algorithm independently. This is advantageous because it increases the chances of the walkers exploring all of the relevant regions of parameters space. With a small number of walkers, they can get stuck in local minima for long periods of time and we may miss some of the posterior distribution. So even though increasing the number of walkers increases the computational time, it increases the number of samples we produce and allows us to converge more quickly. Because each walker proceeds independently, the algorithm is trivial to parallelize.

Each element of the chain depends only on the previous element, either because the proposed step is accepted and  $\Theta'_{i+1}$  tends to be near  $\Theta_i$ , or because the proposed step is rejected so that  $\Theta_{i+1} = \Theta_i$ . Because of this, samples in the chain are correlated with their predecessors. The farther two elements in the chain are from one another, the less likely they are to be correlated. Correlated samples provide less information about the underlying posterior distribution than uncorrelated samples because they are somewhat redundant (they sample the same region of the parameter space). We therefore want to have a sufficiently long chain to have enough uncorrelated samples to sufficiently approximate the posterior distribution.

We can quantify the amount of correlation for each parameter between pairs of samples that are  $t$  iterations apart with the *auto-covariance*<sup>4</sup>  $C_j(t)$  of the parameter  $\theta_j$ . Let us consider an infinitely long chain  $\{\Theta_1 \rightarrow \dots \rightarrow \Theta_i \rightarrow \dots\}$ , where  $\Theta_i = \{\theta_{i,1}, \dots, \theta_{i,j}, \dots, \theta_{i,M}\}$ , so that  $\rho(\Theta) \rightarrow \pi(\Theta)$ . The auto-covariance for the parameter  $\theta_j$  is

$$C_j(t) = \mathbb{E}[(\theta_{i,j} - \bar{\theta}_j)(\theta_{i+t,j} - \bar{\theta}_j)] = \lim_{N \rightarrow \infty} \frac{1}{N} \sum_{i=1}^N [(\theta_{i,j} - \bar{\theta}_j)(\theta_{i+t,j} - \bar{\theta}_j)], \quad (2.17)$$

where  $\mathbb{E}(X)$  denotes the expectation value of the quantity  $X$  and  $\bar{\theta}_j = \lim_{N \rightarrow \infty} \frac{1}{N} \sum_{i=1}^N \theta_{i,j}$ . For sample separation  $t$  such that  $\theta_{i,j}$  and  $\theta_{i+t,j}$  are completely uncorrelated,  $C_j(t) = 0$ . For  $t = 0$  the auto-covariance is at its maximum for all parameters because all pairs are equal and therefore perfectly correlated.

---

<sup>4</sup>Here I define the auto-covariance for each parameter because this is how it is handled in emcee. Often a single auto-covariance is calculated for samples in the parameter space as  $C(t) = \mathbb{E}[(\Theta_i - \bar{\Theta}) \cdot (\Theta_{i+t} - \bar{\Theta})] = \lim_{n \rightarrow \infty} \frac{1}{n} \sum_{i=1}^n (\Theta_i - \bar{\Theta}) \cdot (\Theta_{i+t} - \bar{\Theta})$

We can normalize the auto-covariance by its maximum value to get the *auto-correlation*  $A_j(t)$ :

$$A_j(t) \equiv \frac{C_j(t)}{C_j(0)}. \quad (2.18)$$

The sum of  $A_j(t)$  for all non-zero lags ( $t \neq 0$ ) gives us the *autocorrelation time*  $\tau_j$ :

$$\tau_j = \sum_{t=-\infty}^{t=\infty} A_j(t) - 1 = 2 \sum_{t=1}^{t=\infty} A_j(t) \quad (2.19)$$

which is the number of steps required to produce samples that have independent values of  $\theta$ . This is a useful quantity for determining when the chain has sufficiently converged. Unfortunately, there is actually no guarantee that a finite-length chain will converge for any but the most simple models. A common practice (recommended, e.g., by Goodman & Weare, 2010) is to have a chain at least the length of the autocorrelation time to ensure that we have independent samples. Different parameters may have different autocorrelation times, so to ensure that samples are independent in all parameters, we would ensure that the length of the chain is greater than the greatest autocorrelation time across all parameters.

Calculating  $\tau_j$  exactly requires a chain of infinite length because we need pairs of samples with separation  $t \rightarrow \infty$  to solve Equation (2.19). Estimating  $\tau_j$  from a finite-length chain is more complicated because if we simply replace the upper limit of the sum in Equation (2.19) with the chain length  $N$ , then we have fewer pairs of samples with large  $t$ , so our estimates of  $A_j(t)$  are very noisy. As a result, we need to have a chain much longer than  $\tau_j$  in order to accurately approximate its value. Foreman-Mackey et al. (2013) find that estimating  $\tau_j$  with emcee requires chains of length  $N \geq 50\tau$ . To ensure that we can accurately estimate  $\tau_j$  for each parameter and that our chain is long enough that it has sufficiently explored the parameter space, Foreman-Mackey et al. recommend a convergence criterion that  $N > 50 \max(\{\tau_1, \dots, \tau_M\})$ .

## 2.5 Accounting for Variable Accretion History

In some cases, the accretion rate during an outburst varies enough that it can have an observable impact on the predicted cooling curve as well as the temperature profile in the crust (see Ootes et al., 2016). This can be particularly important when modeling Type I X-ray bursts or superbursts because the temperature at the top of the crust determines the temperature at the base of the envelope

or ocean where the thermonuclear runaway responsible for the X-ray bursts occurs. Notably, for the quasi-persistent transient KS 1731–260, in the months leading up to its observed superburst in 1996 (Kuulkers et al., 2002), the observed accretion rate was significantly higher than its average accretion rate over the outburst. The system maintained this higher accretion rate for approximately 2 years then the accretion rate subsequently dropped below the average rate for another few years before going into quiescence (see Figure 2.1). In order to accurately calculate its temperature profile at the time of the superburst, we need to account for the increase in the accretion rate leading up to the superburst. We also need to be able to account for the decrease in the accretion rate before quiescence in order to properly fit our crust cooling models to the observed cooling curve. Here I will demonstrate how I account for the variable accretion history of KS 1731–260. The process for any other source will be the same.

In principle, there exists some continuous function for the accretion history  $\dot{M}(t)$ . However, we don't have continuous observations of any quasi-persistent transients throughout the duration of an accretion outburst. Instead we have measurements of X-ray counts during observation periods that were often not taken at regular time intervals. If we convert the observed X-ray photon count rates<sup>5</sup> to accretion rates using Equation (2.11), we obtain an array of measured accretion rates with their corresponding times. I will refer to these as  $\dot{M}_i$  and  $t_i$  for the  $i$ th observed accretion rate and its corresponding time respectively.

The time  $t$  can be measured by the modified Julian date (MJD), which is how observations are usually recorded, or by the time relative to the end of the outburst. For consistency with my cooling models, I set  $t = 0$  at the end of the outburst: accretion occurs when  $t < 0$  and quiescent cooling occurs when  $t > 0$ . To approximate the continuous function  $\dot{M}(t)$ , I linearly interpolate between observed accretion rates.

In dStar, accretion is implemented by simulating accretion epochs with constant accretion rates. An outburst can have an arbitrary number of accretion epochs each with a specified duration. In the simplest accretion model, we have only one accretion epoch with an accretion rate equal to

---

<sup>5</sup>Daily averaged X-ray counts and bolometric correction factor for KS 1731–260 received from Laura Ootes, private communication

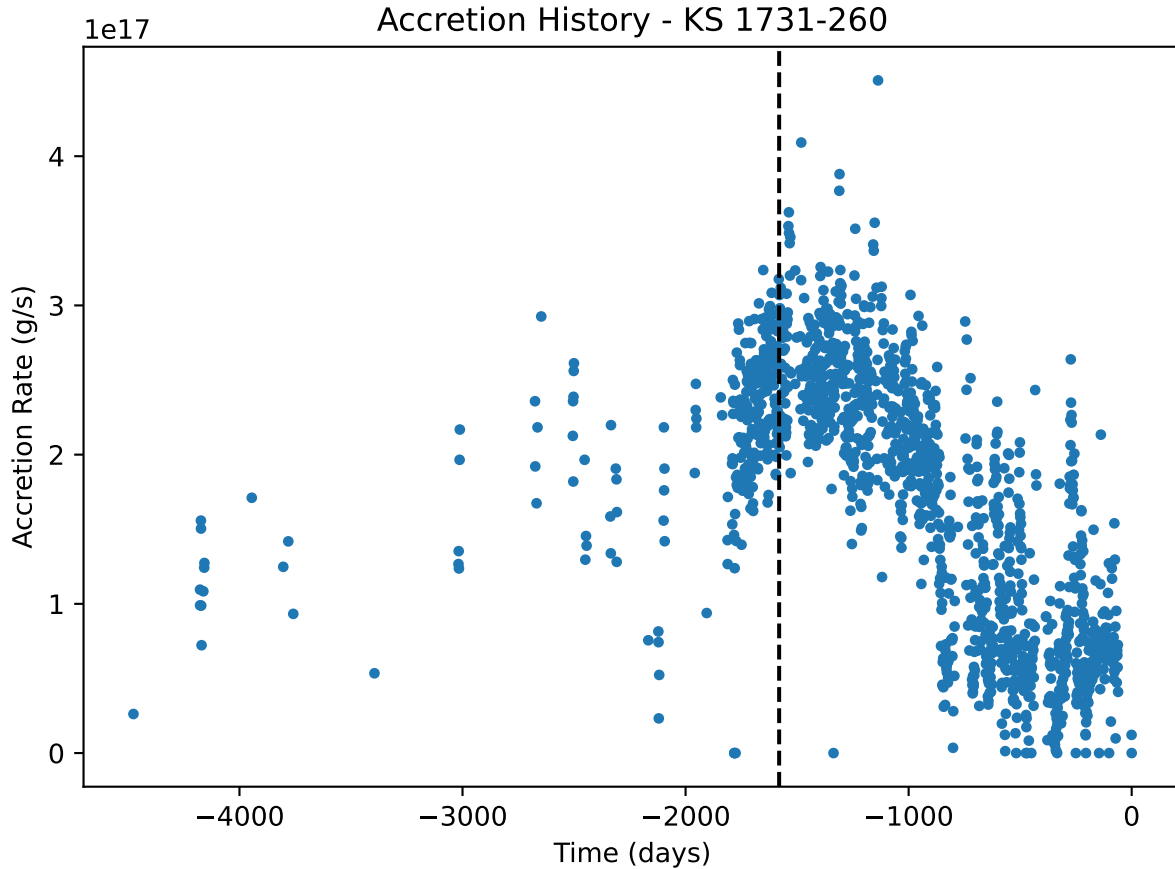


Figure 2.1 Observed accretion rates of KS 1731–260 during its 12.5 year outburst. Accretion rates calculated from X-ray observations taken by *ASM*, *PCA*, *TTM*, and *ART-P*. Distance to KS 1731–260 is assumed to be 7 kpc. The time is measured relative to the end of the outburst. The vertical, dashed line marks the time of the observed superburst.

the average of  $\dot{M}(t)$  over the duration of the outburst. If we want the model to take into account the variability of the accretion rate, we can partition  $\dot{M}(t)$  into accretion epochs where the accretion rate in each epoch is equal to the average of  $\dot{M}(t)$  over its respective time interval.

Increasing the number of accretion epochs increases the computational time of a model. So there is a tradeoff between the accuracy of the model and the computational time. Using an MCMC to fit the model parameters to observational data requires a large number (typically thousands to tens of thousands) of models so reducing the computational time as much as possible while maintaining sufficient accuracy is important.

First, I must determine what is sufficiently accurate for the purposes of the model. My goal is

to fit the crustal cooling model to an observed cooling curve to determine the model parameters described in Section 2.3. Given the same set of model parameters, two models with different accretion history partitioning schemes (i.e, using different numbers of accretion epochs for different accretion time resolution) will produce different cooling curves. Therefore, when fitting the models to the data using an MCMC, the estimated posterior probability distributions of the parameters of the two models will differ. If the difference between the two distributions is negligible (i.e., the best-fit parameters differ by less than the standard deviation of the distributions for each parameter), then the two accretion partitioning schemes are equivalent enough for the purposes of estimating the model parameters.

However, comparing the posterior distribution estimates of many different accretion partitioning schemes is very computationally expensive because it would require running a large number of models for each partitioning scheme. Fortunately, we can take advantage of the fact that the timescale of thermal diffusion to the surface is much less near the surface than at deeper layers in the neutron star (see Brown & Cumming, 2009). When an outburst ends, the heat in the crust diffuses to the surface and to the core. The time for heat to diffuse to the surface increases with the square of the depth in the crust (Henyey & L'Ecuyer, 1969). So the temperature we observe at the surface over time roughly traces the heat produced at different depths in the crust (early temperature observations correspond to the heat produced near the surface; late temperature observations correspond to the heat produced deep in the crust).

This also means that early temperature observations are more sensitive to variations in the accretion rate. For example, if the accretion rate varies on a timescale of 10 days (i.e., the accretion rate increases dramatically for a 10 day period, or the accretion rate fluctuates between a high-accretion rate and a low-accretion rate with 10 day periods), then we would expect the temperature in layers of the crust with  $\tau \lesssim 10$  d (where  $\tau$  is the timescale of diffusion to the surface) to vary with the variation in accretion rate while the temperature in layers with  $\tau \gg 10$  d would stay relatively constant.

Rather than producing a whole posterior distribution for each accretion partitioning scheme,



I focus only on how much the different schemes affect the first temperature observation during quiescence because this corresponds to the layer in the crust whose temperature is most sensitive to variations in the accretion rate.

For KS 1731–260, the first temperature observation during quiescence occurred 64 days after the end of the outburst (Cackett et al., 2006). This means we are largely unable to make constraints on our models that depend on the heat generated in layers with thermal diffusion timescales much less than 64 days. I generate a cooling curve from one high-time-resolution dStar model with  $\dot{M}(t)$  partitioned into 1-day long accretion epochs and use the model parameters found by Merritt et al. (2016). For all other accretion history partitioning scheme, I find the first-point-error which I define as the absolute difference in temperature from the high-resolution model for the first temperature observation:  $\epsilon = |T - T_0|$  where  $T$  and  $T_0$  are the observed surface temperatures of the model being tested and the high-resolution model respectively at  $t = 64$  d.

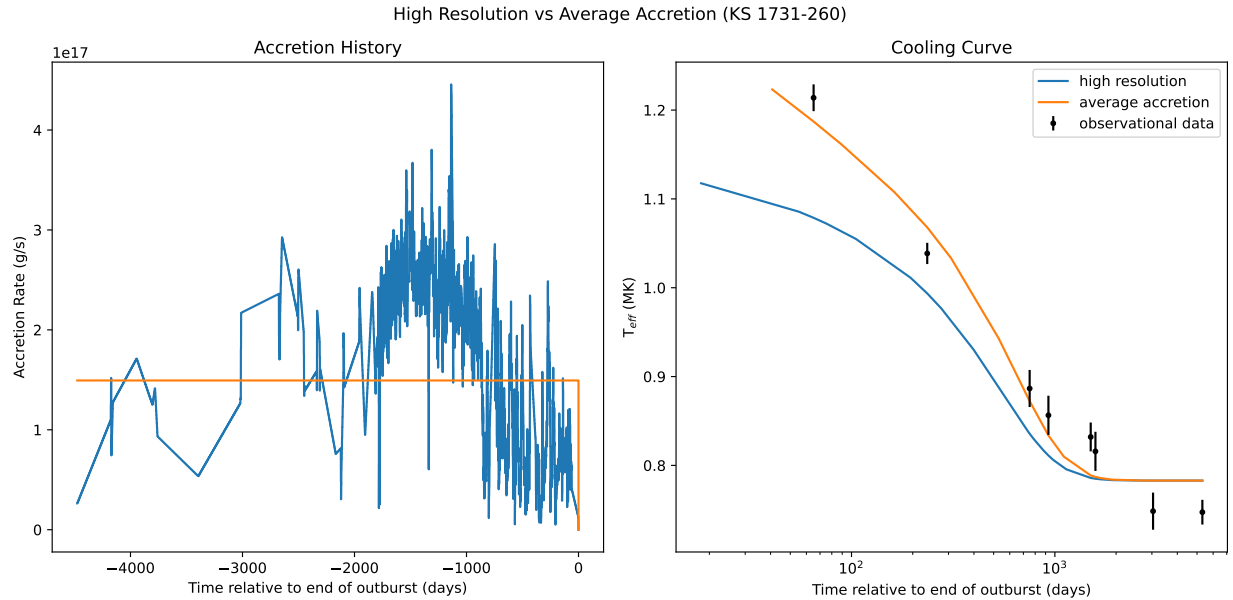


Figure 2.2 Comparison of the high resolution accretion partitioning scheme and the average accretion rate over the duration of the outburst. On the left is the accretion rate as a function of time over the duration of the outburst. On the right are the resulting cooling curves for each partitioning scheme as well as the observed temperatures. The constant accretion rate fits the data better because the dStar model parameters I used came from a fit using the average accretion rate.

The simplest way to partition the accretion history is to split the accretion into  $N$  epochs of

equal duration such that

$$\Delta t_{\text{const}} = \frac{\Delta t_{\text{tot}}}{N} \quad (2.20)$$

where  $\Delta t_{\text{const}}$  is the duration of each epoch and  $\Delta t_{\text{tot}}$  is the total duration of the outburst. By increasing  $N$ , we can increase the time resolution which gives us greater detail in the accretion history, eventually converging on the true accretion rate  $\dot{M}(t)$ . I will refer to this scheme as the *constant* epoch scheme.

However, we don't need to use epochs of equal duration. In fact, intuitively we can reason that we likely want to have higher resolution partitioning near the end of the outburst than earlier in the outburst. This is because the thermal evolution of the crust depends on the thermal diffusion timescales at different depths in the crust. Any layer in the crust will "forget" the accretion on timescales greater than the diffusion timescale at that layer's depth. Since we can only measure the surface temperature after the outburst, we don't need detailed information about the accretion rate at times long before the outburst's end.

I test two different schemes where later accretion epochs are shorter than earlier epochs which I call the *linear* scheme and the *exponential* scheme. In both schemes, the last epoch of the outburst is the shortest. I then go backwards in the outburst (from the end of the outburst to the start of the outburst) and each successive epoch increases in duration relative to its predecessor. The schemes differ in the way in which the epoch duration increases with successive epochs.

In the linear scheme, successive epoch durations grow at a linear rate,

$$\Delta t_{\text{lin},i} = i\Delta t_{\text{lin},1}, \quad (2.21)$$

where  $\Delta t_{\text{lin},1}$  is the duration of the last epoch and  $i$  counts the epochs from the last epoch to the first (so  $\Delta t_{\text{lin},2}$  is the second-to-last epoch,  $\Delta t_{\text{lin},3}$  is the third-to-last-epoch, etc.). The sum of the durations of all epochs must equal the total outburst duration  $\Delta t_{\text{tot}}$ . This gives us the constraint  $\sum_1^N i\Delta t_{\text{lin},1} = \Delta t_{\text{tot}}$ . The duration of the last epoch is thus

$$\Delta t_{\text{lin},1} = \frac{2\Delta t_{\text{tot}}}{N(N+1)}. \quad (2.22)$$

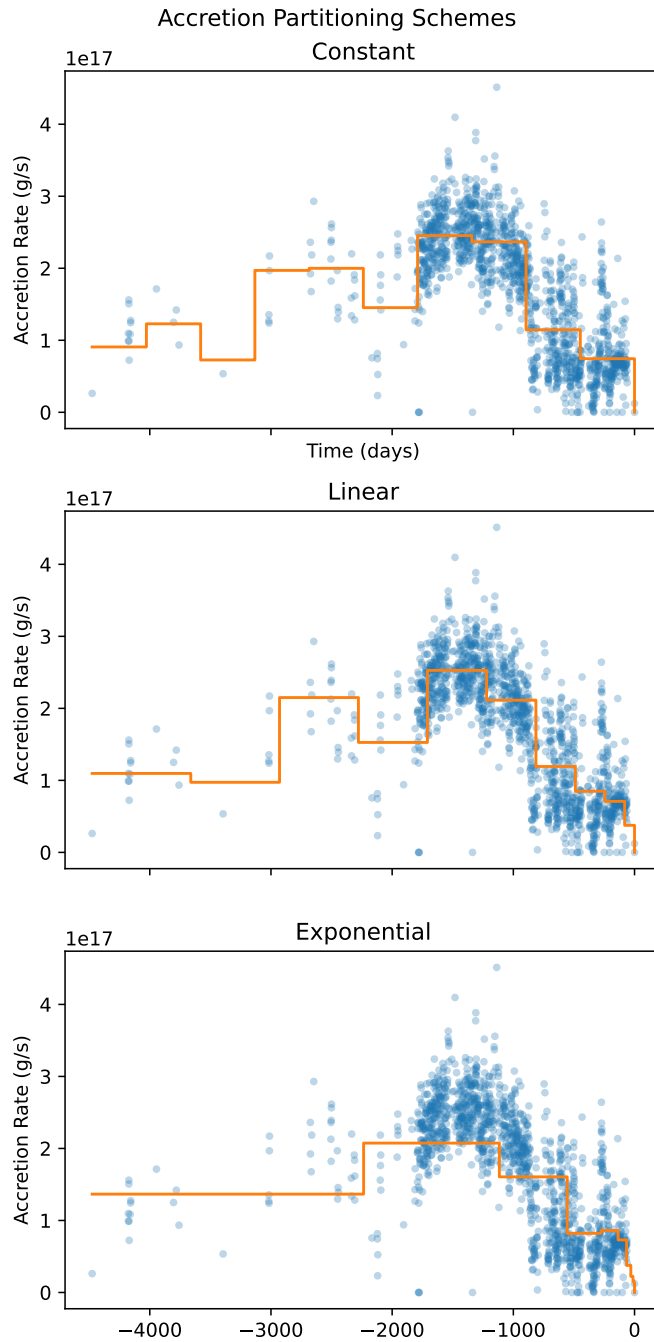


Figure 2.3 Three different partitioning schemes for the accretion rate. The constant scheme has epochs with a constant duration. The linear scheme has epochs that increase in duration linearly when iterating backwards from the last epoch. The exponential scheme has epochs that increase exponentially (with an exponential base of 2 in this example) when iterating backwards from the last epoch. The observed values of the accretion rate are shown in each plot for reference.

In the exponential scheme, successive epoch durations grow at a rate  $r$ :  $\Delta t_i = r\Delta t_{i-1}$ . If the last epoch has a duration  $\Delta t_{\text{exp},1}$ , then

$$\Delta t_{\text{exp},i} = r^{i-1}\Delta t_{\text{exp},1}. \quad (2.23)$$

As with the linear scheme, the sum of the durations of all epochs must equal the total duration of the accretion outburst. This gives us a geometric series:  $\sum_1^N r^{i-1}\Delta t_{\text{exp},1} = \Delta t_{\text{tot}}$ . We can again solve for the duration of the last epoch given  $N$  and  $\Delta t_{\text{tot}}$ :

$$\Delta t_{\text{exp},1} = \frac{(1-r)\Delta t_{\text{tot}}}{1-r^N}. \quad (2.24)$$

The duration of the outburst  $\Delta t_{\text{tot}}$  is an observed value so it is fixed for any partitioning scheme. For any value of  $N$ , I can partition the accretion history using the constant (Equation (2.20)), linear (Equations (2.21) and (2.22)), and exponential (Equations (2.23) and (2.24)) schemes. See Figure 2.3 for the accretion outburst of KS 1731–260 partitioned into 10 epochs with each scheme.

To determine the effectiveness of each scheme, I will compare the first-point error of each scheme as a function of number of epochs. As the number of epochs increases, the error decreases, but the computational time increases. So I will consider the most effective scheme to be the one such that the error falls below a desired threshold with the fewest epochs. The error threshold I use will be the observational uncertainty on the first observed temperature, which for KS 1731–260 is 0.015 MK (Merritt et al., 2016).

Figure 2.4 shows the results of the first-point error tests. In this test, for the exponential scheme, I use an exponential growth rate of  $r = 2.0$  for Equations (2.23) and (2.24). In my high-resolution model, I use 1-day long epochs, so for each partitioning scheme I never use so many epochs that the shortest epoch is less than one day. Note from Equation (2.24), that due to the  $r^N$  in the denominator, the exponential scheme can't have very many epochs before its shortest epoch becomes less than one day. For each scheme, the error decreases with the number of epochs and eventually drops below the error threshold. However, the exponential scheme requires the fewest epochs to drop below the threshold.

Having determined that the exponential scheme yields the best results of the three schemes, I further test the effect of using different growth rates in the range  $1.2 \leq r \leq 3.0$ . Figure 2.5 shows

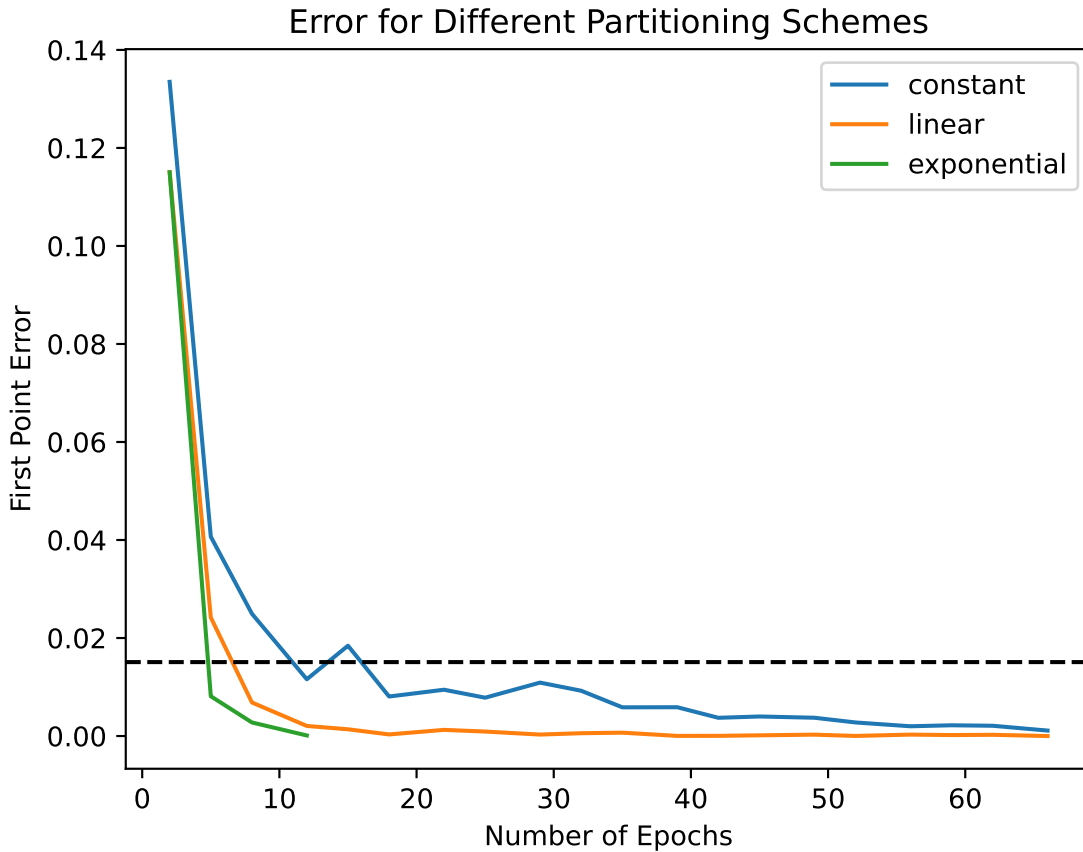


Figure 2.4 Comparison of the first-point error for each partitioning scheme as a function of number of epochs. The first-point error is defined as  $\epsilon = |T - T_0|$  where  $T$  and  $T_0$  are the observed surface temperatures of the model being tested and the high-resolution model respectively. The horizontal dashed line is the error threshold which I set to be equal to the observational uncertainty on the first observed temperature measurement from Merritt et al. (2016).

the results of exponential schemes with  $r = 1.2$ ,  $r = 2.0$ , and  $r = 3.0$ . For a set number of epochs, schemes with larger values of  $r$  will have greater resolution at the end of the outburst and worse resolution at early times in the outburst. When partitioning into only 2–3 epochs, larger values of  $r$  yield lower error because the early quiescent temperature depends more heavily on the accretion history at the end of the outburst. As the number of epochs increases, however, schemes with large  $r$  do not benefit as much as those with small  $r$  because they don't improve the early accretion history resolution as much. I find that  $r = 2.0$  yields error below the error threshold with the fewest epochs and nearly no error with 6–7 epochs.

Given these tests, I conclude that for KS 1731–260, using an exponential partitioning scheme

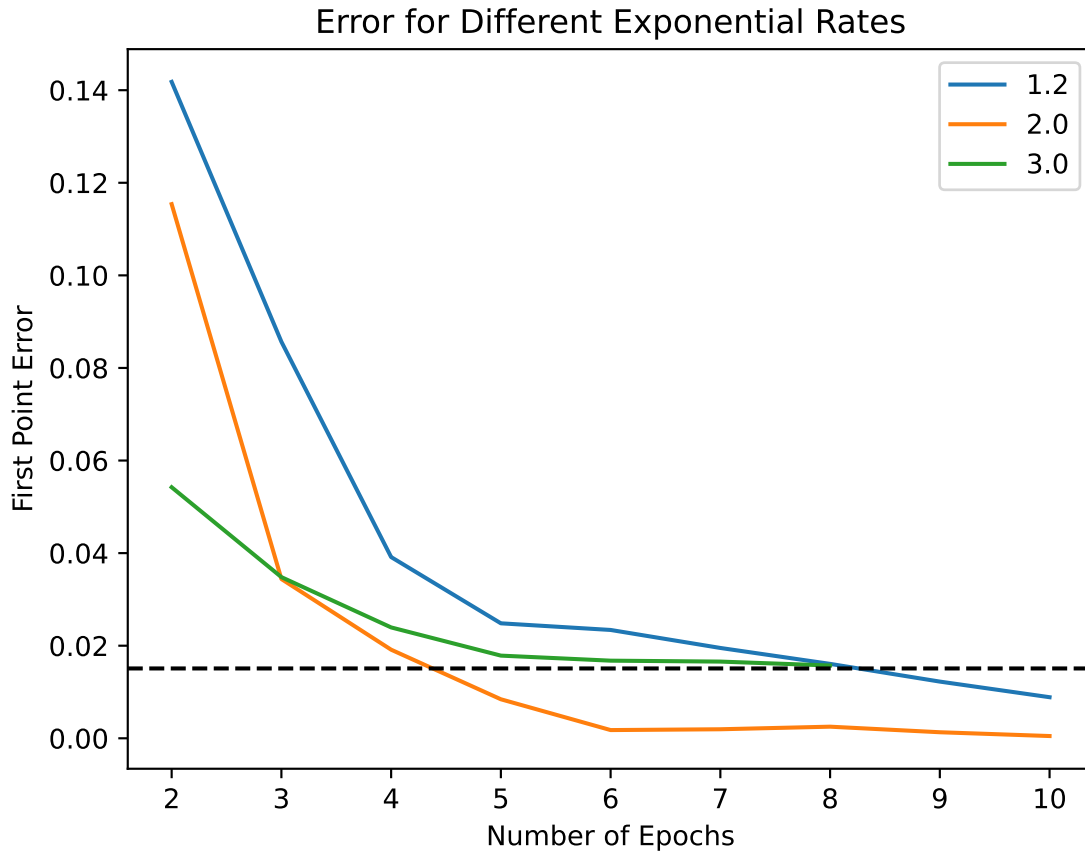


Figure 2.5 Comparison of different growth rates for the exponential partitioning scheme.

for the accretion history with a growth rate of  $r = 2.0$  yields the best results. In all following models of KS 1731–260, I use 6 epochs which yielded a first-point error of less than half of the observational error of the first observed quiescent temperature.

In Chapters 3 and 4, I will use the methodology described here to model crustal cooling of observed LMXBs using dStar.

## CHAPTER 3

### JOINTLY FITTING CRUSTAL PROPERTIES OF MULTIPLE NEUTRON STARS

In Chapter 2, I demonstrated how I can fit dStar models to observed cooling curves to estimate the core temperature, the impurity of the composition, the shallow heating rate, and the inner heating rate of neutron stars in LMXBs.

It is unknown to what extent neutron stars in different LMXBs share crustal properties such as composition and deep crustal heating rates. In Section 2.1, I discussed how accretion induces compositional changes on a layer of matter as it is buried deeper in the crust. The reactions that cause these compositional changes are responsible for the heat produced during an accretion outburst. The reactions and composition may depend significantly on the composition at the top of the crust when accretion occurs. Lau et al. (2018) calculated the reactions, composition, and heat release of several different initial compositions: pure  $^{56}\text{Fe}$ , extreme rp-process ashes (based on the X-ray burst model of Schatz et al., 2001), mixed H/He X-ray burst ashes, and superburst ashes. They found that the deep crustal heating rate does not differ significantly with different initial compositions, but the impurity of the composition in the outer crust does depend strongly on the initial composition.

Additionally, the catalyzed crust composition differs significantly from that of the accreted crust. Because of this, accretion onto a neutron star with a catalyzed crust would induce heating at a different rate than calculated for accreted crusts. As accretion occurs over thousands of years, the catalyzed crust is replaced by the accreted one. At a typical mass accretion rate of  $\dot{M} \sim 10^{17}$  g/s, it would take  $\sim 10^6$  yr to completely replace the crust. During this intermediate period, the neutron star has a partially-accreted crust which would also have different thermal diffusion properties and accretion-induced heating rates from catalyzed crusts and fully accreted crusts (Wijnands et al., 2013).

Because the time to fully replace a catalyzed crust is relatively short compared to the binary evolution timescale, it seems likely that most, if not all, observed LMXB neutron stars have fully accreted crusts. However, since we have only been able to observe accretion for the last several

decades, we can't be certain that the neutron star in any given LMXB has a fully accreted crust rather than a catalyzed or partially-accreted crust.

To test whether or not different observed accreting neutron stars share crust properties, I follow the methodology described in Chapter 2 to fit dStar models to the cooling curves of most of the LMXBs in which crustal cooling has been observed. First, I fit to the cooling curve data of each LMXB independently by estimating the posterior distribution of the core temperature  $T_c$ , impurity  $Q_{\text{imp}}$ , shallow heating rate  $Q_{\text{sh}}$ , and inner heating rate  $Q_{\text{in}}$ . I then jointly fit to the cooling curves of all sources by allowing three of the four dStar parameters to vary independently for each source while the remaining parameter varies jointly across all sources. By comparing the goodness of fit of the independent fits to the joint fits, I can determine whether or not a given parameter may be shared across all of the neutron stars.

### 3.1 Independently Fitting Neutron Star Models to Observations

Crustal cooling following an accretion outburst has been observed in 9 LMXBs: KS 1731–260 (Merritt et al., 2016), MXB 1659–29 (Parikh et al., 2019), XTE J1701–462 (Parikh et al., 2020), EXO 0748–676 (Parikh et al., 2020), IGR J17480–2446 (Degenaar et al., 2013), MAXI J0556–332 (Parikh et al., 2017a), Aql X-1 (Waterhouse et al., 2016), Swift 174805.3–244637 (Degenaar et al., 2015), and 1RXS J180408.9–342058 (Parikh et al., 2017b). Crustal cooling following multiple outbursts has been observed in MXB 1659–29, MAXI J0556–332, and Aql X-1. Masses and radii have not been measured for any of these sources, so canonical values have been used for interpretation of observational data. Most of the sources are estimated to be at a distance of 5–9 kpc away, though MAXI J0556–332 appears to be much farther with an estimated distance of 43.6 kpc. The duration of their outbursts,  $t_{\text{acc}}$ , ranges from weeks to decades and their accretion rates vary by nearly two orders of magnitude. Table 3.1 summarizes these system's properties. In Sections 3.1.1–3.1.9, I will discuss each of these sources in greater detail.

For the accretion rates and quiescent temperatures, I use data reported by the papers cited in Table 3.1. To be consistent with the observational data, I use the masses, radii, and distances used to obtain the observed temperatures. I also set, for all sources, the column depth of the light element



Table 3.1 Neutron stars with observed and fitted crust cooling.

Source	$M$ ( $M_{\odot}$ )	$R$ (km)	$d$ (kpc)	$\langle \dot{M} \rangle$ ( $10^{17} \text{g s}^{-1}$ )	$t_{\text{acc}}$ (yr)	Reference
KS 1731–260	1.4	10	7	1.5	12	1
MXB 1659–29 I	1.6	12	9	1.2	2.5	2
MXB 1659–29 II	1.6	12	9	0.58	1.5	2
XTE J1701–462	1.6	12	8.8	11	1.6	3
EXO 0748–676	1.6	12	7.4	0.3	24	3
IGR J17480–2446	1.4	10	5.5	1.9	0.22	4
MAXI J0556–332 I	1.4	10	43.6	13	1.3	5
MAXI J0556–332 II	1.4	10	43.6	5.1	0.17	5
MAXI J0556–332 III	1.4	10	43.6	7.1	0.25	5
Aql X-1 I	1.6	11	5	4	0.17	6
Aql X-1 II	1.6	11	5	4	0.18	6
Aql X-1 III	1.6	11	5	0.8	0.05	6
Swift 174805.3–244637	1.4	10	5.5	0.96	0.15	7
1RXS J180408.9–342058	1.6	11	5.8	0.58	0.15	8

References: (1) Merritt et al. (2016); (2) Parikh et al. (2019); (3) Parikh et al. (2020); (4) Degenaar et al. (2013); (5) Parikh et al. (2017a); (6) Waterhouse et al. (2016); (7) Degenaar et al. (2015); (8) Parikh et al. (2017b)

layer to  $10^8 \text{g cm}^{-2}$ , which is roughly where  $^4\text{He}$  unstably ignites.

For each source, I perform MCMC analysis using `emcee`, with the following free parameters: core temperature ( $T_c$ ), impurity ( $Q_{\text{imp}}$ ), shallow heating rate ( $Q_{\text{sh}}$ ), and inner heating rate ( $Q_{\text{in}}$ ). I assume Gaussian errors in the observational data to calculate the  $\chi^2$  and determine the likelihood<sup>1</sup>,  $\mathcal{L}(\theta) = \exp(-\chi^2/2)$ , where  $\theta$  is the vector of the parameters,  $\theta = (T_c, Q_{\text{imp}}, Q_{\text{sh}}, Q_{\text{in}})$ .

I use a uniform prior probability distribution that is bounded by a set of minimum and maximum values for each parameter (see Table 3.2). I choose the bounds such that the allowable parameter space is much larger than the values expected based on prior results. This is because for some sources, the data does not constrain one or some of the parameters very well; the MCMC will take too long to converge because the walkers are not bound to any preferential region in parameter space along the unconstrained parameter(s). By choosing a large, but finite prior distribution, the walkers are more bound, allowing the MCMC to converge even if the resulting posterior distribution is

<sup>1</sup>In Section 2.4, I used the notation  $P(\mathbf{D}|\Theta)$  to represent the likelihood for clarity when discussing it in the context of Bayes' Theorem and the underlying principles of Bayesian statistics. From here on, I will use the notation  $\mathcal{L}(\theta)$  to be consistent with most references on quantifying goodness of fit and model selection.

uniform in one or more parameters.

Core temperatures for sources tend to be around  $10^8$  K (see references in Table 3.1), so I choose a minimum and maximum of  $10^7$  K and  $10^9$  K respectively.

The impurity in the outer crust can vary widely depending on the initial composition determined by thermonuclear reactions in the envelope. Lau et al. (2018) found a range of  $Q_{\text{imp}} \approx 4$  for superburst ashes up to  $Q_{\text{imp}} \approx 80$  for extreme X-ray burst ashes. However, in the inner crust, they find that the impurity decreases for all compositions, with the most extreme case yielding  $Q_{\text{imp}} \approx 10$ . Since impurity scattering dominates the thermal conductivity at high densities where phonons are frozen out, our cooling models are most sensitive to  $Q_{\text{imp}}$  in the inner crust. Because of this, I choose a minimum value of  $Q_{\text{imp}} = 0$  and a maximum value of  $Q_{\text{imp}} = 40$ , which gives significant room in parameter space above the expected upper limit of  $Q_{\text{imp}} \approx 10$ .

Similarly, for the shallow heating rate and inner heating rate, I choose minimum values of  $0 \text{ MeV u}^{-1}$  and maximum values of  $40 \text{ MeV u}^{-1}$ . Fits to most outbursts previously have required no more than a few  $\text{MeV u}^{-1}$  of shallow heating (see, e.g., Brown & Cumming, 2009; Page & Reddy, 2013; Merritt et al., 2016). One outburst of MAXI J0556–332 required a shallow heating rate of  $Q_{\text{sh}} \approx 17$  (Deibel et al., 2015; Parikh et al., 2017a), which is the greatest for any observed outburst. The inner heating rate is predicted to be  $1.5\text{--}2 \text{ MeV u}^{-1}$  (Haensel & Zdunik, 2003). In order to quantify how well our data allows us to constrain the inner heating rate, however, I set its upper limit to a much higher value.

Table 3.2 Maximum and minimum values for each parameter. I use a prior probability distribution that is uniform between these values and 0 otherwise.

Parameter	Minimum	Maximum
Core Temperature (K)	$10^6$	$10^9$
Impurity	0	40
Shallow Heating Rate ( $\text{MeV u}^{-1}$ )	0	40
Inner heating rate ( $\text{MeV u}^{-1}$ )	0	40

I let the MCMC chain continue until the number of iterations is equal to 100 times the maximum autocorrelation time (see Section 2.4 for description of MCMC convergence criteria). For my

resulting posterior probability distributions I discard a number of initial samples equal to 2 times the maximum autocorrelation time.

With the set of samples from the converged MCMC chain, I can estimate the posterior probability distribution of the model parameters. I can then determine which parameters are well-constrained, what correlations exist among parameters, and what set of parameters fit the data best (by minimizing  $\chi^2$  and therefore maximizing  $\mathcal{L}(\theta)$ ). The posterior distribution can be estimated using a kernel density estimator to produce a smooth probability distribution function. However, it can be computationally expensive and complicated to perform integration and other operations this way. For simplicity and computational ease, I approximate the posterior distribution with a histogram in parameter space with the MCMC samples.

Due to the random nature of the sampler, it is unlikely to actually sample the best-fit parameters. To find the best-fit parameters, I use the `optimize` subpackage of SciPy (Virtanen et al., 2020) to minimize the  $\chi^2$ . I tested all available minimization methods with different sets of initial parameters for a few different sources and found that the Nelder-Mead method (Nelder & Mead, 1965) most consistently found parameters with a smaller  $\chi^2$  than other minimization methods so that is the method I use for the following results. To determine the initial guess for  $\theta$ , I chose the midpoint of the histogram bin with the greatest number of samples.

To quantify how well-constrained the parameter distributions are, I calculate an approximate 68% credible region (often called a credible interval for unimodal distributions). A 68% credible region is a volume in parameter space that contains 68% of the mass of the posterior probability distribution. The Bayesian interpretation of such a credible region is that there is a 68% probability that the true parameters lie within the region (see Jaynes & Kempthorne, 1976, for a thorough discussion on credible regions and how they differ from their frequentist counterparts, confidence intervals). A credible region is not unique; one can always choose another volume in parameter space that contains the same mass of the probability distribution. The credible region I approximate is the highest density 68% credible region, which is the region of smallest volume that contains 68% of the posterior probability distribution's mass. An advantage of using the highest density

credible region is that it will always contain the set of parameters with the maximum probability.

Because the MCMC chain produces samples such that their density is proportional to posterior probability, a 68% credible region should contain 68% of the MCMC samples. I approximate the highest density 68% credible region by selecting the histogram bins in ascending order of sample counts until I have selected bins containing 68% of the samples. For each parameter, I find the smallest and greatest values among my selected bins and report the difference between these values and the best-fit parameters as the uncertainties in my results. Note that uncertainties reported this way only give a general idea of the spread of the posterior distribution; they don't capture correlations between parameters and can be misleading if a distribution is multimodal. To better convey more complicated features of the distributions, I also include corner plots which show how each pair of parameters correlates.

The reduced chi-squared statistic has been commonly used in astronomy as a measure of how well a model fits observed data (Wall & Jenkins, 2012), including previous analysis of cooling curves (see, e.g., Cackett et al., 2010; Merritt et al., 2016; Parikh et al., 2019; Ootes et al., 2018). The reduced chi-squared statistic is defined as  $\chi^2_\nu = \chi^2/\nu$  where  $\nu$  is the number of degrees of freedom (dof). The number of degrees of freedom is typically calculated as  $\nu = \text{number of observation} - \text{number of parameters}$ . A model that fits the data well should have  $\chi^2_\nu \approx 1$ ; as a general rule, if  $\chi^2_\nu \gtrsim 4$ , then we would consider the model a poor fit (Wall & Jenkins, 2012). The reduced chi-squared statistic is useful because it is simple to calculate and accounts for overfitting by penalizing models with a large number of parameters relative to the number of observations. To provide a general measure of how well my dStar models fit the data to each source and for comparison to other, previous cooling curve analysis, I report  $\chi^2_\nu$  for each model.

However, the reduced chi-squared statistic is often not a good measure of a model's goodness of fit or for model selection or rejection (see Andrae et al., 2010, for a more thorough discussion on the possible pitfalls of the reduced chi-square statistic). In part, this is because the number of degrees of freedom can generally only be directly determined for linear models without priors. For other models, for example when there are strong degeneracies among model parameters, the

degrees of freedom will not simply be the difference between the number of data points and the number of parameters. This means that naively using  $\chi^2_\nu$  to determine whether a model fits well or for choosing between competing models may lead to incorrect results. Despite this, I still calculate and report  $\chi^2_\nu$  as a rough indicator of the goodness of fit and for comparison to previous analysis.

### 3.1.1 KS 1731–260

KS 1731–260 was the first LMXB for which crustal cooling was observed. It was discovered in outburst by the *Mir-Kvant* instrument in August 1989 (Sunyaev, 1989). Further analysis found that the source was in outburst as early as October 1988 (Sunyaev, 1990). Over its 12.5 year outburst, it was observed by many instruments: the *ART-P* (Sunyaev, 1990) and *Sigma* (Barret et al., 1992) telescopes aboard *GRANAT*; *Ginga* (Yamauchi & Koyama, 1990); *ROSAT* (Barret et al., 1998); *RXTE* (Smith et al., 1997; Wijnands & van der Klis, 1997; Munro et al., 2000); *ASCA* (Narita et al., 2001); and *BeppoSAX* (Kuulkers et al., 2002). It went into quiescence some time between January 21 and February 7 of 2001 and the first measurement of its surface temperature in quiescence was taken March 27 2001 (Wijnands et al., 2001).

The source exhibits Type I X-ray bursts (Sunyaev, 1989) and photospheric radius expansion (PRE) bursts (Munro et al., 2000) and has exhibited one superburst (Kuulkers et al., 2002). Assuming the peak luminosity of the PRE burst is equal to the Eddington luminosity, Munro et al. (2000) estimated that the upper limit of the source’s distance is 7 kpc and that it is located near the galactic center. Özel et al. (2012) fit the blackbody emission of the PRE bursts to obtain mass and radius constraints for the neutron star:  $R \leq 12.5$  km and  $M \leq 2.1 M_\odot$ .

I fit my dStar model to the 8 quiescent temperature observations reported by Merritt et al. (2016) who use the neutron star atmosphere (nsa) model of Zavlin et al. (1996) and assume a distance of  $d = 7$  kpc, a mass of  $M = 1.4 M_\odot$ , and a radius of  $R = 10$  km. I partition the accretion history as described in Section 2.5 into six accretion epochs using the exponential scheme with an exponential growth rate of 2.0. Over the 12.4 yr accretion outburst, the average accretion rate was  $1.5 \times 10^{17}$  g/s.

I find best-fit parameters of  $T_c = 6.3^{+1.0}_{-0.6} \times 10^7$  K,  $Q_{\text{imp}} = 6.3^{+4.2}_{-6.3}$ ,  $Q_{\text{sh}} = 1.1^{+1.0}_{-0.6}$  MeV u<sup>-1</sup>, and

$Q_{\text{in}} = 0.0^{+8.1}_{-0.0} \text{ MeV u}^{-1}$  with  $\chi^2_{\nu} = 7.7$  (4 dof). A corner plot of the distribution and the cooling curve using the best-fit parameters is shown in Figure 3.1.

The posterior distributions for the core temperature and the shallow heating rate appear close to normal, centered on their best-fit values. Low values of impurity and inner heating rate appear to be favored. These parameter estimates are all consistent with prior analysis performed on the cooling curve of KS 1731–260 (see, e.g., Brown & Cumming, 2009; Cackett et al., 2010; Merritt et al., 2016; Ootes et al., 2016). Merritt et al. (2016) were able to fit the cooling curve better (with  $\chi^2_{\nu} = 2.0$ ) by setting the column depth of the light element layer to  $10^4 \text{ g cm}^{-2}$ . Though such a shallow light element layer apparently fits the observed data better, it seems physically improbable. If the light elements in the envelope were removed, such as via a type I X-ray burst, it would take  $\sim 1 \text{ s}$  to accumulate a light element layer of  $10^4 \text{ g cm}^{-2}$  at  $\dot{M} = 10^{17} \text{ g s}^{-1}$ . With this system's observed type I X-ray burst recurrence time of a few hours (Galloway et al., 2008), it is very unlikely that accretion shut off within a few seconds of a type I X-ray burst.

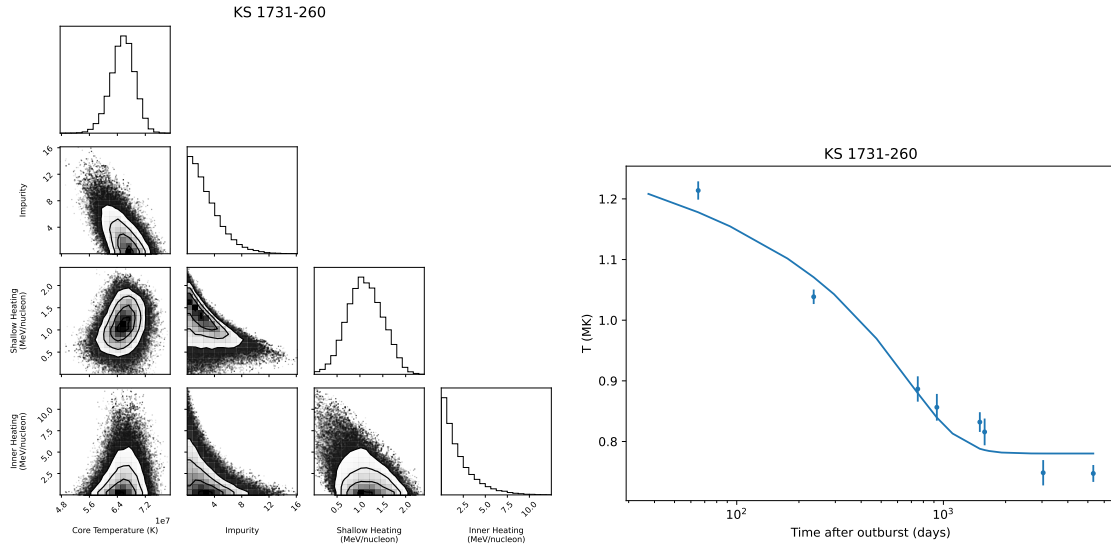


Figure 3.1 Left: Posterior distribution for KS 1731–260. Right: Observed surface effective temperatures (points; Merritt et al., 2016) and the model with best-fit parameters  $T_c = 6.3 \times 10^7 \text{ K}$ ,  $Q_{\text{imp}} = 6.3$ ,  $Q_{\text{sh}} = 1.1 \text{ MeV u}^{-1}$ , and  $Q_{\text{in}} = 0.0 \text{ MeV u}^{-1}$ .

### 3.1.2 MXB 1659–29

MXB 1659–29 was first discovered as an LMXB that exhibited Type I X-ray bursts by Lewin et al. (1976). It was in outburst for about 2–2.5 yr but follow-up observations of the quiescent crust temperature were not possible at the time. It has since had two more accretion outbursts, both of which have had follow-up observations of the crust temperature (Parikh et al., 2019). MXB 1659–29 also exhibits X-ray eclipses with an orbital period of about 7.1 hr and duration of about 900 s (Cominsky & Wood, 1984, 1989; Oosterbroek et al., 2001), suggesting that we are viewing the accretion nearly disk edge-on. The distance to MXB 1659–29 is estimated from PRE bursts Galloway et al. (2008) to be  $9.2 \pm 2$  kpc ( $12 \pm 3$  kpc) for hydrogen-rich (helium-rich) bursts.

The first observed outburst of MXB 1659–29 lasted for 2.5 yr, with the source going into quiescence in September 2001 (Wijnands et al., 2003). Observations taken about 1400 and 2400 d after the end of the outburst were consistent with each other, suggesting that the crust had stopped cooling and returned to thermodynamic equilibrium with the core (Cackett et al., 2008). However, another observation taken 4 yr later showed an unexpected drop in temperature (Cackett et al., 2013). It was unknown at the time if this was due to continued cooling or to an increase in internal absorption (e.g., due to an increase in the height of the accretion disk). Unfortunately, before it could be further investigated, its second accretion outburst began, preventing further observations of the surface. The second outburst lasted for 1.7 yr, going into quiescence in March 2017.

I fit my dStar model to the quiescent temperature data of Parikh et al. (2019) who use the neutron star atmosphere (nsatmos) model of Heinke et al. (2006) and assume a distance of  $d = 9$  kpc, a mass of  $M = 1.6 M_{\odot}$ , and a radius of  $R = 12$  km. Parikh et al. (2019) omit the last temperature observation of the first outburst and argue that it is most likely due to increased internal absorption rather than crustal cooling. I use a single accretion epoch for each outburst. Parikh et al. (2019) report the bolometric fluences for outburst I and outburst II to be  $0.17 \text{ erg cm}^{-2}$  and  $5.18 \times 10^{-2} \text{ erg cm}^{-2}$  respectively. I calculate the average accretion rate using Equation (2.11) with  $\eta = 0.2$  and  $\xi = 1$ . This yields average accretion rates of  $1.18 \times 10^{17} \text{ g s}^{-1}$  and  $5.82 \times 10^{16} \text{ g s}^{-1}$  for outbursts I and II respectively.

I find best-fit parameters of  $T_c = 4.5_{-1.2}^{+0.9} \times 10^7$  K,  $Q_{\text{imp}} = 0.0_{-0.0}^{+4.7}$ ,  $Q_{\text{sh}} = 1.4_{-0.2}^{+0.7}$  MeV u<sup>-1</sup>, and  $Q_{\text{in}} = 17.0_{-17.0}^{+3.6}$  MeV u<sup>-1</sup> with  $\chi^2_{\nu} = 3.5$  (9 dof). A corner plot of the distribution and the cooling curve using the best-fit parameters is shown in Figure 3.2.

These values for  $T_c$ ,  $Q_{\text{imp}}$ , and  $Q_{\text{sh}}$  are similar to the estimates of Parikh et al. (2019) who found that they could use the same shallow heating rate for both outbursts. An inner heating rate of 17 MeV u<sup>-1</sup> is much higher than any theoretical predictions, though the credible region includes values down to 0 MeV u<sup>-1</sup>. The marginal distributions of the core temperature and shallow heating rate are roughly normal, while the marginal distributions of the impurity and inner heating rate are right-skewed, favoring low values. The marginalized distribution for the inner heating rate has a long tail extending to high values, which suggests that the data for MXB 1659–29 does not constrain the inner heating very well. So even though the best-fit parameters include a very high inner heating rate, lower values of inner heating can likely still fit the data well.

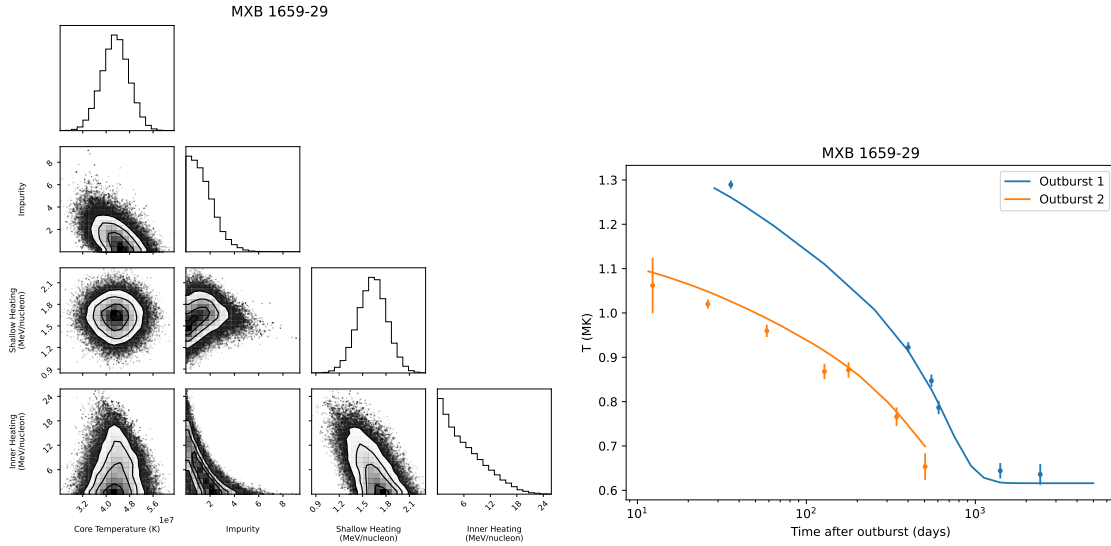


Figure 3.2 Left: Posterior distribution for MXB 1659–29. Right: Observed surface effective temperatures (points; Parikh et al., 2019) and the model with best-fit parameters  $T_c = 4.5 \times 10^7$  K,  $Q_{\text{imp}} = 0.0$ ,  $Q_{\text{sh}} = 1.4$  MeV u<sup>-1</sup>, and  $Q_{\text{in}} = 17.0$  MeV u<sup>-1</sup>.

### 3.1.3 XTE J1701–462

XTE J1701–462 was discovered in outburst in January 2006 by *ASM* (Remillard et al., 2006). This outburst was exceptionally luminous, with a luminosity near the Eddington limit, and lasted



about 1.6 yr, going into quiescence in August 2007 (Homan et al., 2007). Lin et al. (2009) estimated its distance to be 8.8 kpc with 15% error based on analysis of an observed PRE burst.

The system has possibly exhibited some short-term low-level flares during quiescence (Fridriksson et al., 2010). These flares are distinguished from surface temperature observations by their large contribution from the hard, power-law-shaped component of their spectra. Fridriksson et al. (2011) determined that they are unlikely to have had a measurable effect on the overall cooling curve due to their short duration.

Parikh et al. (2020) report cooling curve data 3200 d into quiescence. They use the *nsatmos* atmosphere model and assume a distance of  $d = 8.8$  kpc, a mass of  $M = 1.6 M_{\odot}$ , and a radius of  $R = 12$  km. They omit the observations corresponding to the low-level accretion flares. For my *dStar* model I use a single accretion epoch with a rate of  $1.1 \times 10^{18} \text{ g s}^{-1}$  (Fridriksson et al., 2010).

I find best-fit parameters of  $T_c = 2.38_{-0.65}^{+0.23} \times 10^8 \text{ K}$ ,  $Q_{\text{imp}} = 40.0_{-40.0}^{+0.0}$ ,  $Q_{\text{sh}} = 0.2_{-0.0}^{+0.1} \text{ MeV u}^{-1}$ , and  $Q_{\text{in}} = 0.0_{-0.0}^{+5.0} \text{ MeV u}^{-1}$  with  $\chi^2_{\nu} = 5.5$  (12 dof). A corner plot of the distribution and the cooling curve using the best-fit parameters is shown in Figure 3.3.

The fit suggests that the core is very hot and a small, but non-zero amount of shallow heating occurred during its accretion outburst. A small inner heating rate appears to be favored, which is consistent with theoretical predictions. The impurity is very unconstrained—the credible region extends from the minimum to the maximum of the prior distribution—but appears to slightly favor lower values.

### 3.1.4 EXO 0748–676

EXO 0748–676 was discovered with *EXOSAT* in 1986 (Parmar et al., 1986). It exhibited a 24-year-long accretion outburst at an average luminosity of 5% the Eddington luminosity (Degenaar et al., 2009) which ended in September 2008 (Wolff et al., 2008).

Galloway et al. (2008) found that the spectral properties of observed PRE bursts suggest that ignition occurred in a He-rich environment. Assuming helium-rich bursts, Galloway et al. estimate the distance to be  $7.4 \pm 0.9$  kpc. EXO 0748–676 exhibits X-ray eclipses with 8.3 min duration and a period of 3.82 hr, likely due to obscuration from a thickened region of its accretion disk which

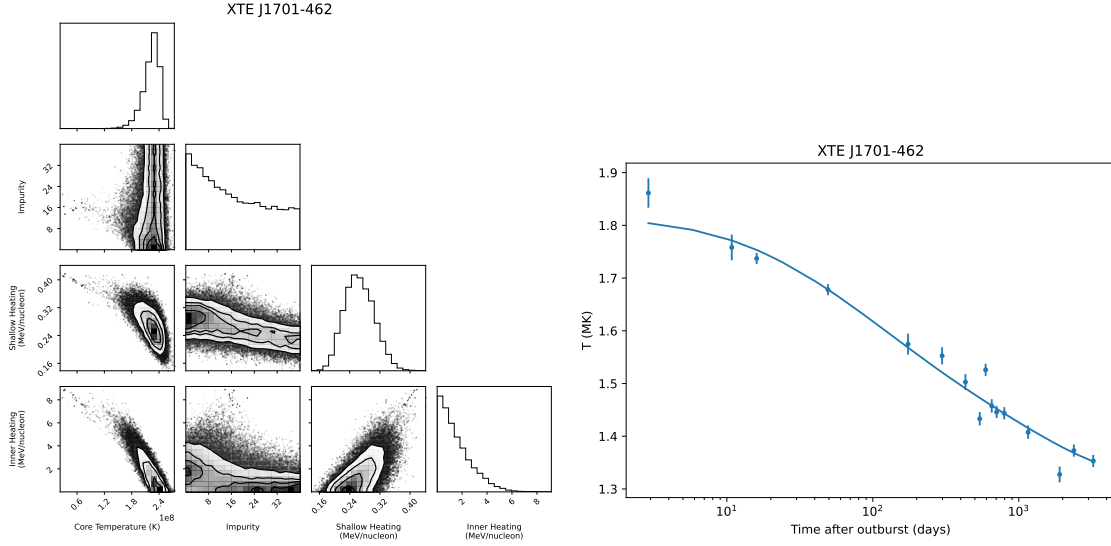


Figure 3.3 Left: Posterior distribution for XTE J1701–462. Right: Observed surface effective temperatures (points; Parikh et al., 2020) and the model with best-fit parameters  $T_c = 2.4 \times 10^8$  K,  $Q_{\text{imp}} = 40.0$ ,  $Q_{\text{sh}} = 0.2 \text{ MeV u}^{-1}$ , and  $Q_{\text{in}} = 0.0 \text{ MeV u}^{-1}$ .

we are viewing nearly edge-on (Parmar et al., 1986).

Parikh et al. (2020) report the observed temperatures nearly 3500 d into quiescence. They use the `nsatmos` atmosphere model and assume a distance of  $d = 7.4$  kpc, a mass of  $M = 1.6 M_\odot$ , and a radius of  $R = 12$  km.

I find best-fit parameters of  $T_c = 2.20 \pm 0.02 \times 10^8$  K,  $Q_{\text{imp}} = 0.0_{-0.0}^{+2.4}$ ,  $Q_{\text{sh}} = 3.7_{-0.4}^{+0.3} \text{ MeV u}^{-1}$ , and  $Q_{\text{in}} = 0.0_{-0.0}^{+0.9} \text{ MeV u}^{-1}$  with  $\chi^2_\nu = 111.5$  (9 dof). A corner plot of the distribution and the cooling curve using the best-fit parameters is shown in Figure 3.4.

Such a large value of  $\chi^2_\nu$  suggests that my dStar model fits the data very poorly. While the cooling curve captures the general decrease in temperature that is observed, there is a significant amount of variation relative to the reported temperature uncertainties. This includes an apparent unexpected increase in temperature around 1000 days into quiescence (as noted by Parikh et al., 2020). This suggests that we are likely not correctly interpreting the temperature data. One possible explanation for this is that the accretion disk may be interfering with direct observations of the neutron star’s surface.

If we accept the parameter estimates as they are, the results are generally consistent with the estimates for other sources. All of the parameters have tight fits. The core temperature is high and

a very low impurity is strongly favored. The shallow heating rate is slightly higher than for most other sources, but fairly tightly constrained. A very low inner heating rate is strongly preferred, which is inconsistent with the theoretical estimates of  $1.5\text{--}2 \text{ MeV u}^{-1}$ . My results are somewhat at odds with those of Degenaar et al. (2014) who estimated a similar core temperature but  $Q_{\text{imp}} = 40$  and  $Q_{\text{sh}} = 1.8 \text{ MeV u}^{-1}$ .

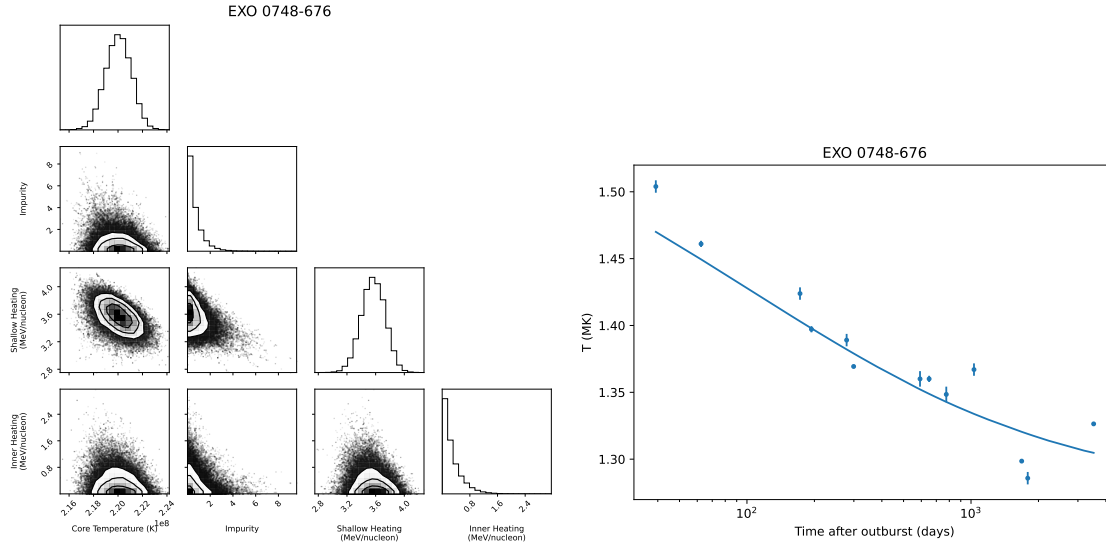


Figure 3.4 Left: Posterior distribution for EXO 0748–676. Right: Observed surface effective temperatures (points; Parikh et al., 2020) and the model with best-fit parameters  $T_c = 2.2 \times 10^8 \text{ K}$ ,  $Q_{\text{imp}} = 0.0$ ,  $Q_{\text{sh}} = 3.7 \text{ MeV u}^{-1}$ , and  $Q_{\text{in}} = 0.0 \text{ MeV u}^{-1}$ .

### 3.1.5 IGR J17480–2446

IGR J17480–2446 is an 11 Hz pulsar with a companion of mass  $M \approx 0.8 M_{\odot}$  (Testa et al., 2012) located in the Terzan 5 globular cluster (Strohmayer et al., 2010) which is a distance of  $d \approx 5.5 \text{ kpc}$  away (Ortolani et al., 2007). An accretion outburst was detected beginning in October 2010 (Bordas et al., 2010; Pooley et al., 2010) that lasted for 79 d with an average accretion rate of  $2 \times 10^{17} \text{ g s}^{-1}$  (Degenaar & Wijnands, 2011).

Quiescent surface effective temperatures for IGR J17480–2446 are taken from Degenaar et al. (2013) which includes data up to 784 d after the end of the outburst. They also include temperature observations taken 2776 days before and 582 days before the end of the outburst. These temperatures estimated during these two observations are consistent with each other and lower than all post-

outburst observations, so I assume that they represent the surface temperature when the crust is in equilibrium with the core. Degenaar et al. use the *nsatmos* atmosphere model and assume a distance of  $d = 5.5$  kpc, a mass of  $M = 1.4 M_{\odot}$ , and a radius of  $R = 10$  km.

I find best-fit parameters of  $T_c = 7.7_{-1.6}^{+2.2} \times 10^7$  K,  $Q_{\text{imp}} = 14.3_{-14.3}^{+25.6}$ ,  $Q_{\text{sh}} = 1.2_{-0.7}^{+0.5}$  MeV u<sup>-1</sup>, and  $Q_{\text{in}} = 7.1_{-7.1}^{+23.6}$  MeV u<sup>-1</sup> with  $\chi^2_{\nu} = 1.0$  (5 dof). A corner plot of the distribution and the cooling curve using the best-fit parameters is shown in Figure 3.5.

The impurity is not well constrained, with a credible region that includes the whole range of the prior. Based on the marginalized distribution, though, it appears that values less than  $Q_{\text{imp}} \lesssim 10$  are disfavored. An inner heating rate in the range of 4–9 MeV u<sup>-1</sup> is favored, which is a factor of several higher than the theoretical estimates of 1.5–2 MeV u<sup>-1</sup>. The shallow heating rate is similar to the fits to most other sources.

Ootes et al. (2019) also fit crust cooling models to the cooling curve of IGR J17480–2446, including three temperature measurements that were taken after those of Degenaar et al. (2013). Ootes et al. fixed the inner heating to  $Q_{\text{in}} = 1.93$  MeV u<sup>-1</sup> and divided the crust into 5 density regimes, each with an impurity that could vary independently with a uniform prior of  $0 \leq Q_{\text{imp}} \leq 1500$ . They found that the impurity was not well constrained in any of the regions of the crust, though the outermost region slightly favored  $Q_{\text{imp}} \lesssim 500$ . Their results for core temperature and shallow heating were similar to mine.

### 3.1.6 MAXI J0556–332

MAXI J0556–332 is an LMXB that was discovered in January 2011 (Matsumura et al., 2011) and exhibited a 16 month accretion outburst that ended in May 2012. Thermonuclear X-ray bursts have not been observed but the source was determined to likely contain a neutron star because its optical-to-X-ray and radio-to-X-ray flux ratios are similar to galactic NS LMXBs (Russell et al., 2011; Coriat et al., 2011). Additionally, analysis of *RXTE* data shows that it is similar to “Z sources”, a class of near- and super-Eddington NS LMXBs that trace out distinct tracks in their X-ray color-color diagrams and hardness-intensity diagrams (Homan et al., 2011). Spectral fits of *XMM-Newton* and *Chandra* data suggest that it is at a distance of  $43.6_{-1.6}^{+0.9}$  kpc (Parikh et al., 2017a).

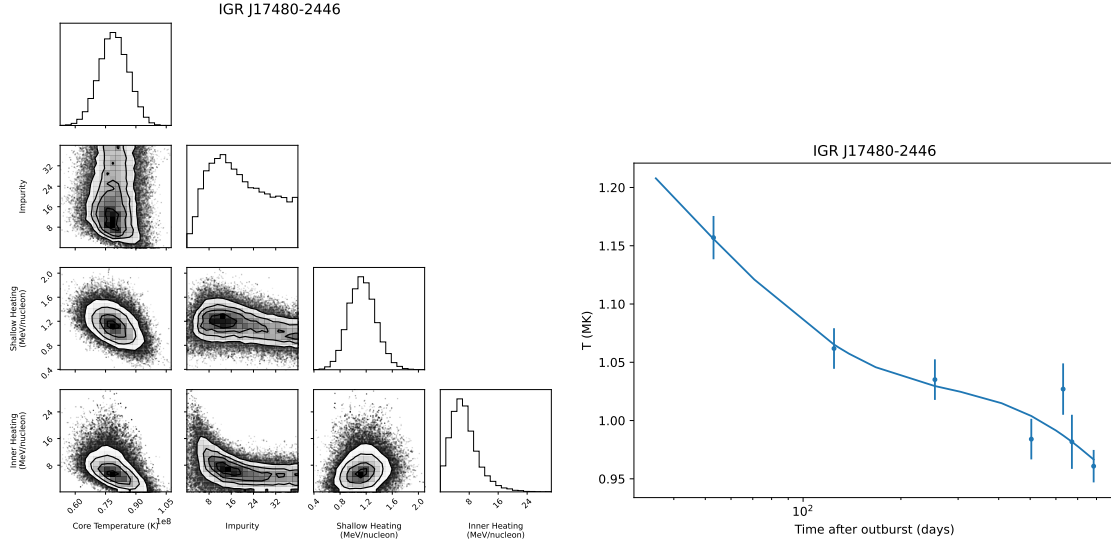


Figure 3.5 Left: Posterior distribution for IGR J17480–2446. Right: Observed surface effective temperatures (points; Degenaar et al., 2013) and the model with best-fit parameters  $T_c = 7.7 \times 10^7$  K,  $Q_{\text{imp}} = 14$ ,  $Q_{\text{sh}} = 1.2 \text{ MeV u}^{-1}$ , and  $Q_{\text{in}} = 7.1 \text{ MeV u}^{-1}$ . The pre-outburst temperature observations are not included in this figure, though they were used to fit the model.

Since its detection, MAXI J0556–332 has exhibited 4 accretion outbursts: from January 2011 to May 2012; from October 2012 to December 2012; from January 2016 to April 2016; and from April 2020 to November 2020.

Though MAXI J0556–332 has exhibited 4 accretion outbursts, there is only one quiescent observation following the fourth outburst. Since this doesn’t provide enough data to reliably fit the model parameters, I only fit to the quiescent data of the first three outbursts which are reported by Parikh et al. (2017a). They use the `nsa` atmosphere model and assume a distance of  $d = 43.6$  kpc, a mass of  $M = 1.4 M_{\odot}$ , and a radius of  $R = 10$  km.

I find best-fit parameters of  $T_c = 1.0^{+23.2}_{-0.0} \times 10^7$  K,  $Q_{\text{imp}} = 0.0^{+40.0}_{-0.0}$ ,  $Q_{\text{sh}} = 3.9^{+0.4}_{-1.2} \text{ MeV u}^{-1}$ , and  $Q_{\text{in}} = 17.9^{+0.6}_{-17.9} \text{ MeV u}^{-1}$  with  $\chi^2_{\nu} = 399.4$  (13 dof). A corner plot of the distribution and the cooling curve using the best-fit parameters is shown in Figure 3.6.

My model fits very poorly to the cooling curves of MAXI J0556–332. The distributions for all parameters exhibit some degree of bimodality. Parikh et al. (2017a) found that fitting the cooling curves of the first three outbursts required different rates of shallow heating for each outburst. They found an exceptionally high heating rate ( $Q_{\text{sh}} \approx 17 \text{ MeV u}^{-1}$ ) for the first outburst (see also Deibel

et al., 2015). The following two outbursts required heating rates similar to the observed outbursts of other sources ( $Q_{\text{sh}} = 2.2 \text{ MeV u}^{-1}$  and  $Q_{\text{sh}} = 0.33 \text{ MeV u}^{-1}$  for the second and third outburst respectively). Page et al. (2022) suggested that the high shallow heating rate of the first outburst can be explained by the occurrence of a “hyperburst”, an extremely rare, energetic burst caused by unstable thermonuclear burning of oxygen and neon in the crust.

The different heating properties of the first outburst explain the bimodal distributions of my fits. I do not allow the model parameters to vary for different outbursts so the MCMC will generally favor parameters that are consistent with either the first outburst or the later outbursts.

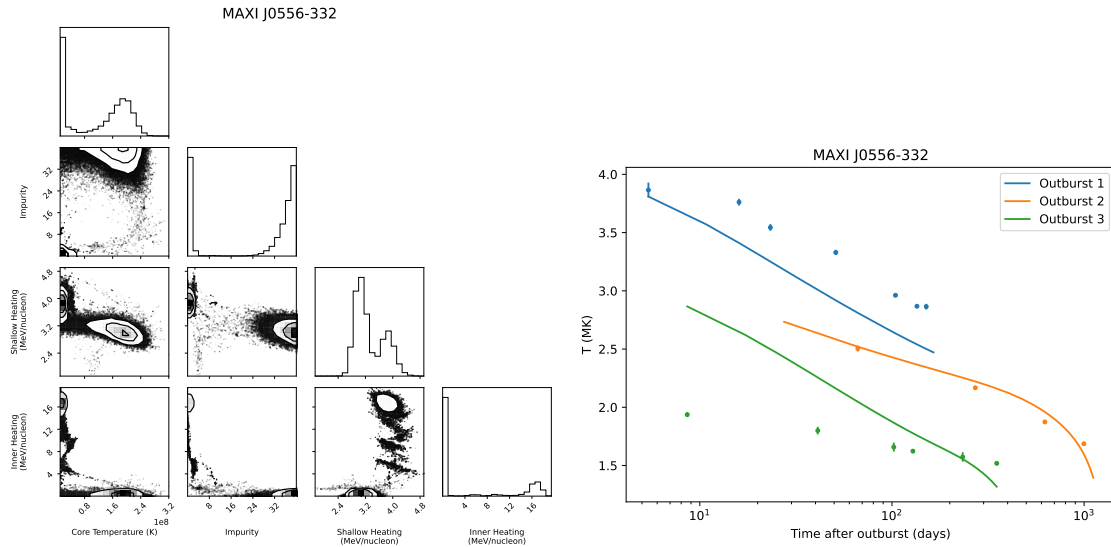


Figure 3.6 Left: Posterior distribution for MAXI J0556–332. Right: Observed surface effective temperatures (points; Parikh et al., 2017a) and the model with best-fit parameters  $T_c = 1.0 \times 10^7 \text{ K}$ ,  $Q_{\text{imp}} = 0$ ,  $Q_{\text{sh}} = 3.9 \text{ MeV u}^{-1}$ , and  $Q_{\text{in}} = 18 \text{ MeV u}^{-1}$ .

### 3.1.7 Aql X-1

Aql X-1 was detected during some of the earliest surveys of X-ray sources using rocket-borne X-ray detectors (Friedman et al., 1967). It exhibits type I X-ray bursts (Koyama et al., 1981) and mHz quasi-periodic oscillations (Revnivtsev et al., 2001; Altamirano et al., 2008). Casella et al. (2008) inferred that it has a spin frequency of 550 mHz, based on observations of pulsations during persistent emission. Rutledge et al. (2001b) estimate its distance to be in the range of 4–6.5 kpc with the best estimate being 5 kpc, based on optical observations of its main sequence companion.

Aql X-1 exhibits accretion outbursts of varying luminosity and duration roughly once per year (Kaluzienski et al., 1977; Kitamoto et al., 1993; Güngör et al., 2014). Campana et al. (2013) examined ASM data from 1996-2012, revealing 20 outbursts with durations ranging from 1–26 weeks with X-ray luminosity ranges from  $10^{35}$ – $10^{36} (D/5.0 \text{ kpc})^2 \text{ erg s}^{-1}$  (where  $D$  is the distance to Aql X-1). During quiescence, Aql X-1 displays non-monotonic variability with the observed flux sometimes increasing by a factor of up to 5 (Cackett et al., 2011; Coti Zelati et al., 2014). This suggests that low-level accretion flares may occur during quiescent periods.

An LMXB is considered to be in a state of outburst when the observed X-ray flux is high enough to be detected by an all-sky monitoring telescope and quiescent otherwise. This means that accretion-induced X-ray emission may still occur in a source in a quiescent state if it is below the detection threshold. When some LMXBs, including Aql X-1, transition from outburst to quiescence, they display a brief, rapid decay in flux followed by a longer, gradual decay (see, e.g., Fridriksson et al., 2011; Campana et al., 2014; Homan et al., 2014). The rapid decay is usually interpreted as a decay in the accretion-induced emission while the gradual decay is interpreted as cooling from the surface. In order to use observations in quiescence to determine the surface temperature and therefore the crustal properties of a neutron star in an LMXB, it may be important to distinguish between observations where X-ray emission is dominated by thermal emission at the surface and observations that contain a significant portion of emission due to accretion.

Waterhouse et al. (2016) analyzed quiescent observations following three accretion outbursts from 2011-2015. Two of these outbursts (in 2011 and 2013) lasted about 8 weeks with average accretion rates of  $4 \times 10^{17} \text{ g s}^{-1}$  while the outburst in 2015 lasted 3 weeks with an accretion rate of  $8 \times 10^{16} \text{ g s}^{-1}$ . The variability of the quiescent emission poses a challenge for analyzing the surface emission of Aql X-1. To account for these complications, Waterhouse et al. (2016) classified each quiescent observation into one of three states: decay, flare, or quiescence. Observations in the decay state occur within the period of rapid decay immediately following the end of the outburst. Observations in the flare state are those whose spectra have a significant contribution from the hard, power-law component and tend to have higher flux values. Observations in the quiescent state are

those whose spectra are dominated by the soft component. The temperatures from the quiescent observations were estimated using the nsatmos atmosphere model and assumed a distance of  $d = 5$  kpc, a mass of  $M = 1.6 M_{\odot}$ , and a radius of  $R = 11$  km.

I find best-fit parameters of  $T_c = 1.17^{+0.67}_{-1.07} \times 10^8$  K,  $Q_{\text{imp}} = 12.8^{+27.2}_{-12.8}$ ,  $Q_{\text{sh}} = 2.2^{+0.8}_{-1.4}$  MeV u<sup>-1</sup>, and  $Q_{\text{in}} = 19.6^{+20.4}_{-19.6}$  MeV u<sup>-1</sup> with  $\chi^2_{\nu} = 8.2$  (10 dof). A corner plot of the distribution and the cooling curve using the best-fit parameters is shown in Figure 3.7.

Only the shallow heating rate appears normally distributed and fairly well-constrained. The core temperature distribution is left-skewed, with a non-negligible tail extending to the bottom of my allowed temperature values. The impurity is poorly constrained with nearly uniform probability above  $Q_{\text{imp}} \gtrsim 4$ . The inner heating peaks at  $Q_{\text{in}} \approx 16$  MeV u<sup>-1</sup>, but the credible includes the maximum and minimum values allowed in the prior.

Waterhouse et al. (2016) fit several models to these outbursts, assuming  $Q_{\text{imp}} = 1$ . In some models, they fixed the core temperature to  $T_c = 1.75 \times 10^8$  K, based on the lowest temperature ever reported for Aql X-1. Their models varied by whether they included the observations in the flare state and whether they fixed the core temperature. They found shallow heating rates ranging from 1–3 MeV u<sup>-1</sup>, which is consistent with my results.

Ootes et al. (2018) modelled 23 of the observed outbursts of Aql X-1. They fit to the quiescent observation data of 5 of the outbursts, allowing the shallow heating rate and depth of shallow heating to vary across these outbursts (shallow heating was fixed for all other outbursts at  $Q_{\text{sh}} = 1.5$  MeV u<sup>-1</sup> at a depth of  $\rho = 4 \times 10^8$  g cm<sup>-3</sup>). They found that they needed to allow the shallow heating rate, depth of shallow heating, and envelope composition to vary between these outbursts to properly fit the data. The heating properties could not be connected to spectral properties of the outbursts. They also found that there is not enough time between outbursts for the crust to return to equilibrium with the core, so that the crust heating during any outburst affects the cooling curve of subsequent outbursts.



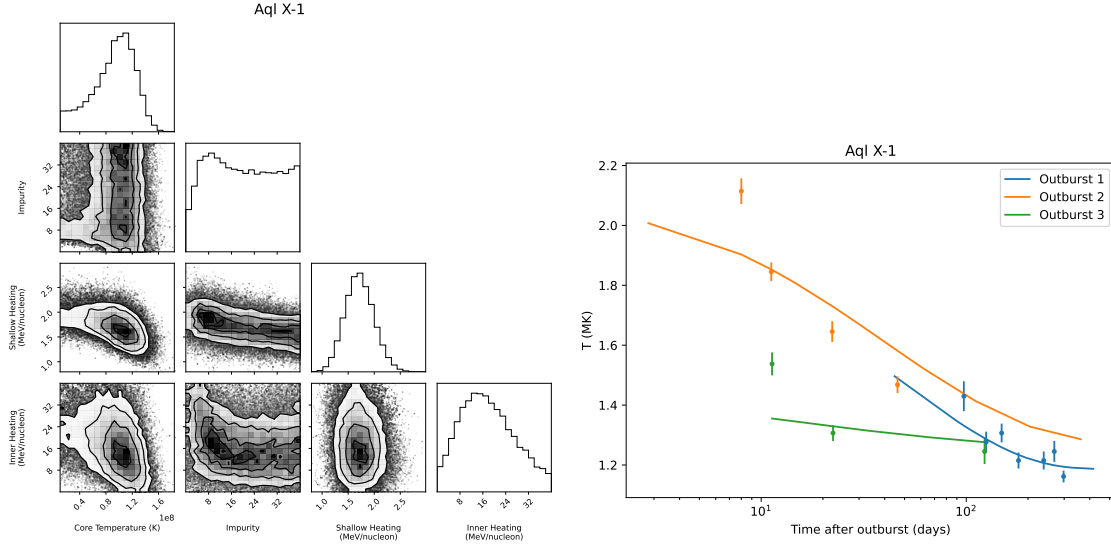


Figure 3.7 Left: Posterior distribution for Aql X-1. Right: Observed surface effective temperatures (points; Waterhouse et al., 2016) and the model with best-fit parameters  $T_c = 1.2 \times 10^8$  K,  $Q_{\text{imp}} = 13$ ,  $Q_{\text{sh}} = 2.2 \text{ MeV u}^{-1}$ , and  $Q_{\text{in}} = 20 \text{ MeV u}^{-1}$ .

### 3.1.8 Swift 174805.3–244637

Swift 174805.3–244637 was discovered in outburst in early July 2012 and was last detected in outburst late August 2012. It was identified as a NS LMXB due to the detection of a Type I X-ray burst during this accretion outburst (Bahramian et al., 2014). Analysis of archival *Chandra* data showed that, prior to the accretion outburst, the source’s X-ray spectrum was dominated by a thermal component with a constant temperature from 2003-2012 which suggests that accretion was either not active or occurred only at a low level (Bahramian et al., 2014).

Degenaar et al. (2015) analyzed quiescent observations up 668 d after the end of the outburst. They assumed a distance of  $d = 5.5$  kpc, a mass of  $M = 1.4 M_{\odot}$ , and a radius of  $R = 10$  km and modelled the atmosphere using *nsatmos*. They include seven pre-outburst quiescent observations, which provides a stronger constraint on the core temperature.

I find best-fit parameters of  $T_c = 1.18 \pm 0.17 \times 10^8$  K,  $Q_{\text{imp}} = 0.0^{+40.0}_{-0.0}$ ,  $Q_{\text{sh}} = 0.9^{+0.7}_{-0.9} \text{ MeV u}^{-1}$ , and  $Q_{\text{in}} = 0.0^{+40.0}_{-0.0} \text{ MeV u}^{-1}$  with  $\chi^2_{\nu} = 0.2$  (8 dof). A corner plot of the distribution and the cooling curve using the best-fit parameters is shown in Figure 3.8.

Both the core temperature and the shallow heating are well constrained. The fit favors a low, but non-zero shallow heating rate. This is consistent with the results of Degenaar et al. (2015),

who found that adding any shallow heating increased the  $\chi^2$  of the fit. The impurity is very poorly constrained, with a nearly uniform probability across my allowed range. A low inner heating rate is favored, though the constraint is not very strong.

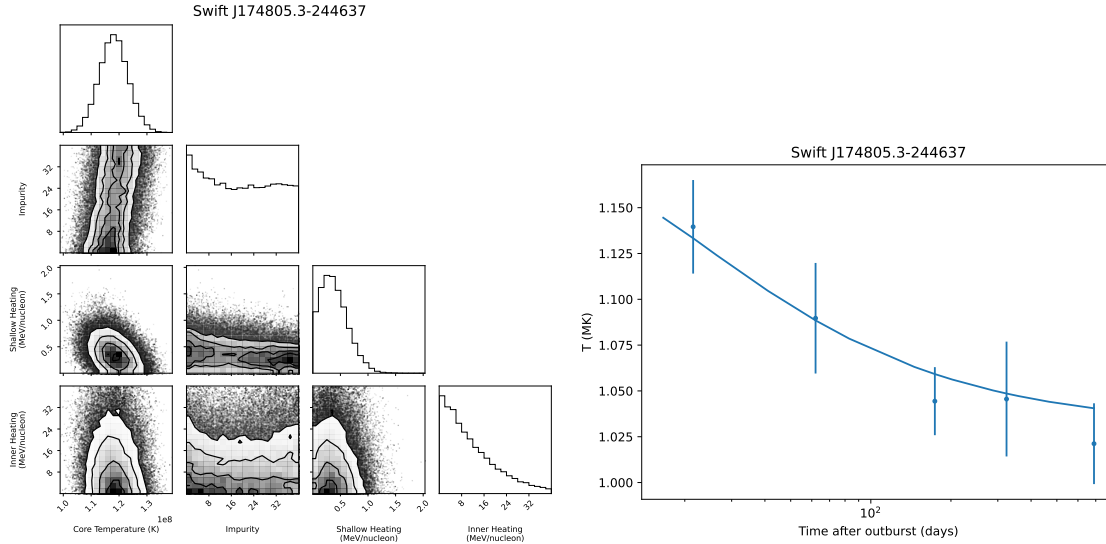


Figure 3.8 Left: Posterior distribution for Swift 174805.3–244637. Right: Observed surface effective temperatures (points; Degenaar et al., 2015) and the model with best-fit parameters  $T_C = 1.2 \times 10^8$  K,  $Q_{\text{imp}} = 0$ ,  $Q_{\text{sh}} = 0.9 \text{ MeV u}^{-1}$ , and  $Q_{\text{in}} = 0.0 \text{ MeV u}^{-1}$ . The pre-outburst temperature observations are not included in this figure, though they were used to fit the model.

### 3.1.9 1RXS J180408.9–342058

1RXS J180408.9–342058 was discovered in 1990 by *ROSAT* (Voges et al., 1999). It was detected again when it briefly went into outburst in 2012 and exhibited a Type I X-ray burst (Chenevez et al., 2012). Based on the brightness of this burst, Chenevez et al. (2012) estimated the upper limit of its distance to be 5.8 kpc. Shortly afterwards it went into quiescence again (Kaur & Heinke, 2012). It then was observed in outburst again in 2015 (Barthelmy et al., 2015a,b; Krimm et al., 2015) and remained in outburst for 4.5 months.

Parikh et al. (2017b) monitored 1RXS J180408.9–342058 during quiescence following its 2015 outburst. They found that the quiescent spectra were dominated by the thermal component, but all observations still required a power-law component to fit. Though their fits show decrease in temperature over time, the non-negligible contribution of the power-law component suggests that low-level accretion may have continued throughout the observed quiescent period. This makes

the interpretation of the observations difficult because we cannot determine whether the apparent decrease in temperature is due to crustal cooling or some other process. Because of this uncertainty, I do not fit a dStar model to the cooling curve data of 1RXS J180408.9–342058.

### 3.1.10 Summary of Independently Fit Cooling Curves

It is likely that my results for EXO 0748–676, MAXI J0556–332, and Aql X-1 are not accurately representative of the crustal heating and cooling that has been observed. In the case of EXO 0748–676, the amount of variability relative to the reported measurement uncertainties is very high, suggesting that we may not be correctly interpreting the temperature data. MAXI J0556–332 and Aql X-1 both exhibited multiple outbursts and the shallow heating rate between outbursts likely varies. The first outburst of MAXI J0556–332 requires an extremely high shallow heating rate. This may be due to a hyperburst, though confirmation requires further investigation. For Aql X-1, the short recurrence time of the outbursts prevents the crust from thermally relaxing to equilibrium with the core, which makes it difficult to estimate the core temperature without modelling many outbursts simultaneously. Additionally, the variability in the spectral state of its quiescent observations makes it difficult to interpret the temperature data.

Table 3.3 gives a summary of the best-fit parameters of all models. Excluding EXO 0748–676, MAXI J0556–332, and Aql X-1, I find good constraints for all core temperatures; for all sources, the IQR is less than 15% of the median value. KS 1731–260 and MXB 1659–29 are the only sources with good constraints on the impurity, and each of these sources favors low values ( $Q_{\text{imp}} \lesssim 4$ ). The impurity is very unconstrained for the remaining sources. The shallow heating rate is well-constrained for all sources and consistent with previous results. For KS 1731–260, MXB 1659–29, and IGR J17480–2446,  $Q_{\text{sh}} \approx 1\text{--}4 \text{ MeV u}^{-1}$ . Both XTE J1701–462 and Swift 174805.3–244637 seem to favor nearly no shallow heating. There is some variation in the predicted values of the inner heating rate. IGR J17480–2446 and Swift 174805.3–244637 favor significantly greater inner heating than the theoretical estimate of  $1.5\text{--}2 \text{ MeV u}^{-1}$ . The middle 50% of the distributions for KS 1731–260, MXB 1659–29, and XTE J1701–462 overlap with the range of theoretical values.

One potential shortcoming of my analysis is that I do not analyze the observational data directly

Table 3.3 Best-fit parameters for all independent models.

Source	$T_c$ ( $10^7$ K)	$Q_{\text{imp}}$	$Q_{\text{sh}}$ ( $\text{MeV u}^{-1}$ )	$Q_{\text{in}}$ ( $\text{MeV u}^{-1}$ )	$\chi^2_\nu$ ( $\chi^2$ )
KS 1731–260	$6.3^{+1.0}_{-0.6}$	$6.3^{+4.2}_{-6.3}$	$1.1^{+1.0}_{-0.6}$	$0.0^{+8.1}_{-0.0}$	7.7 (30.7)
MXB 1659–29	$4.5^{+0.9}_{-1.2}$	$0.0^{+4.7}_{-0.0}$	$1.4^{+0.7}_{-0.2}$	$17.0^{+3.6}_{-17.0}$	3.5 (31.7)
XTE J1701–462	$23.8^{+2.3}_{-6.5}$	$40.0^{+0.0}_{-40.0}$	$0.2^{+0.1}_{-0.0}$	$0.0^{+5.0}_{-0.0}$	5.5 (66.5)
EXO 0748–676	$22.0 \pm 0.2$	$0.0^{+2.4}_{-0.0}$	$3.7^{+0.3}_{-0.4}$	$0.0^{+0.9}_{-0.0}$	111.5 (1004)
IGR J17480–2446	$7.7^{+2.2}_{-1.6}$	$14.3^{+25.6}_{-14.3}$	$1.2^{+0.5}_{-0.7}$	$7.1^{+23.6}_{-7.1}$	1.0 (4.9)
MAXI J0556–332	$1.0^{+23.2}_{-0.0}$	$0.0^{+40.0}_{-0.0}$	$3.9^{+0.4}_{-1.2}$	$17.9^{+0.6}_{-17.9}$	399.4 (5191.6)
Aql X-1	$11.7^{+6.7}_{-10.7}$	$12.8^{+27.2}_{-12.8}$	$2.2^{+0.8}_{-1.4}$	$19.6^{+20.4}_{-19.6}$	8.2 (82.1)
Swift 174805.3–244637	$11.8 \pm 1.7$	$0.0^{+40.0}_{-0.0}$	$0.9^{+0.7}_{-0.9}$	$0.0^{+40.0}_{-0.0}$	0.2 (1.7)

to estimate the quiescent temperatures and quantify the uncertainties myself. Instead, I fit my dStar models to the temperature measurements reported from other sources. Temperature values are not directly measured; counts of X-ray photons are measured. The temperature is then estimated using neutron star atmosphere models which depend on assumptions about the neutron star’s mass, radius, distance, and atmosphere properties and the column of hydrogen between the LMXB and the observing telescope. These quantities may be highly uncertain, but for the purpose of obtaining temperature estimates, they are typically fixed to a specific value. The resulting uncertainties that are reported are due statistical uncertainties of X-ray photon counts, but don’t include uncertainties about the neutron star’s properties. Accounting for the uncertainties in the atmosphere models would lead to systematic errors in the cooling curves. Additionally, there are other possible sources of apparent temperature variation such as low-level accretion flares or eclipses due to the accretion disk that can be difficult to identify. As a result, the temperatures uncertainties are likely smaller than they should be.

To be consistent with the reported temperatures, I use the same neutron star mass, radius, and distance in dStar as the values that were used for the temperature estimates (see Table 3.1). Unfortunately, without doing a more thorough analysis of how all of the observational uncertainties impact the uncertainty in individual temperature measurements, I have no way to determine what more realistic uncertainties should be. The primary effect of fitting to data with artificially small

errors is that the  $\chi^2$  for any set of parameters is higher and the width of my posterior distributions for each parameter is narrower than they otherwise would be.

### 3.2 Testing Whether All Sources Share Crust Properties

In previous studies, as well as in my results above, each source was modeled independently with no assumptions that the shallow heating rates or compositions of the crusts were the same for all neutron stars. To test this assumption, I will model multiple sources under the constraint that they must share the same value of  $Q_{\text{imp}}$ ,  $Q_{\text{sh}}$ , or  $Q_{\text{in}}$ . I will compare the overall goodness-of-fit of the models fit to each source independently (the “independent fits”) to the models which are fit to all sources simultaneously with one parameter varying jointly (the “joint fits”).

My dStar models fit observations of KS 1731–260, MXB 1659–29, XTE J1701–462, IGR J17480–2446, and Swift 174805.3–244637 with reasonable  $\chi^2_{\nu}$ , but are inadequate for describing EXO 0748–676, MAXI J0556–332, and Aql X-1. For my joint fit tests, I will only include the 5 sources that I can fit well independently.

To perform the joint fits, I perform MCMC analysis similar to that described in Section 3.1. For each step of the MCMC, I run one dStar model each for KS 1731–260, MXB 1659–29, XTE J1701–462, IGR J17480–2446, and Swift 174805.3–244637. Three of the dStar parameters for each source vary independently while 1 parameter is shared among all sources. When jointly fitting  $Q_{\text{imp}}$ , for example, I run dStar for each model with independent values of  $T_c$ ,  $Q_{\text{sh}}$ , and  $Q_{\text{in}}$ , but every source has the same value of  $Q_{\text{imp}}$ . The parameter vector can be represented as

$$\boldsymbol{\theta} = (T_c^1, Q_{\text{sh}}^1, Q_{\text{in}}^1, \dots, T_c^5, Q_{\text{sh}}^5, Q_{\text{in}}^5, Q_{\text{imp}}).$$

Here the superscripts 1, . . . , 5 are labels for the sources KS 1731–260, . . . , Swift 174805.3–244637. For each source, I calculate  $\chi^2$  and then calculate a total  $\chi^2$ ,  $\chi_{\text{tot}}^2 = \sum_i \chi_i^2$ , where  $i$  is a label for each source. I calculate the overall likelihood of  $\boldsymbol{\theta}$  as  $\mathcal{L}(\boldsymbol{\theta}) = \exp(-\chi_{\text{tot}}^2/2)$ . The MCMC then proceeds in the same way as for the independent fits by proposing steps in  $\boldsymbol{\theta}$ , then accepting or rejecting steps based on  $\mathcal{L}(\boldsymbol{\theta})$ . I stop the MCMC using the same convergence condition as for the independent fits: when the number of iterations is equal to 100 times the maximum autocorrelation time.

I use the same accretion rates, masses, and radii as for the independent fits (see Table 3.1) and the same priors for all parameters (see Table 3.2), and the same convergence condition. I also find the best-fit parameters using the Nelder-Mead method the same way I did with the independent fits.

When jointly fitting  $Q_{\text{sh}}$  or  $Q_{\text{in}}$ , the procedure is the same. The only difference is that in the parameter vector  $\theta$ ,  $Q_{\text{sh}}$  or  $Q_{\text{in}}$  will be the shared parameter while there will be independent variables representing the other parameters for each source.

### 3.2.1 Model Selection

There are many possible metrics which can be used for comparing the independent fits to the joint fits such as: the likelihood ratio test (Wilks, 1938), the Akaike information criterion (Akaike, 1974), the Bayesian information criterion (Schwarz, 1978), and the Bayes factor (Kass & Raftery, 1995) (see also Hogg et al., 2010, for a review on fitting models to data and quantifying their goodness of fit). Of these, the Bayes factor, which is the ratio of the marginal likelihoods (or evidences,  $P(\mathbf{D})$ ) of the two models, would be the most effective since it accounts for all possible parameters in the parameter space rather than just the best-fit parameters. Estimating the marginal likelihood is computationally expensive, however, which is one of the motivations for using MCMC methods in the first place (see Section 2.4).

Due to the fact that my models don't properly capture all of the uncertainties of the data (as discussed in Section 3.1.10), even the most rigorous statistical methods may not produce reliable criteria for model selection. Because of this, I will simply compare the independent fits to the joint fits by calculating the overall  $\chi^2_\nu$  ( $\chi^2_\nu = \chi^2_{\text{tot}}/\nu$ ) of each set of fits, using their respective best-fit parameters. Across all 5 sources, there are 58 total observations. For the independent fits, there are 20 total parameters (4 for each of the 5 sources) which gives  $\nu = 38$ . For the joint fits, there are 16 total parameters (3 independent parameters for each of the 5 sources and 1 shared parameter) which gives  $\nu = 42$ .

As with the independent fits, I find the best-fit parameters for the joint fits using the Nelder-Mead minimization method. However, it is not practical to create a 16-dimensional histogram to identify the bin with the greatest number of samples to act as the initial guess for minimization. Instead, I

use the MCMC sample with the lowest  $\chi^2$  as the initial guess.

Since my primary objective is to compare the independent fits to the joint fits using the  $\chi^2$  of the best-fit parameters, I only find the best-fit parameters of the joint fits and don't estimate any credible intervals. As a result, I don't quantify any uncertainty on my parameter estimates.

### 3.2.2 Joint Fits Results

The best-fit parameters and  $\chi^2$  for the independent fits and each of the joint fits for KS 1731–260, MXB 1659–29, XTE J1701–462, IGR J17480–2446, and Swift 174805.3–244637 are summarized in Table 3.4. For the independent fits,  $\chi^2_{\text{tot}} = 135.6$  and  $\chi^2_{\nu} = 3.6$  (38 dof). When jointly fitting  $Q_{\text{imp}}$ ,  $\chi^2_{\text{tot}} = 198.0$  and  $\chi^2_{\nu} = 4.7$  (42 dof). When jointly fitting  $Q_{\text{sh}}$ ,  $\chi^2_{\text{tot}} = 776.4$  and  $\chi^2_{\nu} = 18.5$  (42 dof). When jointly fitting  $Q_{\text{in}}$ ,  $\chi^2_{\text{tot}} = 167.1$  and  $\chi^2_{\nu} = 4.0$  (42 dof). The cooling curves for each source and each set of fits are in Figure 3.9.

The  $Q_{\text{imp}}$  and  $Q_{\text{in}}$  joint fits do not fit significantly worse than the independent fits. Most of the best-fit parameters of the joint fits fall within the 68% credible regions of the independent fits for each source. The  $Q_{\text{imp}}$  joint fits favor very low impurity. This is not surprising because the independent fits for most sources either favor very low  $Q_{\text{imp}}$  or do not constrain  $Q_{\text{imp}}$  very well so their data allow low values. Similarly for the  $Q_{\text{in}}$  joint fits, a very low value is favored because the independent fits either favor low values or do not constrain  $Q_{\text{in}}$  very well. The cooling curves of the  $Q_{\text{imp}}$  and  $Q_{\text{in}}$  joint fits do not differ significantly from the cooling curves of the independent fits for most sources.

This is consistent with theoretical calculations of the crust composition which predict that the composition and accretion-induced heating rate in the inner crust do not depend strongly on the initial composition of the matter in the envelope (see, e.g., Lau et al., 2018).

Jointly fitting  $Q_{\text{sh}}$  substantially increases  $\chi^2_{\nu}$  compared to the independent fits. All of the sources fit poorly, with  $\chi^2$  values significantly greater than for the independent fits and the other joint fits. When fit independently, XTE J1701–462 favors a very low value of  $Q_{\text{sh}}$  and all other sources favor  $Q_{\text{sh}} \approx 1 \text{ MeV u}^{-1}$ . However, unlike for  $Q_{\text{imp}}$  and  $Q_{\text{in}}$ , all sources have fairly good constraints on  $Q_{\text{sh}}$  so there is no value of  $Q_{\text{sh}}$  that can fit the data for all sources well simultaneously. As a result,

Table 3.4 Best-fit parameters and  $\chi^2_{\nu}$  for the independent fits and each set of joint fits. The values for  $\chi^2$  are the individual contribution of each source to the overall  $\chi^2_{\text{tot}}$ .

Source	$T_c$ ( $10^7$ K)	$Q_{\text{imp}}$	$Q_{\text{sh}}$ ( $\text{MeV u}^{-1}$ )	$Q_{\text{in}}$ ( $\text{MeV u}^{-1}$ )	$\chi^2$
Independent fits					
KS 1731–260	$6.3^{+1.0}_{-0.6}$	$6.3^{+4.2}_{-6.3}$	$1.1^{+1.0}_{-0.6}$	$0.0^{+8.1}_{-0.0}$	30.7
MXB 1659–29	$4.5^{+0.9}_{-1.2}$	$0.0^{+4.7}_{-0.0}$	$1.4^{+0.7}_{-0.2}$	$17.0^{+3.6}_{-17.0}$	31.7
XTE J1701–462	$23.8^{+2.3}_{-6.5}$	$40.0^{+0.0}_{-40.0}$	$0.2^{+0.1}_{-0.0}$	$0.0^{+5.0}_{-0.0}$	66.5
IGR J17480–2446	$7.7^{+2.2}_{-1.6}$	$14.3^{+25.6}_{-14.3}$	$1.2^{+0.5}_{-0.7}$	$7.1^{+23.6}_{-7.1}$	4.9
Swift 174805.3–244637	$11.8 \pm 1.7$	$0.0^{+40.0}_{-0.0}$	$0.9^{+0.7}_{-0.9}$	$0.0^{+40.0}_{-0.0}$	1.7
$\chi^2_{\nu} = 3.6$ (38 dof)					
$Q_{\text{imp}}$ fit jointly					
KS 1731–260	6.6	0.0	1.6	6.0	43.1
MXB 1659–29	4.5	0.0	1.4	16.8	32.0
XTE J1701–462	22.4	0.0	0.3	2.1	78.8
IGR J17480–2446	8.6	0.0	1.1	18.8	41.6
Swift 174805.3–244637	11.8	0.0	0.9	3.4	2.5
$\chi^2_{\nu} = 4.7$ (42 dof)					
$Q_{\text{sh}}$ fit jointly					
KS 1731–260	6.2	8.0	0.4	2.1	47.3
MXB 1659–29	4.2	0.0	0.4	29.0	432.7
XTE J1701–462	13.7	10.9	0.4	6.3	207.3
IGR J17480–2446	9.0	40.0	0.4	3.8	83.0
Swift 174805.3–244637	12.0	32.4	0.4	0.6	6.0
$\chi^2_{\nu} = 18.5$ (42 dof)					
$Q_{\text{in}}$ fit jointly					
KS 1731–260	6.4	5.2	1.2	0.0	31.8
MXB 1659–29	4.3	2.7	1.7	0.0	39.3
XTE J1701–462	23.8	38.8	0.2	0.0	67.0
IGR J17480–2446	8.5	38.4	1.1	0.0	25.6
Swift 174805.3–244637	12.0	29.0	0.5	0.0	3.5
$\chi^2_{\nu} = 4.0$ (42 dof)					



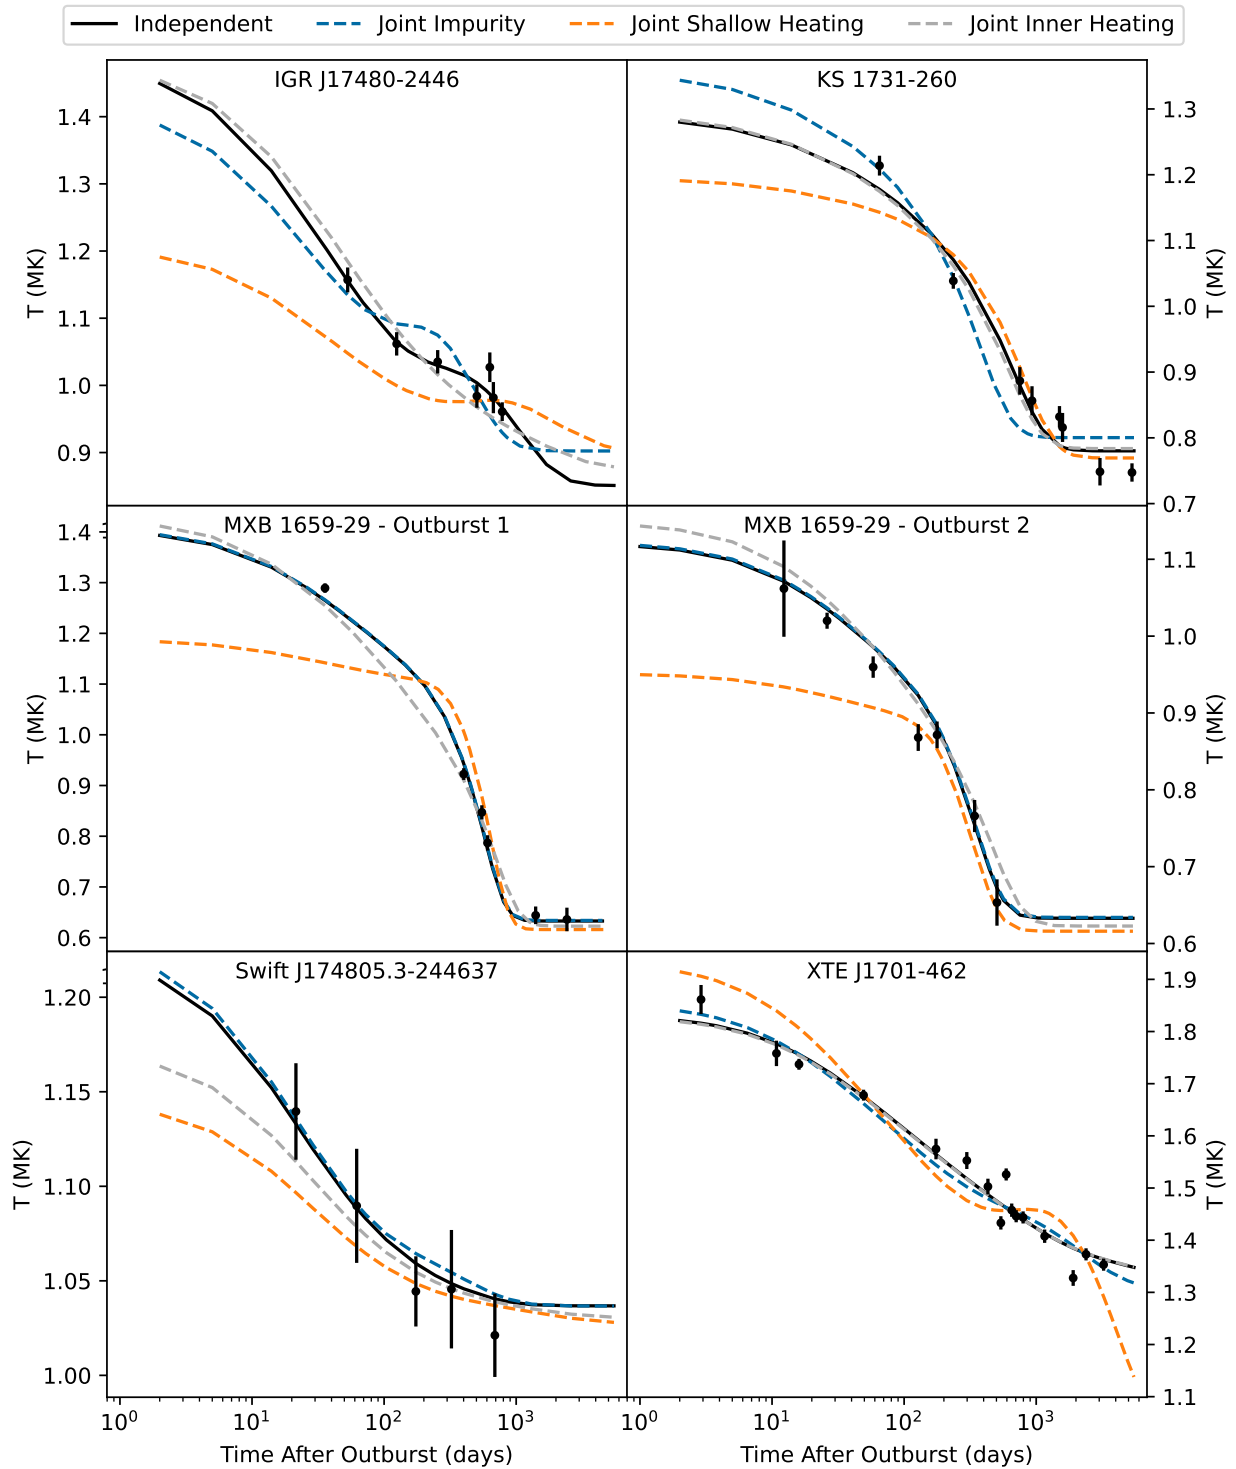


Figure 3.9 Cooling curves of KS 1731–260, MXB 1659–29, XTE J1701–462, IGR J17480–2446, and Swift 174805.3–244637, using the best-fit parameters of the independent and joint fits.

the best-fit value of  $Q_{\text{sh}}$  for the joint fit is outside the 68% credible regions for all sources. This can be seen in the cooling curves where the temperatures in the first  $\sim 100$  days of quiescence deviate significantly from the cooling curves of the independent fits.

This may suggest that the physical process responsible for heating in the shallow layers of the crust can differ between different neutron stars or it may depend on properties of the accretion flow onto the neutron star's surface. If nuclear processes are not the dominant physical process responsible for shallow heating, the heating rate may not actually be proportional to the accretion rate. In this case, the way in which shallow heating has typically been modelled may be incorrect, which could explain why the shallow heating can't be fit jointly across all sources.

### 3.3 Conclusion

I have fit dStar models to observations of KS 1731–260, MXB 1659–29, XTE J1701–462, EXO 0748–676, IGR J17480–2446, MAXI J0556–332, Aql X-1, and Swift 174805.3–244637. Using emcee, I have estimated the posterior probability distributions of the core temperature, impurity of the crustal composition, and the accretion-induced heating rates in the shallow layers of the crust and deep in the inner crust.

I first fit my models to all sources independently, so that my parameter estimates apply only to their respective sources. My models fit well to the cooling curves of KS 1731–260, MXB 1659–29, XTE J1701–462, IGR J17480–2446, and Swift 174805.3–244637 but do not fit well to the cooling curves of EXO 0748–676, MAXI J0556–332, and Aql X-1. MAXI J0556–332 and Aql X-1 have each exhibited multiple outbursts that likely have different shallow heating rates between outbursts. Low-level accretion flares on Aql X-1 and eclipsing due to the accretion disk of EXO 0748–676 make it difficult to determine with a high level of confidence when we observe the surface temperature directly. Accurately modelling these 3 sources will require more sophisticated crustal thermal evolution models than I have used and a better understanding of the accretion properties to properly interpret the observational data.

To test whether the crust properties of accreting neutron stars can be shared by all neutron stars, I perform joint fits to the cooling curves of KS 1731–260, MXB 1659–29, XTE J1701–462,

IGR J17480–2446, and Swift 174805.3–244637. In the joint fits, I fit dStar models to all 5 of these sources simultaneously with the value of  $Q_{\text{imp}}$ ,  $Q_{\text{sh}}$ , or  $Q_{\text{in}}$  being shared across all sources. I compare the goodness-of-fit of each set of joint fits to the independent fits by measuring the  $\chi^2_{\nu}$  of the best-fit parameters of each set.

Jointly fitting  $Q_{\text{imp}}$  and  $Q_{\text{in}}$  does not significantly impact the quality of the fits to the observed cooling curves when compared to the independent fits. This suggests that the neutron stars of these sources either have the same impurity and accretion-induced heating rates in their inner crusts or that these quantities differ from source-to-source, but the data we have available does not constrain the parameters of our models sufficiently to differentiate among sources.

Jointly fitting  $Q_{\text{sh}}$  decreases the quality of the fits significantly for all 5 sources. This suggests that different values of  $Q_{\text{sh}}$  are required to accurately model the crusts of different sources. This may suggest that the physical mechanism responsible for shallow heating differs for different sources or that the heating rate depends on some other properties of the accretion outburst.

It is important to note that further analysis is necessary to draw conclusions with greater confidence. I had to exclude 3 of the 8 LMXBs from my analysis due to poorly fit models. My modelling method is too simplistic to accurately model some systems that have exhibited multiple outbursts and for some systems, it is difficult to interpret whether the observational data is capturing the surface temperature or other features of accretion.

Additionally, the surface temperature measurements to which I compare my models do not fully reflect the true uncertainties in the observational data. The estimated surface temperature values depend on neutron star atmosphere models which require assumptions about the neutron star's mass, radius, distance, and atmosphere properties as well as the density of the hydrogen column between the neutron star and the Earth. These quantities often have high uncertainties, but have been fixed to specific values in order to obtain well-constrained surface temperature estimates. As a result, the reported temperatures have systematic uncertainties that are not reported. Accounting for all the observational uncertainties in the data would likely change the posterior probability distributions of the model parameters.

One way to account for all of the observational uncertainties would be to combine the atmosphere models with the thermal evolution models. For each observation of an X-ray photon count rate and a set of model parameters, the atmosphere model would produce an estimate of the temperature. This temperature estimate would then be compared to the prediction of the thermal evolution model using the same set of parameters. Each time the parameters vary, both the atmosphere model and the thermal evolution model would produce new estimates of the temperature that would be compared.

Because the MCMC may take tens or hundreds of thousands of steps to converge, estimating the temperature from the atmosphere model for each step may not be computationally practical. Instead, it may be possible to use other methods to estimate how the temperature uncertainty depends on the atmosphere model's parameters independent of the thermal evolution models. Bayesian hierarchical models combine sub-models to form a hierarchy of models so that the posterior distribution of the model parameters account for all uncertainties of the sub-models (see Sharma, 2017, for a review of MCMC methods and how to apply Bayesian hierarchical modelling).

## CHAPTER 4

### THE EFFECT OF NEUTRON DIFFUSION ON ACCRETION-INDUCED CRUSTAL HEATING

In Section 2.1, I described how accretion induces nuclear reactions in the crust of a neutron star. These reactions heat the crust and determine its composition. Typically the heat release and composition have been calculated (see, e.g., Haensel & Zdunik, 1990) by tracking the composition and reactions within a closed, Lagrangian fluid element as it is buried deeper in the crust over the course of accretion. These calculations have been performed both by using a one-component approximation (Haensel & Zdunik, 1990, 2003, 2008) and by using reaction networks that allow mixtures from which the impurity can be directly calculated (Steiner, 2012; Lau et al., 2018; Shchechilin & Chugunov, 2019). In these calculations, the reactions are determined under the condition of charge neutrality and conservation of baryons in a Lagrangian fluid element. In the outer crust, we expect the mass number of nuclei to remain constant because the nuclear reactions are limited to electron capture reactions. At neutron drip, neutron emissions and neutron captures may occur locally so the number of baryons stays constant within a fluid element. In the inner crust, however, the local enforcement of baryon conservation implies that the free neutrons in a fluid element move with the nuclei in the lattice.

#### 4.1 Neutron Diffusion

Except for a narrow layer at the outer-inner crust boundary and possibly near the crust-core transition, we expect the free neutrons in the inner crust to be superfluid (see, e.g., Chamel & Haensel, 2008). The superfluid neutrons move with a velocity  $v_{\text{sn}}$ , governed by the linearized superfluid equation (Schmitt, 2015)  $m_n \partial v_{\text{sn}} / \partial t = -\nabla \mu_n^\infty$ , where  $\mu_n^\infty \equiv \mu_n \exp(\phi/c^2)$  is the redshifted neutron chemical potential and  $\phi$  is the gravitational potential. Hydrostatic equilibrium implies that  $\mu_n^\infty = \text{const}$  in the whole superfluid region. Gusakov & Chugunov (2020) argue that conserving the baryon number in the WS cell violates this condition (see also Chugunov & Shchechilin, 2020; Gusakov & Chugunov, 2021; Shchechilin et al., 2021; Gusakov et al., 2021) and instead propose an EOS that is consistent with neutron hydrostatic and diffusion equilibrium (which they call “nHD”).

They find that in the inner crust, the nHD EOS produces a crust that is very similar in composition to the catalyzed crust.

Using the nHDEOS, Gusakov & Chugunov (2021) calculated the accretion-induced heat release to be  $0.21\text{--}0.53 \text{ MeV u}^{-1}$ , a factor of a few less than the  $1.5\text{--}2.0 \text{ MeV u}^{-1}$  predicted by the non-diffusive models. Additionally, some nuclear physics models predict the existence of a pasta phase at the bottom of the inner crust (Pethick & Potekhin, 1998; Chamel & Haensel, 2008). Molecular dynamics simulations predict that a pasta layer should have low electrical and thermal conductivity due to its disordered structure (Horowitz et al., 2015). The inner heating rate and the thermal conductivity (which depends on the impurity of the composition) impact the thermal evolution of the crust of an accreting neutron star. This means that whether a thermal evolution model uses the heat release and composition of the nHD model or of a non-diffusive model and whether a pasta layer is included may produce observationally distinguishable cooling curves.

In this chapter, I use `dStar` to compare the nHD model to non-diffusive models. I perform model selection for each source by estimating the Bayes factor using nested sampling to determine whether the observational data favors models the nHD model or a non-diffusive model. I also compare models that include a pasta layer to those that don't include a pasta layer. For each nHD and non-diffusive model, I perform model selection to determine whether the data favors the pasta or the non-pasta models.

## 4.2 Incorporating Neutron Diffusion Into Cooling Models

Using methodology similar to that described in Section 3.1, I model the observed cooling of KS 1731–260, MXB 1659–29, XTE J1701–462, IGR J17480–2446, and Swift 174805.3–244637. I use the same neutron star masses, radii, distances, and accretion rates and fit to the same observed temperature data (see Table 3.1). I exclude EXO 0748–676, MAXI J0556–332, and Aql X-1 from my analysis for the reasons I described in Section 3.1.10. I create one set of `dStar` models for the nHD crust, and three sets of models for non-diffusive crusts, each corresponding to different initial conditions for the non-diffusive crust calculations. For all models,  $T_c$  and  $Q_{sh}$  are free parameters.  $Q_{in}$  and  $Q_{imp}$  are set according to the calculated values of the respective models.

For the nHD crust dStar models, I set  $Q_{\text{in}} = 0.37 \text{ MeV u}^{-1}$ . In the outer crust, the heat released from electron capture reactions is  $0.3 \text{ MeV u}^{-1}$  for  $26.86 \leq \log[P/(\text{dyn cm}^{-2})] \leq 30.13$ . As predicted by Gusakov & Chugunov (2021), nearly all the heat due to deep crustal heating is released very near neutron drip. I set the depth of the inner heating to occur at a pressure range of  $29.9 \leq \log[P/\text{dyn cm}^{-2}] \leq 30.0$ . It is a good assumption that the catalyzed crust of a neutron star is highly pure because the matter is allowed to reach its absolute ground state (see, e.g., Baym et al., 1971; Baym et al., 1971). Since the nHD crust composition is similar to that of the catalyzed crust, I set  $Q_{\text{imp}} = 0$ .

For the non-diffusive dStar models, I use compositions and heating profiles calculated by the reaction network of Lau et al. (2018). The composition, impurity, and heat release as a function of depth depend on the initial composition of the WS cell which represents matter at the top of the crust. I use the results of three different initial compositions<sup>1</sup>:  $^{56}\text{Fe}$ , superburst ashes, and extreme rp-process ashes (based on the X-ray burst model of Schatz et al., 2001).

The reaction network calculations yield the nuclear composition, impurity, and cumulative heat release per accreted nucleon as a function of pressure (which can be treated as a proxy for depth). The calculations are performed up to a density of  $6.36 \times 10^{12} \text{ g cm}^{-3}$  for the initial  $^{56}\text{Fe}$  and superburst ash compositions and  $5.43 \times 10^{12} \text{ g cm}^{-3}$  for the initial rp-process ashes. Above those densities, I use the composition of Haensel & Zdunik (1990). The mean charge and mass numbers and impurity of each reaction network as functions of pressure are shown in Figure 4.1.

For each reaction network calculation, I load the composition and impurity directly into their respective dStar models. Each reaction that releases heat could be treated as a very narrow heating zone at the precise pressure where the reaction occurs. However, there are hundreds of heating zones for each reaction network and the computational time increases with the number of heating zones.

To speed up the computation, I approximate the heat release by dividing the crust into three different heating zones where the heat is released uniformly over the whole zone. For each reaction

---

<sup>1</sup>Results of calculations received from Rahul Jain, private communication




Figure 4.1 Profiles of mean charge number (top), mean mass number (middle), and impurity (bottom) of initial compositions of pure <sup>56</sup>Fe, superburst ashes, and rp-process ashes.



network, I create a piecewise function  $Q(P)$  to represent the cumulative heat release as a function of pressure. The domain of  $Q(P)$  is the same as the domain of pressure in the reaction networks.  $Q(P)$  is continuous and its value at the maximum  $P$  in the domain is equal to the cumulative energy calculated in the reaction network. Within 3 subdomains (the heating zones),  $dQ/dP$  is constant and positive; outside of the heating zones,  $Q(P)$  is constant. I choose heating zone subdomains and heating rates (quantified by  $dQ/dP$  within the heating zones) such that the root-mean square error relative to the reaction network cumulative heat release is minimized.

Figure 4.2 shows my approximations of  $Q(P)$  compared to the reaction network calculations for each of the three initial compositions. Note that I have plotted the heating profile on a  $\log P$  scale so that it is easier to visualize, so  $Q(P)$  does not have a constant slope in the heating zones. The heating rates and boundaries of the heating zones for  $Q(P)$  are summarized in Table 4.1.

Table 4.1 The heating rates per accreted nucleon and pressure ranges for each heating zone. In each heating zone, heat is released uniformly over the domain  $P_{\text{low}} \leq P \leq P_{\text{high}}$ .

Initial Composition	Heating Rate (MeV u <sup>-1</sup> )	$\log P_{\text{low}}$ (dyn cm <sup>-2</sup> )	$\log P_{\text{high}}$ (dyn cm <sup>-2</sup> )
<sup>56</sup> Fe	0.17	27.90	28.11
	0.37	28.39	30.62
	1.56	30.63	31.21
Superburst	0.12	27.48	28.22
	0.38	28.22	30.61
	1.59	30.63	31.20
Rp-Process	0.13	27.36	27.71
	0.40	28.16	30.50
	1.22	30.50	31.12

#### 4.2.1 Adding a Pasta Layer

Horowitz et al. (2015) model nuclear pasta using a molecular dynamics simulation to estimate the effective impurity at a baryon density of  $n = 0.05 \text{ fm}^{-3}$ , which is about 1/3 nuclear saturation density and is a typical density where nuclear pasta is expected. They estimate that  $Q_{\text{imp}} \approx 40$ , but note that there is a large degree of uncertainty.

I create two sets of models for the nHD crust and for each reaction network calculation: one without a pasta layer and one with a pasta layer. For the set of models with no pasta layer  $Q_{\text{imp}}$

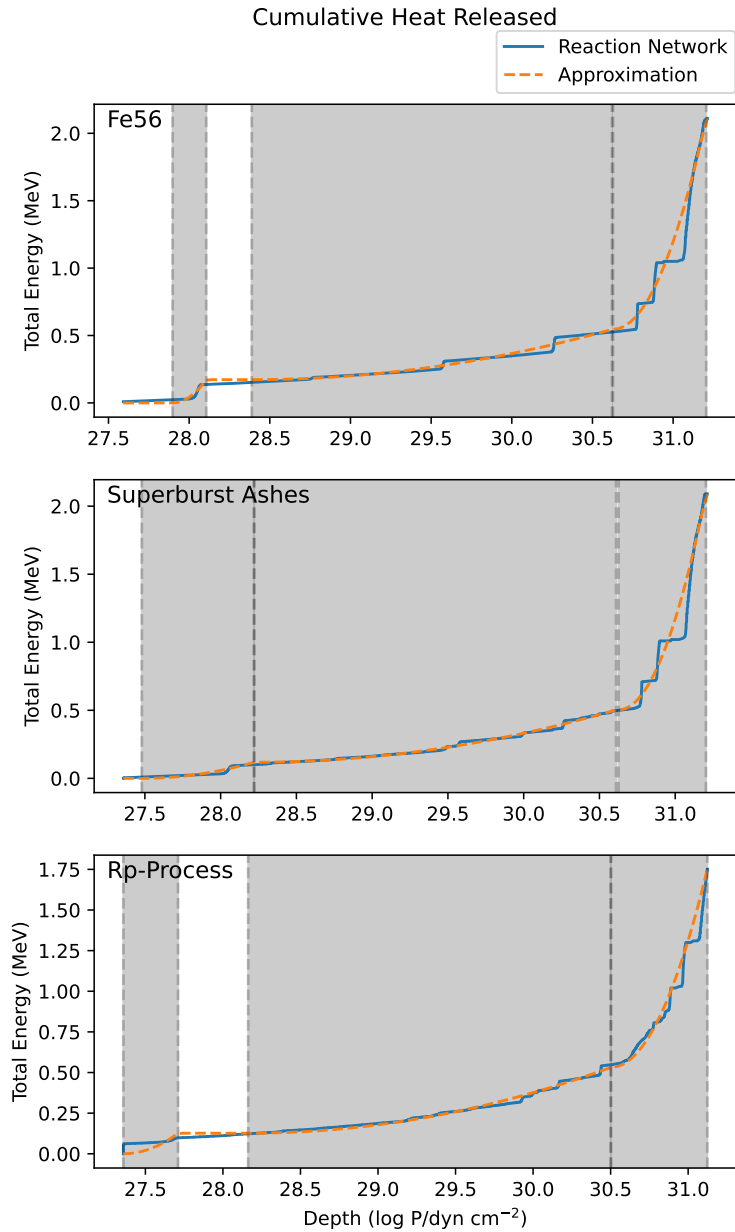


Figure 4.2 Cumulative heat release in the crust based on the reaction network calculations of Lau et al. (2018) for initial compositions of pure  $^{56}\text{Fe}$  (top), superburst ashes (middle), and rp-process ashes (bottom). The solid blue lines are the cumulative energy released up to a depth of  $\log P$  calculated by the reaction network. The dashed yellow lines are my approximation of the heat release where heat is released uniformly over three heating zones. The shaded regions indicate the locations of the heating zones.

throughout the crust is determined by the reaction network calculations or is fixed at  $Q_{\text{imp}} = 0$  for the nHD crust. For the set of models with a pasta layer, at baryon densities greater than  $n = 0.05 \text{ fm}^{-3}$ , the impurity is equal to the pasta impurity,  $Q_{\text{imp,pasta}}$ . Because of the uncertainty in the value of  $Q_{\text{imp,pasta}}$ , I leave it as a free parameter with a uniform prior ranging from 20–200. This allows the pasta impurity to vary significantly around the estimate of  $Q_{\text{imp,pasta}} \approx 40$ , but constrains it to remain at a level much higher than is expected in the rest of the crust.

### 4.3 Model Selection

For each source, I select from three sets of competing models: first, the three non-diffusive reaction network models against each other; second, the non-diffusive models against the nHD models; and third, models without a pasta layer against models with a pasta layer. To select between competing models, I will calculate the Bayes factor (Kass & Raftery, 1995). For two competing models, the Bayes factor is the ratio of the marginal likelihoods of the two models. As introduced in Section 2.4, evidence is the marginal likelihood of all possible model parameters,  $\mathcal{Z} = \int_{\Omega_{\Theta}} \mathcal{L}(\theta)\pi(\Theta)d\Theta$ , where  $\Omega_{\Theta}$  represents the whole parameter space. This provides more information about the overall fits of competing models than many model selection methods which only account for the maximum likelihoods of the competing models (for example, the likelihood ratio test, the Akaike information criterion, or the Bayesian information criterion). One disadvantage of the Bayes factor is that it is usually impossible to calculate directly and is typically very difficult to approximate numerically.

Most commonly used MCMC methods, such as the Metropolis-Hastings algorithm and its variants, are effective at estimating the posterior probability distribution of a model given a set of data but do not estimate  $\mathcal{Z}$ , which is the quantity we need for estimating the Bayes factors for competing models. To estimate  $\mathcal{Z}$ , I use another type of MCMC sampling method called “nested sampling” (Skilling, 2004).

Estimating  $\mathcal{Z}$  by numerically integrating  $\mathcal{L}(\theta)\pi(\Theta)d\Theta$  over the parameter space is impractical for more than a small number of dimensions due to the curse of dimensionality. However, we can

reformulate the integral by introducing the prior mass,

$$X(\lambda) \equiv \int_{\Theta: \mathcal{L}(\theta) > \lambda} \pi(\Theta) d\Theta. \quad (4.1)$$

$X(\lambda)$  is the mass of the prior in the parameter space for which  $\mathcal{L}(\theta) > \lambda$  and decreases from 1 to 0 as  $\lambda$  increases from 0 to  $\infty$ . Using  $dX = \pi(\Theta) d\Theta$ , we can evaluate the evidence,

$$\mathcal{Z} = \int_{\Omega_{\Theta}} \mathcal{L}(\theta) \pi(\Theta) d\Theta = \int_0^1 \mathcal{L}(X) dX, \quad (4.2)$$

where  $\mathcal{L}(X) \equiv \mathcal{L}(X(\lambda)) = \lambda$ . By evaluating iso-likelihood contours  $\mathcal{L}_i \equiv \mathcal{L}(X_i)$  for prior masses ranging from 0 to 1, we can numerically integrate Equation (4.2).

The complete nested sampling algorithm and its mathematical justification is beyond the scope of this chapter (see Skilling, 2004), but works by estimating the prior masses of iso-likelihood contours based on samples from the prior. Dynamic nested sampling is a generalization of the nested sampling algorithm that improves sampling efficiency for more accurate estimates of both the evidence and the posterior (Higson et al., 2019).

I use the code `dynesty` to estimate  $\mathcal{Z}$  for each source and each crust model. For two competing models, I calculate the Bayes factor  $K = \mathcal{Z}_1 / \mathcal{Z}_2$ . If  $K > 0$ , the data favors model 1 over model 2, with larger values of  $K$  indicating stronger favor towards model 1 over model 2. The thresholds of  $\log K$  for model selection are in Table 4.2.

Table 4.2 Thresholds of  $\log K$  for model selection and hypothesis testing suggested by Kass & Raftery (1995).

$\log K$	Strength of Evidence
0 to 0.5	Not worth more than a bare mention
0.5 to 1	Substantial
1 to 2	Strong
> 2	Decisive

#### 4.4 Comparison of Models

There are eight total models for each source: four models with no pasta layer for each of the crust compositions ( $^{56}\text{Fe}$ , superburst ashes, rp-process ashes, and nHD) and four models with a pasta layer for each of the crust models. The estimated evidence for each model, calculated using

dynesty, is in Table 4.3. Cooling curves for nHD models and the most probable non-diffusive model for each source are shown in Figure 4.3. The posterior distributions of each model can be found in the Appendix.

Table 4.3 Estimated evidence  $\ln \mathcal{Z}$  for each of the non-diffusive and nHD models. The greatest evidence (most favored model) for each source is in bold.

Source	Crust Model	No Pasta	Pasta
KS 1731–260	$^{56}\text{Fe}$	$-25.410 \pm 0.062$	$-28.165 \pm 0.104$
	Superburst ashes	<b><math>-23.334 \pm 0.056</math></b>	$-27.741 \pm 0.084$
	Rp-process ashes	$-25.314 \pm 0.063$	$-43.832 \pm 0.084$
	nHD	$-27.011 \pm 0.067$	$-28.977 \pm 0.079$
MXB 1659–29	$^{56}\text{Fe}$	$-26.855 \pm 0.065$	<b><math>-22.269 \pm 0.069</math></b>
	Superburst ashes	$-27.687 \pm 0.060$	$-24.593 \pm 0.064$
	Rp-process ashes	$-113.806 \pm 0.064$	$-128.232 \pm 0.079$
	nHD	$-34.952 \pm 0.068$	$-37.470 \pm 0.070$
XTE J1701–462	$^{56}\text{Fe}$	$-60.591 \pm 0.074$	$-62.126 \pm 0.079$
	Superburst ashes	$-62.613 \pm 0.071$	$-64.188 \pm 0.076$
	Rp-process ashes	$-57.268 \pm 0.075$	$-58.046 \pm 0.077$
	nHD	<b><math>-49.793 \pm 0.254</math></b>	$-50.222 \pm 0.254$
IGR J17480–2446	$^{56}\text{Fe}$	$-33.251 \pm 0.069$	$-30.013 \pm 0.070$
	Superburst ashes	$-31.166 \pm 0.060$	$-28.407 \pm 0.064$
	Rp-process ashes	$-24.816 \pm 0.067$	<b><math>-23.514 \pm 0.069</math></b>
	nHD	$-35.850 \pm 0.226$	$-34.658 \pm 0.229$
Swift 174805.3–244637	$^{56}\text{Fe}$	$-9.624 \pm 0.062$	$-9.838 \pm 0.063$
	Superburst ashes	<b><math>-8.402 \pm 0.057</math></b>	$-8.478 \pm 0.058$
	Rp-process ashes	$-10.198 \pm 0.063$	$-10.368 \pm 0.064$
	nHD	$-9.971 \pm 0.216$	$-9.765 \pm 0.214$

For any given source, the most favored model yields the greatest value of  $\ln \mathcal{Z}$ . Any pair of models can be compared by calculating  $\ln K = \ln \mathcal{Z}_1 - \ln \mathcal{Z}_2$ ; when using the model selection thresholds from Table 4.2,  $\ln \mathcal{Z}_1$  is the more favorable of the two models. There are three sets of model comparisons I would like to discuss with these sources. First, I will compare the non-diffusive models to each other. The initial compositions of the reaction network calculations represent the composition at the top of the crust during accretion. This depends on the types of thermonuclear burning that have occurred in the envelope and may depend on the accretion history. Second, I will compare the non-diffusive models to the nHD models. Third, I will compare the pasta models with the no-pasta models.

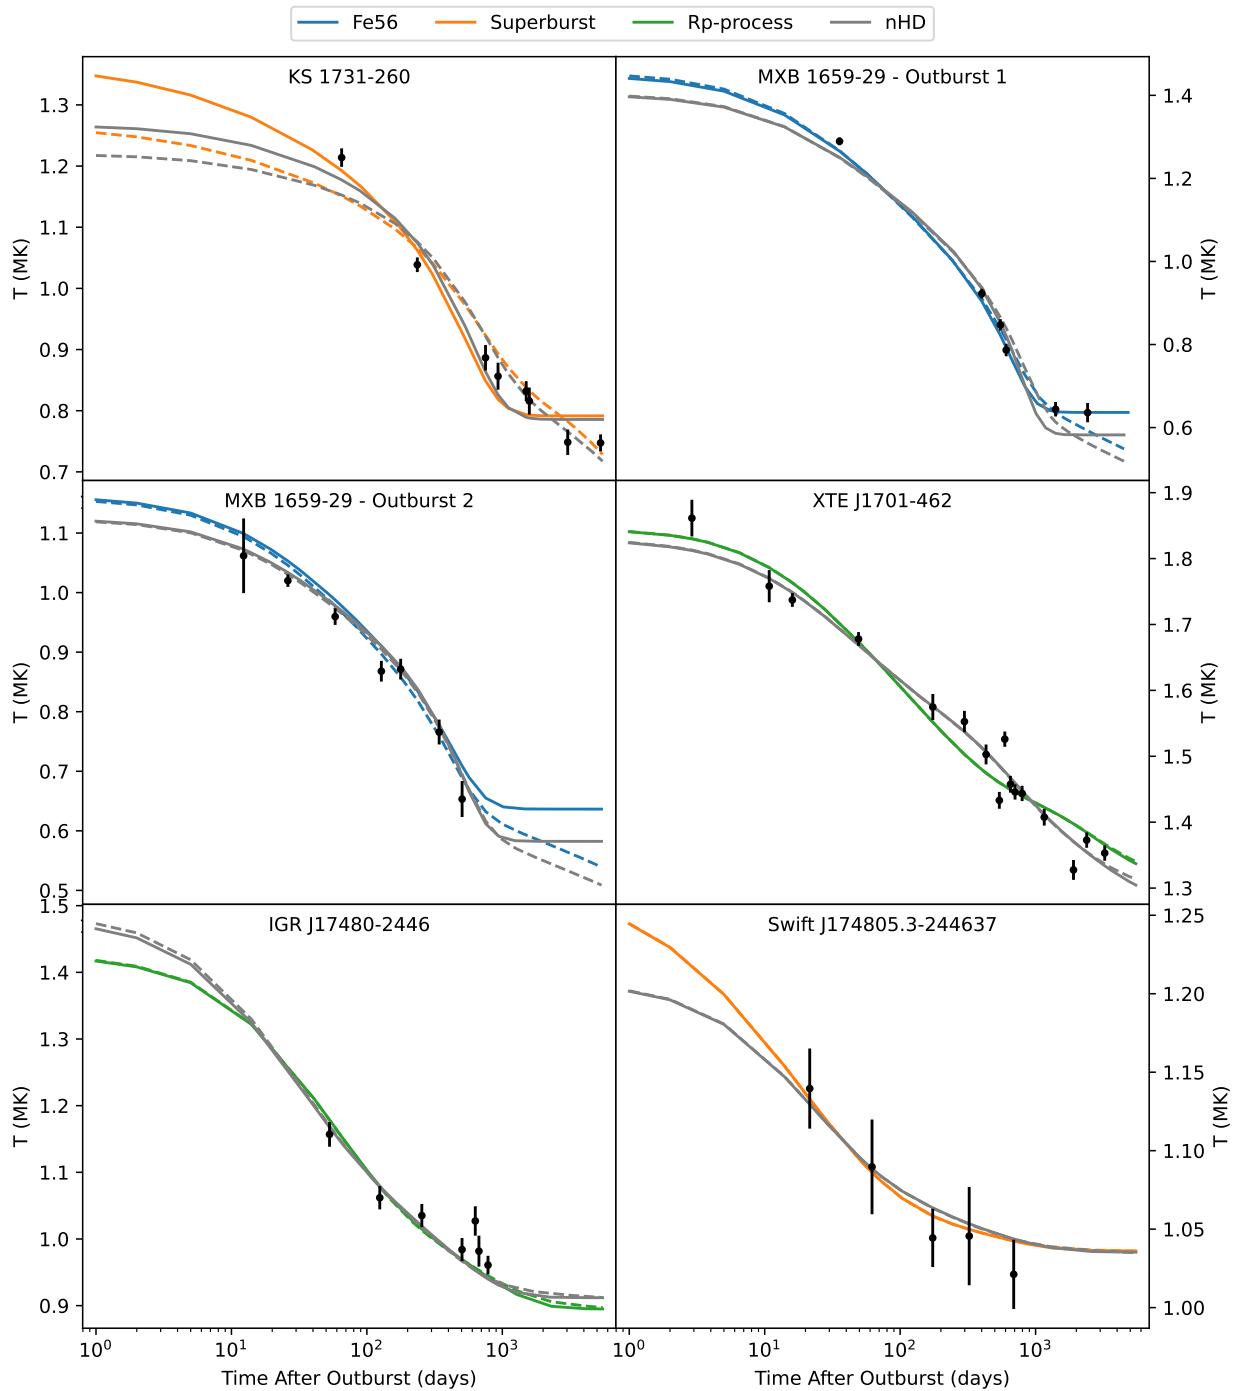


Figure 4.3 Cooling curves of KS 1731–260, MXB 1659–29, XTE J1701–462, IGR J17480–2446, and Swift 174805.3–244637, using the best-fit parameters of the nHD models and the non-diffusive models. Only the cooling curve of the non-diffusive model with the greatest evidence is shown. Solid lines represent models with no pasta layer; dashed lines represent models with a pasta layer.

There does not appear to be one non-diffusive model that is favored consistently by all the sources. KS 1731–260 and Swift 174805.3–244637 favor the initial superburst ashes composition; MXB 1659–29 favors the initial  $^{56}\text{Fe}$  composition; XTE J1701–462 and IGR J17480–2446 favor the initial rp-process ashes composition. The preferences among the non-diffusive models for each source are consistent between the pasta and no-pasta models. We don't necessarily expect different neutron stars to have the same crust composition since the accretion history varies from source-to-source. So it is not unexpected for different sources to favor different reaction network compositions.

However, when comparing the pasta models to the no-pasta models, we would expect all sources to favor one or the other. Whether or not pasta phases exist depends on the physics of the dense matter at the bottom of the crust, which should not differ between neutron stars. However, the sources here do not consistently favor either pasta or no-pasta models. KS 1731–260 and XTE J1701–462 favor no pasta layer for all crust compositions (with only a slight preference for XTE J1701–462 with the rp-process and nHD compositions). IGR J17480–2446 favors a pasta layer for all crust compositions. MXB 1659–29 favors a pasta layer for the  $^{56}\text{Fe}$  and superburst ashes compositions and no pasta layer otherwise. The differences between the no-pasta and pasta models for each composition for Swift 174805.3–244637 are too small to favor one over the other.

Similarly, when comparing the non-diffusive models to the nHD models, we do expect all sources to favor one or the other. This is because the difference between the non-diffusive models and the nHD models is in how the free neutrons are treated. A comparison between the sets of models is, in principle, a test of which treatment better captures the behavior of the free neutrons in the inner crust. Regardless of the inclusion of a pasta layer, XTE J1701–462 favors the nHD models, while the other sources all favor one of the non-diffusive models. In the cases of KS 1731–260 with a pasta layer and Swift 174805.3–244637 with or without pasta, the preference for the superburst ashes composition is not decisive ( $\log K < 2$ ).

One interesting feature of the pasta models is that for KS 1731–260 and MXB 1659–29, the data appears to favor very low  $T_c$  compared to the no-pasta models (see the posterior distributions

in Figures A.2 and A.4). The no-pasta models predict that the crusts have finished cooling and have returned to thermal equilibrium with the cores, while the pasta models predict that the crusts are still cooling (see Figure 4.3). Further X-ray observations could help determine if these two sources are still cooling which could be an effective test of the existence of a pasta layer. If KS 1731–260 and MXB 1659–29 did in fact have significantly cooler cores, this would imply a larger inferred specific heat capacity (Cumming et al., 2017) and/or a more efficient neutrino emissivity (Brown et al., 2018).

#### 4.5 Conclusion

I have compared neutron star crust models with different crust compositions and heating profiles using the observed cooling of KS 1731–260, MXB 1659–29, XTE J1701–462, IGR J17480–2446, and Swift 174805.3–244637. I have also compared models with a low thermal conductivity pasta layer at the bottom of the crust to models with no pasta layer. I do not find any consistent preference across all sources for any set of models. Different accreting neutron stars have different accretion histories and exhibit different types of thermonuclear burning in their envelopes, so we expect that the composition at the top of the crust may differ from source-to-source. However, when comparing non-diffusive models to the nHD model or the pasta models with the no-pasta models, we expect the data for all sources to prefer one or the other.

The cause for the unexpected results may be due to the fact that we are not fully accounting for all the uncertainties in the data (as discussed in Sections 3.1.10 and 3.3). The observed  $T_{\text{eff}}$  values to which I compare my models may have significant systematic errors that I do not account for and their random errors are likely greater than reported. Properly quantifying these errors would change the estimated  $\mathcal{Z}$  values of each of my tested models and may change which models are favored by the data. Alternatively, it may lead to estimated  $\mathcal{Z}$  values that are too close to distinguish between different models. Coming to a proper conclusion will require more accurately quantifying the observational errors.



## CHAPTER 5

### CONCLUSIONS

Accreting neutron stars provide a window into the physics of the dense matter that is present in a neutron star. By fitting models of the thermal evolution of the crust to the cooling observed during quiescence following an accretion outburst, we can estimate the core temperatures, accretion-induced heating rates, and impurity of the composition. Using Bayesian MCMC methods, we can approximate the posterior probability distributions of model parameters and compare competing models.

#### 5.1 Jointly Fitting Crustal Properties of Multiple Neutron Stars

There are nine LMXBs in which crustal cooling has been observed following an accretion outburst: KS 1731–260, MXB 1659–29, XTE J1701–462, EXO 0748–676, IGR J17480–2446, MAXI J0556–332, Aql X-1, Swift 174805.3–244637, and 1RXS J180408.9–342058. Several of these sources have been analyzed independently prior to this work. Joint fits for multiple sources, in which it is assumed that multiple sources share the same heating rates or compositional impurity, had never been performed before. It was unknown whether different neutron star crusts share the same properties or if they differ based on other factors such as their accretion history.

I fit `dStar` models to the cooling curves of all of these sources except 1RXS J180408.9–342058. All quiescent observations of 1RXS J180408.9–342058 exhibit spectral properties that suggest that low-level accretion may be occurring, so we can't reliably interpret the observational data as thermal emission from the neutron star's surface. For the remaining sources, I estimated the posterior probability distributions of  $T_c$ ,  $Q_{\text{imp}}$ ,  $Q_{\text{sh}}$ , and  $Q_{\text{in}}$  using `emcee`. I found that my models fit poorly to the observed cooling of EXO 0748–676, MAXI J0556–332, and Aql X-1, as measured by the  $\chi^2_{\nu}$  values for the best-fit parameters. EXO 0748–676 exhibits a high degree of variability relative to the reported temperature uncertainty, which may be due to interference from its accretion disk. Both MAXI J0556–332 and Aql X-1 have exhibited multiple accretion outbursts with shallow heating rates that likely vary between outbursts. Additionally, low-level accretion was likely present for many of the quiescent observations of Aql X-1.

My dStar models fit well the cooling curves of KS 1731–260, MXB 1659–29, XTE J1701–462, IGR J17480–2446, and Swift 174805.3–244637 so I use these sources for all remaining tests. I performed joint fits where all sources share the same value of  $Q_{\text{imp}}$ ,  $Q_{\text{sh}}$ , or  $Q_{\text{in}}$ . I compared the goodness-of-fit of these joint fits to the overall goodness-of-fit of these sources when fit independently. I found that jointly fitting  $Q_{\text{imp}}$  or  $Q_{\text{in}}$  does not significantly decrease the goodness-of-fit, which suggests that all sources may have the same impurity and inner heating rates. However, I found that jointly fitting  $Q_{\text{sh}}$  significantly decreases the goodness-of-fit. This suggests that different accreting neutron stars very likely have different shallow heating rates. The physical process responsible for the shallow heating is unknown, but its variability suggests that it is likely not primarily due to nuclear processes.

## 5.2 The Effect of Neutron Diffusion on Accretion-Induced Crustal Heating

Most calculations of the composition and deep crustal heating rates do not account for the diffusion of neutrons throughout the inner crust. Gusakov & Chugunov (2020) argue that because the unbound neutrons are superfluid, they should diffuse easily and that this should reduce the heating rate and drive the composition closer to that of the catalyzed crust. I made dStar models using the calculated heat release and impurity of three non-diffusive reaction network calculations using different initial compositions (Lau et al., 2018) and the nHD calculation of Gusakov & Chugunov (2021). For each of these heat release and impurity calculations, I made one model with a low-conductivity pasta layer at the bottom of the inner crust and one model with no pasta layer. I fit these models to the cooling curves of KS 1731–260, MXB 1659–29, XTE J1701–462, IGR J17480–2446, and Swift 174805.3–244637 and calculated the Bayesian evidence  $\mathcal{Z}$  using *dynesty*. I compared the estimated values of  $\mathcal{Z}$  for each model to determine which models are preferred by the data for each source.

I found that there are no models that are generally preferred across all sources. It is not surprising to find that the data from different sources favor different initial compositions among the non-diffusive models. The initial composition in the reaction network calculation represents the composition at the top of the crust, which we expect to vary depending on the history of accretion

and thermonuclear burning in the envelope.

However, we expected to find that all sources favored either a non-diffusive model or the nHD model. Whether or not free neutrons in the inner crust diffusive significantly should depend on the underlying physics of neutrons in the high density regime and should not vary for different neutron stars. Similarly, we expected all sources to favor either a model with a pasta layer or a model with no pasta layer. The existence of pasta phases at the bottom of the crust depends on the EOS of the matter just below nuclear saturation density, which should be different for different neutron stars.

### **5.3 Future Work**

The results I have reported in this work are subject to several shortcomings that could be improved in future work. I was only able to fit my dStar models to five of the nine sources well, which is particularly problematic when attempting to determine whether some crust properties are global or differ for different neutron stars. Fitting to MAXI J0556–332 and Aql X-1 requires varying the shallow heating rate for different outbursts (which has been done by, e.g., Parikh et al., 2017a; Ootes et al., 2018; Degenaar et al., 2019). Properly interpreting the quiescent observations of Aql X-1 and 1RXS J180408.9–342058 requires properly distinguishing between surface emission and low-level accretion or other disk emission. This may also necessary for the quiescent emission of EXO 0748–676, which exhibits X-ray eclipses and high estimated temperature variability.

It will also be necessary to more accurately account for all observational uncertainties when fitting models to available data. In my work and all prior analysis of quiescent crust cooling, thermal evolution models have been fit to observed surface temperatures that have been estimated from neutron star atmosphere models based on assumptions about the neutron star’s distance, mass, and radius, and the intervening column of hydrogen between the neutron star and the Earth. The values of each of these parameters are uncertain, but the temperature estimates and their corresponding uncertainties all assume precise values. By fitting thermal evolution models to these temperatures, we are only able to estimate our model parameters and quantify the goodness-of-fit on the condition that the parameters assumed in the surface temperature calculations are true.

Future analysis of cooling curves should quantify the uncertainties in the estimates of the surface

temperature. One method that may be effective for this is Bayesian hierarchical modelling, where we combine sub-models to form a hierarchical model that accounts for the uncertainty in each sub-model (see, e.g., Sharma, 2017).

## BIBLIOGRAPHY

- Akaike, H. 1974, IEEE Transactions on Automatic Control, 19, 716, doi: 10.1109/TAC.1974.1100705
- Altamirano, D., van der Klis, M., Wijnands, R., & Cumming, A. 2008, ApJ, 673, L35, doi: 10.1086/527355
- Altamirano, D., Keek, L., Cumming, A., et al. 2012, MNRAS, 426, 927, doi: 10.1111/j.1365-2966.2012.21769.x
- Andrae, R., Schulze-Hartung, T., & Melchior, P. 2010, arXiv e-prints, arXiv:1012.3754, doi: 10.48550/arXiv.1012.3754
- Antoniadis, J., Freire, P. C. C., Wex, N., et al. 2013, Science, 340, 448, doi: 10.1126/science.1233232
- Baade, W., & Zwicky, F. 1934, Proceedings of the National Academy of Science, 20, 259, doi: 10.1073/pnas.20.5.259
- Bahramian, A., Heinke, C. O., Sivakoff, G. R., et al. 2014, ApJ, 780, 127, doi: 10.1088/0004-637X/780/2/127
- Baiko, D. A., Kaminker, A. D., Potekhin, A. Y., & Yakovlev, D. G. 1998, Phys. Rev. Lett., 81, 5556
- Baiko, D. A., Potekhin, A. Y., & Yakovlev, D. G. 2001, Phys. Rev. E, 64, 057402
- Barret, D., Motch, C., & Predehl, P. 1998, A&A, 329, 965
- Barret, D., Bouchet, L., Mandrou, P., et al. 1992, ApJ, 394, 615, doi: 10.1086/171615
- Barthelmy, S. D., Gehrels, N., Page, K. L., & Ukwatta, T. N. 2015a, GRB Coordinates Network, 17367, 1
- Barthelmy, S. D., Burrows, D. N., Gronwall, C., et al. 2015b, GRB Coordinates Network, 17416, 1
- Baym, G., Bethe, H. A., & Pethick, C. J. 1971, Nucl. Phys. A, 175, 225
- Baym, G., Pethick, C., & Sutherland, P. 1971, ApJ, 170, 299
- Bisnovatyi-Kogan, G. S., & Chechetkin, V. M. 1979, Soviet Phys.-Uspekhi, 127, 263
- Bordas, P., Kuulkers, E., Alfonso-Garzón, J., et al. 2010, The Astronomer's Telegram, 2919, 1
- Brown, E. F. 2015, dStar: Neutron star thermal evolution code. <http://ascl.net/1505.034>

Brown, E. F., & Cumming, A. 2009, ApJ, 698, 1020

Brown, E. F., Cumming, A., Fattoyev, F. J., et al. 2018, Phys. Rev. Lett., 120, 182701, doi: 10.1103/PhysRevLett.120.182701

Burrows, A., & Lattimer, J. M. 1986, ApJ, 307, 178

Cackett, E. M., Brown, E. F., Cumming, A., et al. 2013, ApJ, 774, 131, doi: 10.1088/0004-637X/774/2/131

—. 2010, ApJ, 722, L137

Cackett, E. M., Fridriksson, J. K., Homan, J., Miller, J. M., & Wijnands, R. 2011, MNRAS, 414, 3006, doi: 10.1111/j.1365-2966.2011.18601.x

Cackett, E. M., Wijnands, R., Linares, M., et al. 2006, MNRAS, 372, 479

Cackett, E. M., Wijnands, R., Miller, J. M., Brown, E. F., & Degenaar, N. 2008, ApJ, 687, L87

Cameron, A. G. 1959, ApJ, 130, 884, doi: 10.1086/146780

Campana, S., Brivio, F., Degenaar, N., et al. 2014, MNRAS, 441, 1984, doi: 10.1093/mnras/stu709

Campana, S., Coti Zelati, F., & D'Avanzo, P. 2013, MNRAS, 432, 1695, doi: 10.1093/mnras/stt604

Casella, P., Altamirano, D., Patruno, A., Wijnands, R., & van der Klis, M. 2008, ApJ, 674, L41, doi: 10.1086/528982

Chabrier, G., & Potekhin, A. Y. 1998, Phys. Rev. E, 58, 4941

Chadwick, J. 1932, Nature, 129, 312, doi: 10.1038/129312a0

Chamel, N., & Haensel, P. 2008, Living Reviews in Relativity, 11, 10, doi: 10.12942/lrr-2008-10

Chenevez, J., Kuulkers, E., Brandt, S., et al. 2012, The Astronomer's Telegram, 4050, 1

Cheng, K. S., Yao, C. C., & Dai, Z. G. 1997, Phys. Rev. C, 55, 2092, doi: 10.1103/PhysRevC.55.2092

Chugunov, A. I., & Shchepochin, N. N. 2020, MNRAS, 495, L32, doi: 10.1093/mnrasl/slaa055

Cominsky, L. R., & Wood, K. S. 1984, ApJ, 283, 765, doi: 10.1086/162361

- . 1989, *ApJ*, 337, 485, doi: [10.1086/167117](https://doi.org/10.1086/167117)
- Coriat, M., Tzioumis, T., Corbel, S., et al. 2011, *The Astronomer's Telegram*, 3119, 1
- Coti Zelati, F., Campana, S., D'Avanzo, P., & Melandri, A. 2014, *MNRAS*, 438, 2634, doi: [10.1093/mnras/stt2384](https://doi.org/10.1093/mnras/stt2384)
- Cumming, A., Brown, E. F., Fattoyev, F. J., et al. 2017, *Phys. Rev. C*, 95, 025806, doi: [10.1103/PhysRevC.95.025806](https://doi.org/10.1103/PhysRevC.95.025806)
- Degenaar, N., Brown, E. F., & Wijnands, R. 2011, *MNRAS*, 418, L152, doi: [10.1111/j.1745-3933.2011.01164.x](https://doi.org/10.1111/j.1745-3933.2011.01164.x)
- Degenaar, N., & Wijnands, R. 2011, *Monthly Notices of the Royal Astronomical Society: Letters*, 412, L68, doi: [10.1111/j.1745-3933.2011.01007.x](https://doi.org/10.1111/j.1745-3933.2011.01007.x)
- Degenaar, N., Wijnands, R., Wolff, M. T., et al. 2009, *MNRAS*, 396, L26
- Degenaar, N., Wijnands, R., Brown, E. F., et al. 2013, *ApJ*, 775, 48, doi: [10.1088/0004-637X/775/1/48](https://doi.org/10.1088/0004-637X/775/1/48)
- Degenaar, N., Medin, Z., Cumming, A., et al. 2014, *ApJ*, 791, 47, doi: [10.1088/0004-637X/791/1/47](https://doi.org/10.1088/0004-637X/791/1/47)
- Degenaar, N., Wijnands, R., Bahramian, A., et al. 2015, *MNRAS*, 451, 2071, doi: [10.1093/mnras/stv1054](https://doi.org/10.1093/mnras/stv1054)
- Degenaar, N., Ootes, L. S., Page, D., et al. 2019, *MNRAS*, 488, 4477, doi: [10.1093/mnras/stz1963](https://doi.org/10.1093/mnras/stz1963)
- Deibel, A., Cumming, A., Brown, E. F., & Page, D. 2015, *ApJ*, 809, L31, doi: [10.1088/2041-8205/809/2/L31](https://doi.org/10.1088/2041-8205/809/2/L31)
- Demorest, P. B., Pennucci, T., Ransom, S. M., Roberts, M. S. E., & Hessels, J. W. T. 2010, *Nature*, 467, 1081
- Douchin, F., & Haensel, P. 2000, *Physics Letters B*, 485, 107, doi: [10.1016/S0370-2693\(00\)00672-9](https://doi.org/10.1016/S0370-2693(00)00672-9)
- Douchin, F., Haensel, P., & Meyer, J. 2000, *Nuclear Physics A*, 665, 419, doi: [https://doi.org/10.1016/S0375-9474\(99\)00397-8](https://doi.org/10.1016/S0375-9474(99)00397-8)
- Dubus, G., Hameury, J.-M., & Lasota, J.-P. 2001, *A&A*, 373, 251
- Dubus, G., Lasota, J.-P., Hameury, J.-M., & Charles, P. 1999, *MNRAS*, 303, 139, doi: [10.1046/](https://doi.org/10.1046/)

j.1365-8711.1999.02212.x

Foreman-Mackey, D., Hogg, D. W., Lang, D., & Goodman, J. 2013, *PASP*, 125, 306, doi: 10.1086/670067

Fridriksson, J. K., Homan, J., Wijnands, R., et al. 2010, *ApJ*, 714, 270

Fridriksson, J. K., Homan, J., Wijnands, R., et al. 2011, *The Astrophysical Journal*, 736, 162, doi: 10.1088/0004-637X/736/2/162

Friedman, H., Byram, E. T., & Chubb, T. A. 1967, *Science*, 156, 374, doi: 10.1126/science.156.3773.374

Galloway, D. K., Muno, M. P., Hartman, J. M., Psaltis, D., & Chakrabarty, D. 2008, *ApJS*, 179, 360

Giacconi, R., Gursky, H., Paolini, F. R., & Rossi, B. B. 1962, *Phys. Rev. Lett.*, 9, 439

Gold, T. 1968, *Nature*, 218, 731, doi: 10.1038/218731a0

Goodman, J., & Weare, J. 2010, *Communications in Applied Mathematics and Computational Science*, 5, 65, doi: 10.2140/camcos.2010.5.65

Güngör, C., Güver, T., & Ekşi, K. Y. 2014, *MNRAS*, 439, 2717, doi: 10.1093/mnras/stu128

Gusakov, M. E., & Chugunov, A. I. 2020, *Phys. Rev. Lett.*, 124, 191101, doi: 10.1103/PhysRevLett.124.191101

—. 2021, *Phys. Rev. D*, 103, L101301, doi: 10.1103/PhysRevD.103.L101301

Gusakov, M. E., Kantor, E. M., & Chugunov, A. I. 2021, *Phys. Rev. D*, 104, L081301, doi: 10.1103/PhysRevD.104.L081301

Haensel, P., Potekhin, A. Y., & Yakovlev, D. G. 2007, *Neutron Stars 1: Equation of State and Structure*, Vol. 326 (New York: Springer)

Haensel, P., & Zdunik, J. L. 1990, *A&A*, 227, 431

—. 2003, *A&A*, 404, L33, doi: 10.1051/0004-6361:20030708

—. 2008, *A&A*, 480, 459, doi: 10.1051/0004-6361:20078578

Hansen, C. J., & Van Horn, H. M. 1975, *ApJ*, 195, 735

Hashimoto, M., Seki, H., & Yamada, M. 1984, *Progress of Theoretical Physics*, 71, 320, doi: 10.



1143/PTP.71.320

- Hastings, W. K. 1970, *Biometrika*, 57, 97, doi: 10.1093/biomet/57.1.97
- Heinke, C. O., Rybicki, G. B., Narayan, R., & Grindlay, J. E. 2006, *ApJ*, 644, 1090
- Heney, L., & L'Ecuyer, J. L. 1969, *ApJ*, 156, 549
- Hewish, A., Bell, S. J., Pilkington, J. D. H., Scott, P. F., & Collins, R. A. 1968, *Nature*, 217, 709
- Higson, E., Handley, W., Hobson, M., & Lasenby, A. 2019, *Statistics and Computing*, 29, 891, doi: 10.1007/s11222-018-9844-0
- Hogg, D. W., Bovy, J., & Lang, D. 2010, arXiv e-prints, arXiv:1008.4686, doi: 10.48550/arXiv.1008.4686
- Homan, J., Fridriksson, J. K., Wijnands, R., et al. 2014, *ApJ*, 795, 131, doi: 10.1088/0004-637X/795/2/131
- Homan, J., Linares, M., van den Berg, M., & Fridriksson, J. 2011, *The Astronomer's Telegram*, 3650, 1
- Homan, J., van der Klis, M., Wijnands, R., et al. 2007, *ApJ*, 656, 420
- Horowitz, C. J., Berry, D. K., Briggs, C. M., et al. 2015, *Physical Review Letters*, 114, 031102, doi: 10.1103/PhysRevLett.114.031102
- Horowitz, C. J., Dussan, H., & Berry, D. K. 2008, *Phys. Rev. C*, 77, 045807
- Hughto, J., Schneider, A. S., Horowitz, C. J., & Berry, D. K. 2011, *Phys. Rev. E*, 84, 016401, doi: 10.1103/PhysRevE.84.016401
- Hulse, R. A., & Taylor, J. H. 1975, *ApJ*, 195, L51, doi: 10.1086/181708
- Inogamov, N. A., & Sunyaev, R. A. 1999, *Astronomy Letters*, 25, 269
- . 2010, *Astronomy Letters*, 36, 848, doi: 10.1134/S1063773710120029
- Itoh, N., & Kohyama, Y. 1993, *ApJ*, 404, 268
- Jaynes, E. T., & Kempthorne, O. 1976, *Confidence Intervals vs Bayesian Intervals* (Dordrecht: Springer Netherlands), 175–257, doi: 10.1007/978-94-010-1436-6\_6
- Kaluzienski, L. J., Holt, S. S., Boldt, E. A., & Serlemitsos, P. J. 1977, *Nature*, 265, 606, doi: 10.1038/265606a0

- Kass, R. E., & Raftery, A. E. 1995, *Journal of the American Statistical Association*, 90, 773, doi: 10.1080/01621459.1995.10476572
- Kaur, R., & Heinke, C. 2012, *The Astronomer's Telegram*, 4085, 1
- King, A. R. 2004, in *Compact Stellar X-ray Sources*, ed. W. H. G. Lewin & M. Van Der Klis (Cambridge: Cambridge University Press), 507
- King, A. R., & Ritter, H. 1998, *MNRAS*, 293, L42, doi: 10.1046/j.1365-8711.1998.01295.x
- Kitamoto, S., Tsunemi, H., Miyamoto, S., & Roussel-Dupre, D. 1993, *ApJ*, 403, 315, doi: 10.1086/172204
- Koyama, K., Inoue, H., Makishima, K., et al. 1981, *ApJ*, 247, L27, doi: 10.1086/183582
- Krimm, H. A., Barthelmy, S. D., Baumgartner, W., et al. 2015, *The Astronomer's Telegram*, 6997, 1
- Kuulkers, E., in 't Zand, J. J. M., van Kerkwijk, M. H., et al. 2002, *A&A*, 382, 503
- Lau, R., Beard, M., Gupta, S. S., et al. 2018, *ApJ*, 859, 62, doi: 10.3847/1538-4357/aabfe0
- Lewin, W. H. G., Hoffman, J. A., Doty, J., & Liller, W. 1976, *IAU Circ.*, 2994, 2
- Lin, D., Altamirano, D., Homan, J., et al. 2009, *ApJ*, 699, 60, doi: 10.1088/0004-637X/699/1/60
- Mackay, D. J. C. 2003, *Information Theory, Inference and Learning Algorithms* (Cambridge University Press)
- Mackie, F. D., & Baym, G. 1977, *Nucl. Phys. A*, 285, 332
- Maruyama, T., Tatsumi, T., Voskresensky, D. N., Tanigawa, T., & Chiba, S. 2005, *Phys. Rev. C*, 72, 015802, doi: 10.1103/PhysRevC.72.015802
- Matsumura, T., Negoro, H., Suwa, F., et al. 2011, *The Astronomer's Telegram*, 3102, 1, doi: 10.48550/arXiv.1105.4136
- Medin, Z., & Cumming, A. 2011, *ApJ*, 730, 97, doi: 10.1088/0004-637X/730/2/97
- Medin, Z., & Cumming, A. 2014, *The Astrophysical Journal Letters*, 783, L3, doi: 10.1088/2041-8205/783/1/L3
- . 2015, *The Astrophysical Journal*, 802, 29, doi: 10.1088/0004-637X/802/1/29

- Merritt, R. L., Cackett, E. M., Brown, E. F., et al. 2016, *ApJ*, 833, 186, doi: [10.3847/1538-4357/833/2/186](https://doi.org/10.3847/1538-4357/833/2/186)
- Metropolis, N., Rosenbluth, A. W., Rosenbluth, M. N., Teller, A. H., & Teller, E. 1953, *The Journal of Chemical Physics*, 21, 1087, doi: [10.1063/1.1699114](https://doi.org/10.1063/1.1699114)
- Möller, P., Sierk, A. J., Ichikawa, T., & Sagawa, H. 2016, *Atomic Data and Nuclear Data Tables*, 109, 1, doi: [10.1016/j.adt.2015.10.002](https://doi.org/10.1016/j.adt.2015.10.002)
- Muno, M. P., Fox, D. W., Morgan, E. H., & Bildsten, L. 2000, *ApJ*, 542, 1016
- Narita, T., Grindlay, J. E., & Barret, D. 2001, *ApJ*, 547, 420, doi: [10.1086/318326](https://doi.org/10.1086/318326)
- Negele, J. W., & Vautherin, D. 1973, *Nucl. Phys. A*, 207, 298
- Nelder, J. A., & Mead, R. 1965, *The Computer Journal*, 7, 308, doi: [10.1093/comjnl/7.4.308](https://doi.org/10.1093/comjnl/7.4.308)
- Oosterbroek, T., Parmar, A. N., Sidoli, L., in 't Zand, J. J. M., & Heise, J. 2001, *A&A*, 376, 532
- Ootes, L. S., Page, D., Wijnands, R., & Degenaar, N. 2016, *MNRAS*, 461, 4400, doi: [10.1093/mnras/stw1799](https://doi.org/10.1093/mnras/stw1799)
- Ootes, L. S., Wijnands, R., Page, D., & Degenaar, N. 2018, *MNRAS*, 477, 2900, doi: [10.1093/mnras/sty825](https://doi.org/10.1093/mnras/sty825)
- Ootes, L. S., Vats, S., Page, D., et al. 2019, *MNRAS*, 487, 1447, doi: [10.1093/mnras/stz1406](https://doi.org/10.1093/mnras/stz1406)
- Oppenheimer, J. R., & Volkoff, G. M. 1939, *Phys. Rev.*, 55, 374
- Ortolani, S., Barbuy, B., Bica, E., Zoccali, M., & Renzini, A. 2007, *A&A*, 470, 1043, doi: [10.1051/0004-6361:20066628](https://doi.org/10.1051/0004-6361:20066628)
- Oyamatsu, K., Hashimoto, M. a., & Yamada, M. 1984, *Prog. Theor. Phys.*, 72, 373, doi: [10.1143/ptp.72.373](https://doi.org/10.1143/ptp.72.373)
- Oyamatsu, K., & Iida, K. 2007, *Phys. Rev. C*, 75, 015801, doi: [10.1103/PhysRevC.75.015801](https://doi.org/10.1103/PhysRevC.75.015801)
- Özel, F., & Freire, P. 2016, *ARA&A*, 54, 401, doi: [10.1146/annurev-astro-081915-023322](https://doi.org/10.1146/annurev-astro-081915-023322)
- Özel, F., Gould, A., & Güver, T. 2012, *ApJ*, 748, 5, doi: [10.1088/0004-637X/748/1/5](https://doi.org/10.1088/0004-637X/748/1/5)
- Pacini, F. 1967, *Nature*, 216, 567, doi: [10.1038/216567a0](https://doi.org/10.1038/216567a0)
- Page, D., Homan, J., Nava-Callejas, M., et al. 2022, arXiv e-prints, arXiv:2202.03962. <https://arxiv.org/abs/2202.03962>

- Page, D., & Reddy, S. 2013, Phys. Rev. Lett., 111, 241102, doi: [10.1103/PhysRevLett.111.241102](https://doi.org/10.1103/PhysRevLett.111.241102)
- Parikh, A. S., Wijnands, R., Homan, J., et al. 2020, A&A, 638, L2, doi: [10.1051/0004-6361/202038198](https://doi.org/10.1051/0004-6361/202038198)
- Parikh, A. S., Homan, J., Wijnands, R., et al. 2017a, ApJ, 851, L28, doi: [10.3847/2041-8213/aa9e03](https://doi.org/10.3847/2041-8213/aa9e03)
- Parikh, A. S., Wijnands, R., Degenaar, N., et al. 2017b, MNRAS, 466, 4074, doi: [10.1093/mnras/stw3388](https://doi.org/10.1093/mnras/stw3388)
- Parikh, A. S., Wijnands, R., Ootes, L. S., et al. 2019, A&A, 624, A84, doi: [10.1051/0004-6361/201834412](https://doi.org/10.1051/0004-6361/201834412)
- Parmar, A. N., White, N. E., Giommi, P., & Gottwald, M. 1986, ApJ, 308, 199
- Pethick, C. J., & Potekhin, A. Y. 1998, Physics Letters B, 427, 7, doi: [10.1016/S0370-2693\(98\)00341-4](https://doi.org/10.1016/S0370-2693(98)00341-4)
- Pethick, C. J., Ravenhall, D. G., & Lorenz, C. P. 1995, Nucl. Phys. A, 584, 675
- Pooley, D., Homan, J., Heinke, C., et al. 2010, The Astronomer's Telegram, 2974, 1
- Potekhin, A. Y., & Chabrier, G. 2000, Phys. Rev. E, 62, 8554
- Potekhin, A. Y., & Chabrier, G. 2010, Contributions to Plasma Physics, 50, 82
- Potekhin, A. Y., Chabrier, G., & Yakovlev, D. G. 1997, A&A, 323, 415
- Potekhin, A. Y., Pons, J. A., & Page, D. 2015, Space Sci. Rev., 191, 239, doi: [10.1007/s11214-015-0180-9](https://doi.org/10.1007/s11214-015-0180-9)
- Ravenhall, D. G., Pethick, C. J., & Wilson, J. R. 1983, Phys. Rev. Lett., 50, 2066, doi: [10.1103/PhysRevLett.50.2066](https://doi.org/10.1103/PhysRevLett.50.2066)
- Remillard, R. A., Lin, D., ASM Team at MIT, & NASA/GSFC. 2006, The Astronomer's Telegram, 696, 1
- Revnivtsev, M., Churazov, E., Gilfanov, M., & Sunyaev, R. 2001, A&A, 372, 138, doi: [10.1051/0004-6361:20010434](https://doi.org/10.1051/0004-6361:20010434)
- Ritter, H., Zhang, Z. Y., & Kolb, U. 2000, A&A, 360, 959, doi: [10.48550/arXiv.astro-ph/0005480](https://doi.org/10.48550/arXiv.astro-ph/0005480)

- Russell, D. M., Lewis, F., Doran, R., & Roberts, S. 2011, *The Astronomer's Telegram*, 3116, 1
- Rutledge, R. E., Bildsten, L., Brown, E. F., Pavlov, G. G., & Zavlin, V. E. 2001b, *ApJ*, 559, 1054
- Rutledge, R. E., Bildsten, L., Brown, E. F., et al. 2002, *ApJ*, 580, 413
- Sandage, A., Osmer, P., Giacconi, R., et al. 1966, *ApJ*, 146, 316, doi: 10.1086/148892
- Sato, K. 1979, *Prog. Theor. Physics*, 62, 957
- Schatz, H., Aprahamian, A., Barnard, V., et al. 2001, *Phys. Rev. Lett.*, 86, 3471, doi: 10.1103/PhysRevLett.86.3471
- Schiesser, W. E. 1991, *The Numerical Method of Lines: Integration of Partial Differential Equations* (Academic Press)
- Schmitt, A. 2015, *Lecture Notes in Physics*, Vol. 888, *Introduction to Superfluidity* (Springer), doi: 10.1007/978-3-319-07947-9
- Schwarz, G. 1978, *The Annals of Statistics*, 6, 461, doi: 10.1214/aos/1176344136
- Sharma, S. 2017, *Annual Review of Astronomy and Astrophysics*, 55, 213, doi: 10.1146/annurev-astro-082214-122339
- Shchechilin, N. N., & Chugunov, A. I. 2019, *MNRAS*, 490, 3454, doi: 10.1093/mnras/stz2838
- Shchechilin, N. N., Gusakov, M. E., & Chugunov, A. I. 2021, *MNRAS*, 507, 3860, doi: 10.1093/mnras/stab2415
- Shklovsky, I. S. 1967, *ApJ*, 148, L1+
- Shternin, P. S., & Yakovlev, D. G. 2006, *Phys. Rev. D*, 74, 043004
- Shternin, P. S., Yakovlev, D. G., Haensel, P., & Potekhin, A. Y. 2007, *MNRAS*, 382, L43
- Skilling, J. 2004, in *American Institute of Physics Conference Series*, Vol. 735, *Bayesian Inference and Maximum Entropy Methods in Science and Engineering: 24th International Workshop on Bayesian Inference and Maximum Entropy Methods in Science and Engineering*, ed. R. Fischer, R. Preuss, & U. V. Toussaint (AIP), 395–405, doi: 10.1063/1.1835238
- Skyrme, T. 1958, *Nuclear Physics*, 9, 615, doi: [https://doi.org/10.1016/0029-5582\(58\)90345-6](https://doi.org/10.1016/0029-5582(58)90345-6)
- Smith, D. A., Morgan, E. H., & Bradt, H. 1997, *ApJ*, 479, L137, doi: 10.1086/310604

- Steiner, A. W. 2012, *Phys. Rev. C*, 85, 055804, doi: [10.1103/PhysRevC.85.055804](https://doi.org/10.1103/PhysRevC.85.055804)
- Strohmayer, T. E., & Bildsten, L. 2004, in *Compact Stellar X-ray Sources*, ed. W. Lewin & M. van der Klis (Cambridge: Cambridge University Press), 113
- Strohmayer, T. E., Markwardt, C. B., Pereira, D., & Smith, E. A. 2010, *The Astronomer's Telegram*, 2946, 1
- Sunyaev, R. 1989, *IAU Circ.*, 4839, 1
- . 1990, *IAU Circ.*, 5104, 1
- Taylor, J. H., & Weisberg, J. M. 1982, *ApJ*, 253, 908, doi: [10.1086/159690](https://doi.org/10.1086/159690)
- . 1989, *ApJ*, 345, 434, doi: [10.1086/167917](https://doi.org/10.1086/167917)
- Testa, V., di Salvo, T., D'Antona, F., et al. 2012, *A&A*, 547, A28, doi: [10.1051/0004-6361/201219904](https://doi.org/10.1051/0004-6361/201219904)
- Timmes, F. X., & Swesty, F. D. 2000, *ApJS*, 126, 501
- Tolman, R. C. 1939, *Physical Review*, 55, 364
- Ushomirsky, G., & Rutledge, R. E. 2001, *MNRAS*, 325, 1157
- Virtanen, P., Gommers, R., Oliphant, T. E., et al. 2020, *Nature Methods*, 17, 261, doi: [10.1038/s41592-019-0686-2](https://doi.org/10.1038/s41592-019-0686-2)
- Voges, W., Aschenbach, B., Boller, T., et al. 1999, *A&A*, 349, 389, doi: [10.48550/arXiv.astro-ph/9909315](https://doi.org/10.48550/arXiv.astro-ph/9909315)
- Wall, J. V., & Jenkins, C. R. 2012, *Practical Statistics for Astronomers* (Cambridge University Press)
- Wang, M., Huang, W., Kondev, F., Audi, G., & Naimi, S. 2021, *Chinese Physics C*, 45, 030003, doi: [10.1088/1674-1137/abddaf](https://doi.org/10.1088/1674-1137/abddaf)
- Watanabe, G., & Iida, K. 2003, *Phys. Rev. C*, 68, 045801, doi: [10.1103/PhysRevC.68.045801](https://doi.org/10.1103/PhysRevC.68.045801)
- Waterhouse, A. C., Degenaar, N., Wijnands, R., et al. 2016, *Monthly Notices of the Royal Astronomical Society*, 456, 4001, doi: [10.1093/mnras/stv2959](https://doi.org/10.1093/mnras/stv2959)
- Wijnands, R., Degenaar, N., & Page, D. 2013, *MNRAS*, 432, 2366, doi: [10.1093/mnras/stt599](https://doi.org/10.1093/mnras/stt599)
- . 2017, *Journal of Astrophysics and Astronomy*, 38, 49, doi: [10.1007/s12036-017-9466-5](https://doi.org/10.1007/s12036-017-9466-5)

- Wijnands, R., Miller, J. M., Markwardt, C., Lewin, W. H. G., & van der Klis, M. 2001, *ApJ*, 560, L159, doi: 10.1086/324378
- Wijnands, R., Nowak, M., Miller, J. M., et al. 2003, *ApJ*, 594, 952, doi: 10.1086/377122
- Wijnands, R. A. D., & van der Klis, M. 1997, *ApJ*, 482, L65, doi: 10.1086/310682
- Wilks, S. S. 1938, *The Annals of Mathematical Statistics*, 9, 60, doi: 10.1214/aoms/1177732360
- Wolff, M. T., Ray, P. S., & Wood, K. S. 2008, *The Astronomer's Telegram*, 1736, 1
- Yakovlev, D. G., Kaminker, A. D., Gnedin, O. Y., & Haensel, P. 2001, *Phys. Rep.*, 354, 1
- Yakovlev, D. G., & Pethick, C. J. 2004, *ARA&A*, 42, 169
- Yamauchi, S., & Koyama, K. 1990, *PASJ*, 42, L83
- Zavlin, V. E., Pavlov, G. G., & Shibano, Y. A. 1996, *A&A*, 315, 141

## APPENDIX

### DETAILED RESULTS FOR MODELS WITH AND WITHOUT NEUTRON DIFFUSION

Included here are the posterior distributions of the non-diffusive and nHD models discussed in Chapter 4. For each source, there are eight models, distinguished by the heating profile and impurity of the crust and by the exclusion or inclusion of a pasta layer.

The heating profile and impurity of the crust are set by the crust composition, which may be one of the non-diffusive, reaction network calculations of Lau et al. (2018) or the nHD crust of Gusakov & Chugunov (2020). The non-diffusive crusts differ in their initial composition:  $^{56}\text{Fe}$ , superburst ashes, or rp-process ashes. Both  $Q_{\text{sh}}$  and  $T_{\text{c}}$  are left as free parameters for all models.

Pasta may be included in a model by adding a layer at the bottom of the crust where the impurity is equal to the pasta impurity,  $Q_{\text{imp,pasta}}$ . Although uncertain, the value of  $Q_{\text{imp,pasta}}$  is expected to be high (see, e.g., Horowitz et al., 2015, who estimate  $Q_{\text{imp,pasta}} \approx 40$ ). I leave  $Q_{\text{imp,pasta}}$  as a free parameter in the pasta models under the constraint that  $Q_{\text{imp,pasta}} \geq 20$ .

For models without pasta,  $T_{\text{c}}$  and  $Q_{\text{sh}}$  are well-constrained with normal marginal distributions. For the models with pasta,  $Q_{\text{sh}}$  is well-constrained and normally distributed for all sources and models.  $Q_{\text{imp,pasta}}$  is not well-constrained by most of the models. For KS 1731–260, all crust compositions favor  $Q_{\text{imp,pasta}} \lesssim 25$ . For MXB 1659–29 with initial rp-process ashes composition  $Q_{\text{imp,pasta}} \lesssim 30$  is favored, while for all other compositions,  $Q_{\text{imp,pasta}} \approx 40\text{--}60$  appears to be favored.

One notable feature of the pasta models for KS 1731–260 and MXB 1659–29 is that they tend to favor very low  $T_{\text{c}}$ . The marginalized distributions of  $T_{\text{c}}$  all appear to be truncated at the minimum value allowed by the prior ( $\log [T_{\text{c}}/(\text{K})] = 6.5$ ). A core temperature this low seems very unlikely, because it is inconsistent with the sources being in a long-term steady-state. For example, Brown et al. (2018) estimate that, given its observed accretion history, MXB 1659–29 would reach a steady-state core temperature in a few centuries, which is much shorter than the accretion lifetime of a LMXB; further, MXB 1659–29 would need a neutrino luminosity of  $3 \times 10^{34} \text{ erg s}^{-1}$  to maintain a steady-state core temperature. For a direct Urca process, operating over the entire



neutron star core (assuming a radius of 11 km), to supply this luminosity requires  $T_c \gtrsim 10^7$  K.

While the posterior distributions can inform us of the parameter values that are favored by the data as well as how well-constrained the parameters are, they do not inform us of the overall goodness-of-fit of each model. That is quantified with the evidence  $\mathcal{Z}$ , which is reported in Table 4.3 and discussed in Section 4.4.

## KS 1731-260 (No Pasta)

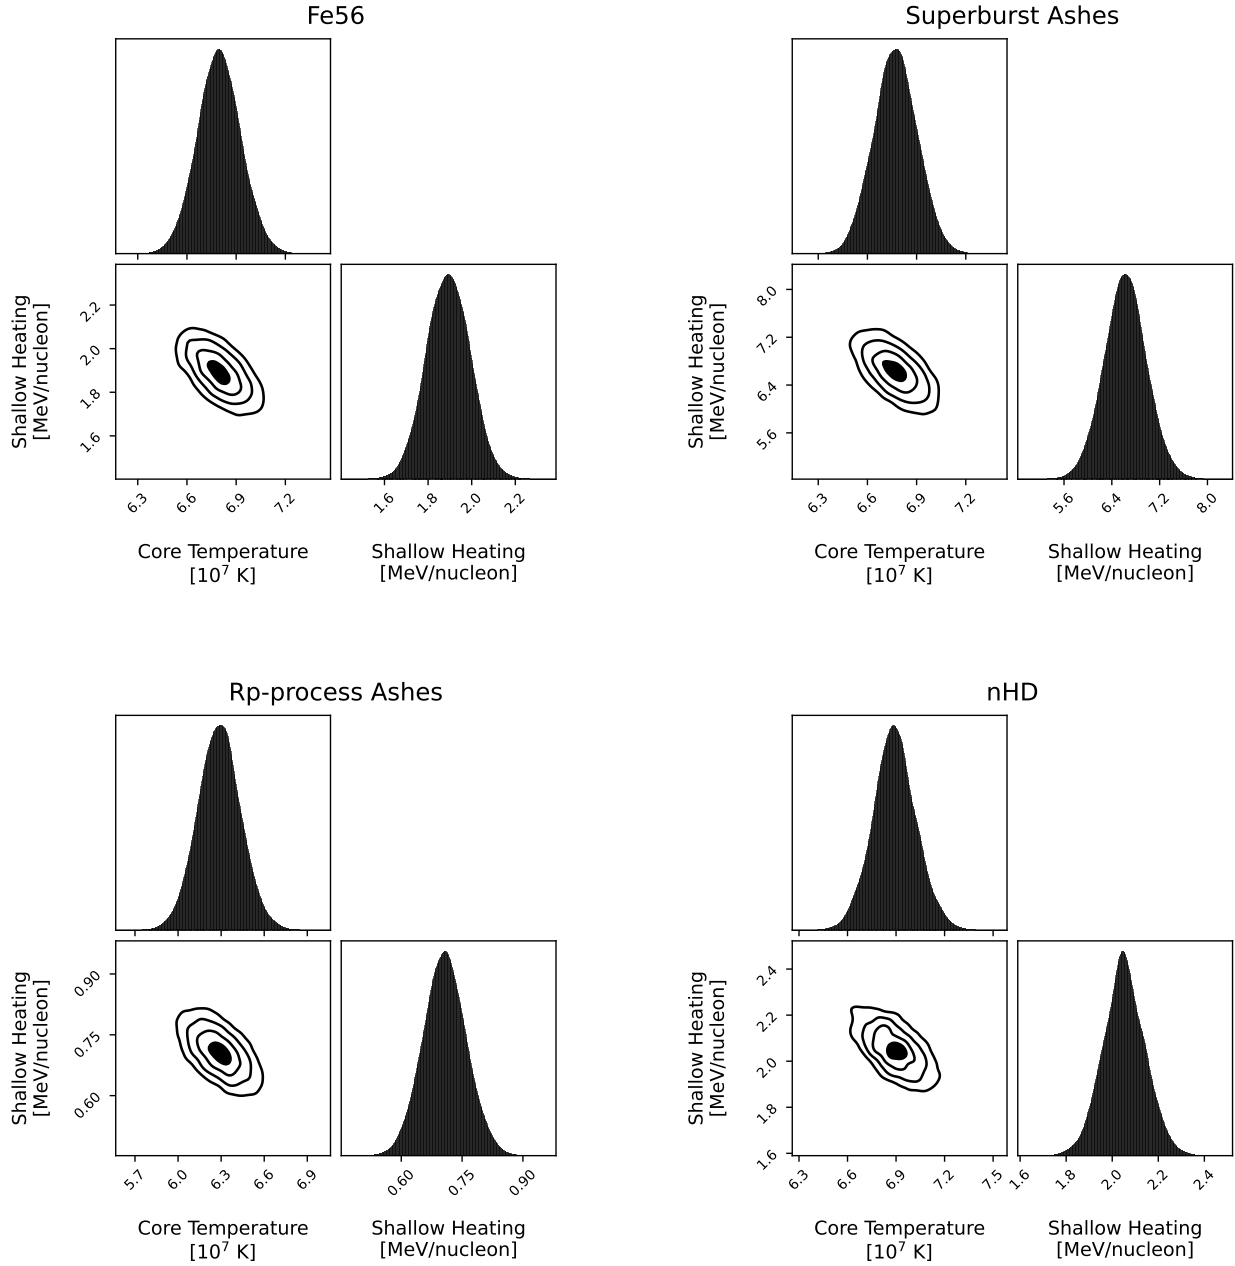


Figure A.1 Posterior distributions of  $T_c$  and  $Q_{sh}$  for KS 1731–260. The non-diffusive and nHD models without a pasta layer are shown.

KS 1731-260 (Pasta)

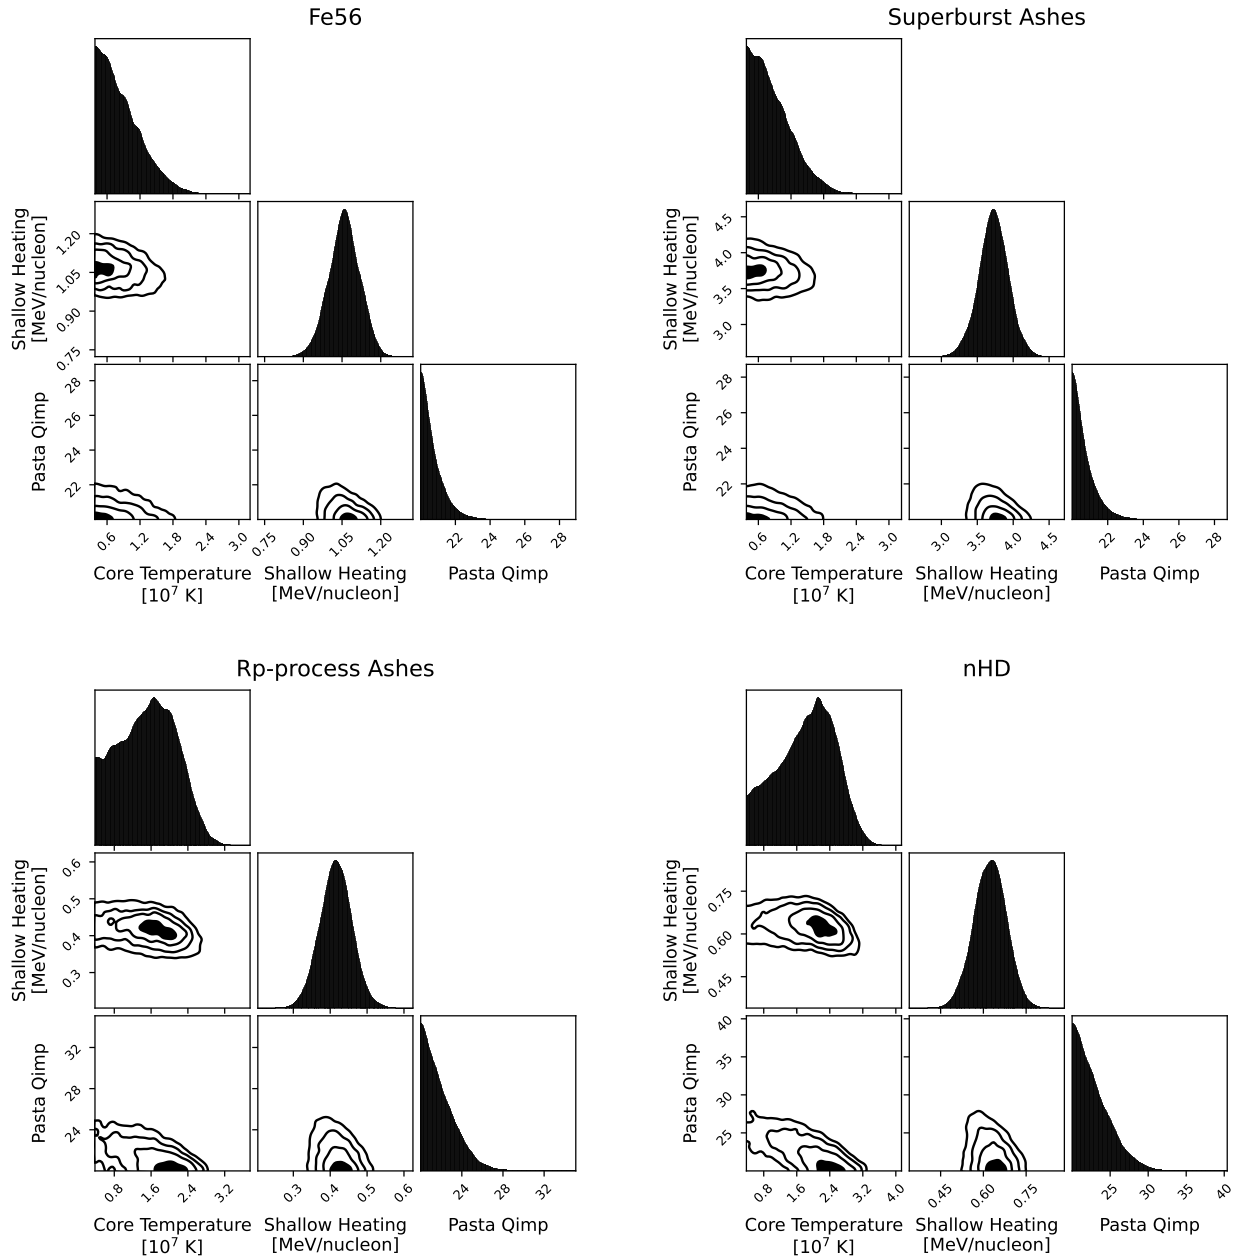


Figure A.2 Posterior distributions of  $T_c$ ,  $Q_{sh}$ , and  $Q_{imp,pasta}$  for MXB 1659–29. The plots display the non-diffusive and nHD models with a pasta layer.

### MXB 1659-29 (No Pasta)

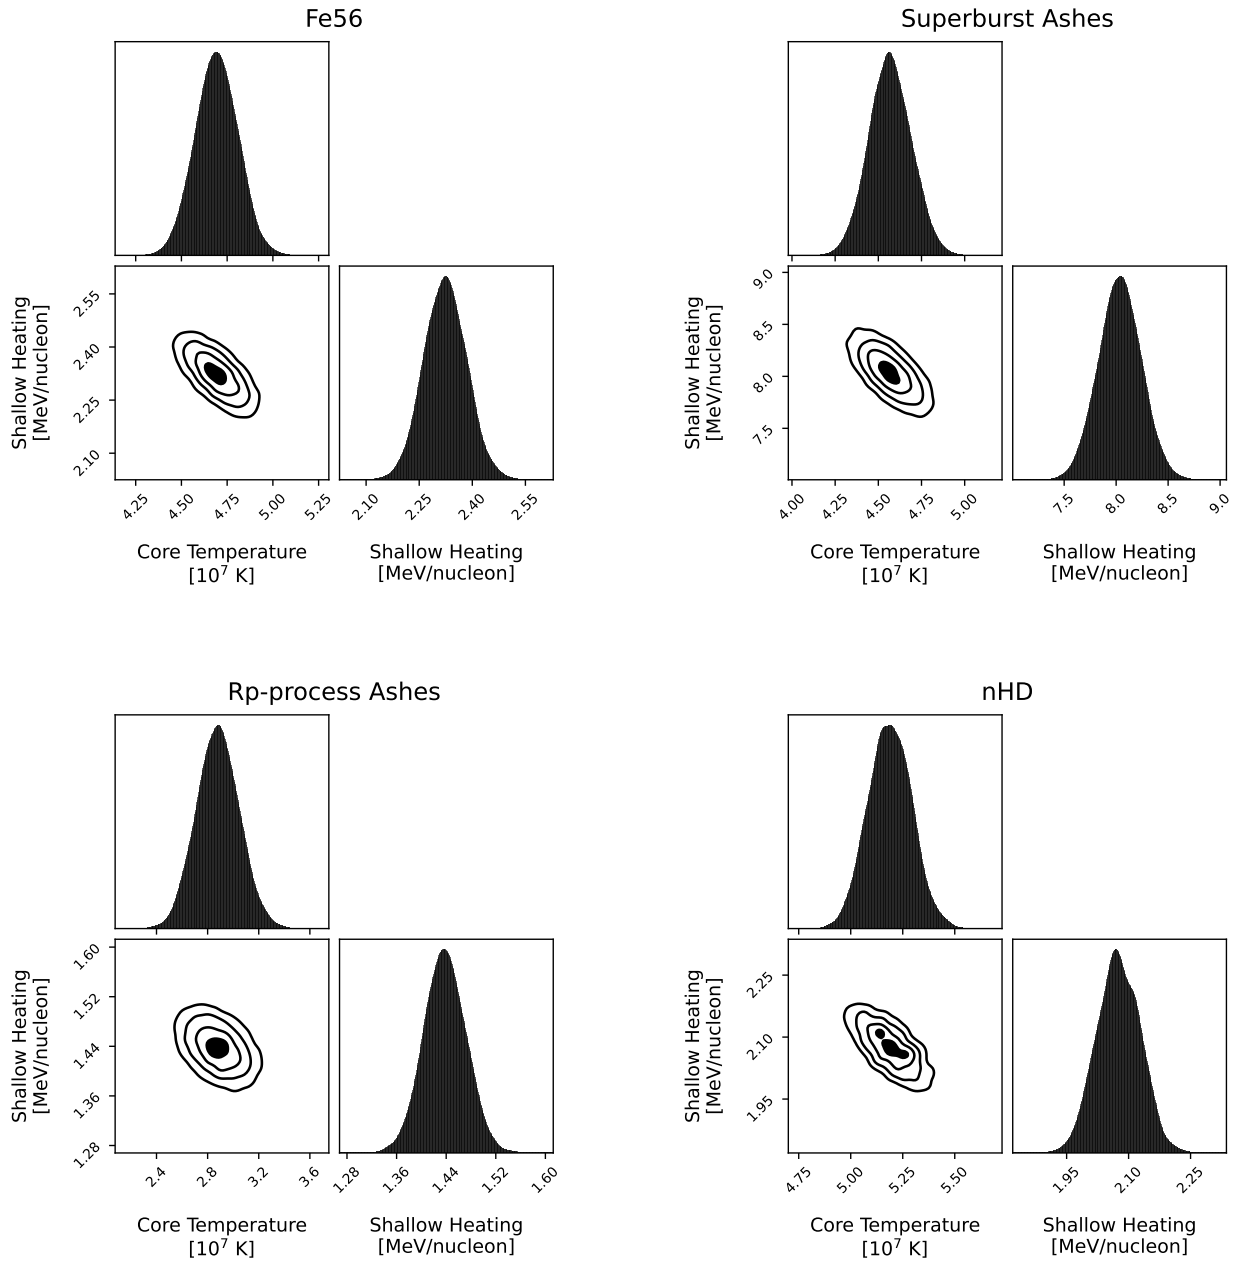


Figure A.3 Posterior distributions of  $T_c$  and  $Q_{sh}$  for MXB 1659-29. The non-diffusive and nHD models without a pasta layer are shown.

### MXB 1659-29 (Pasta)

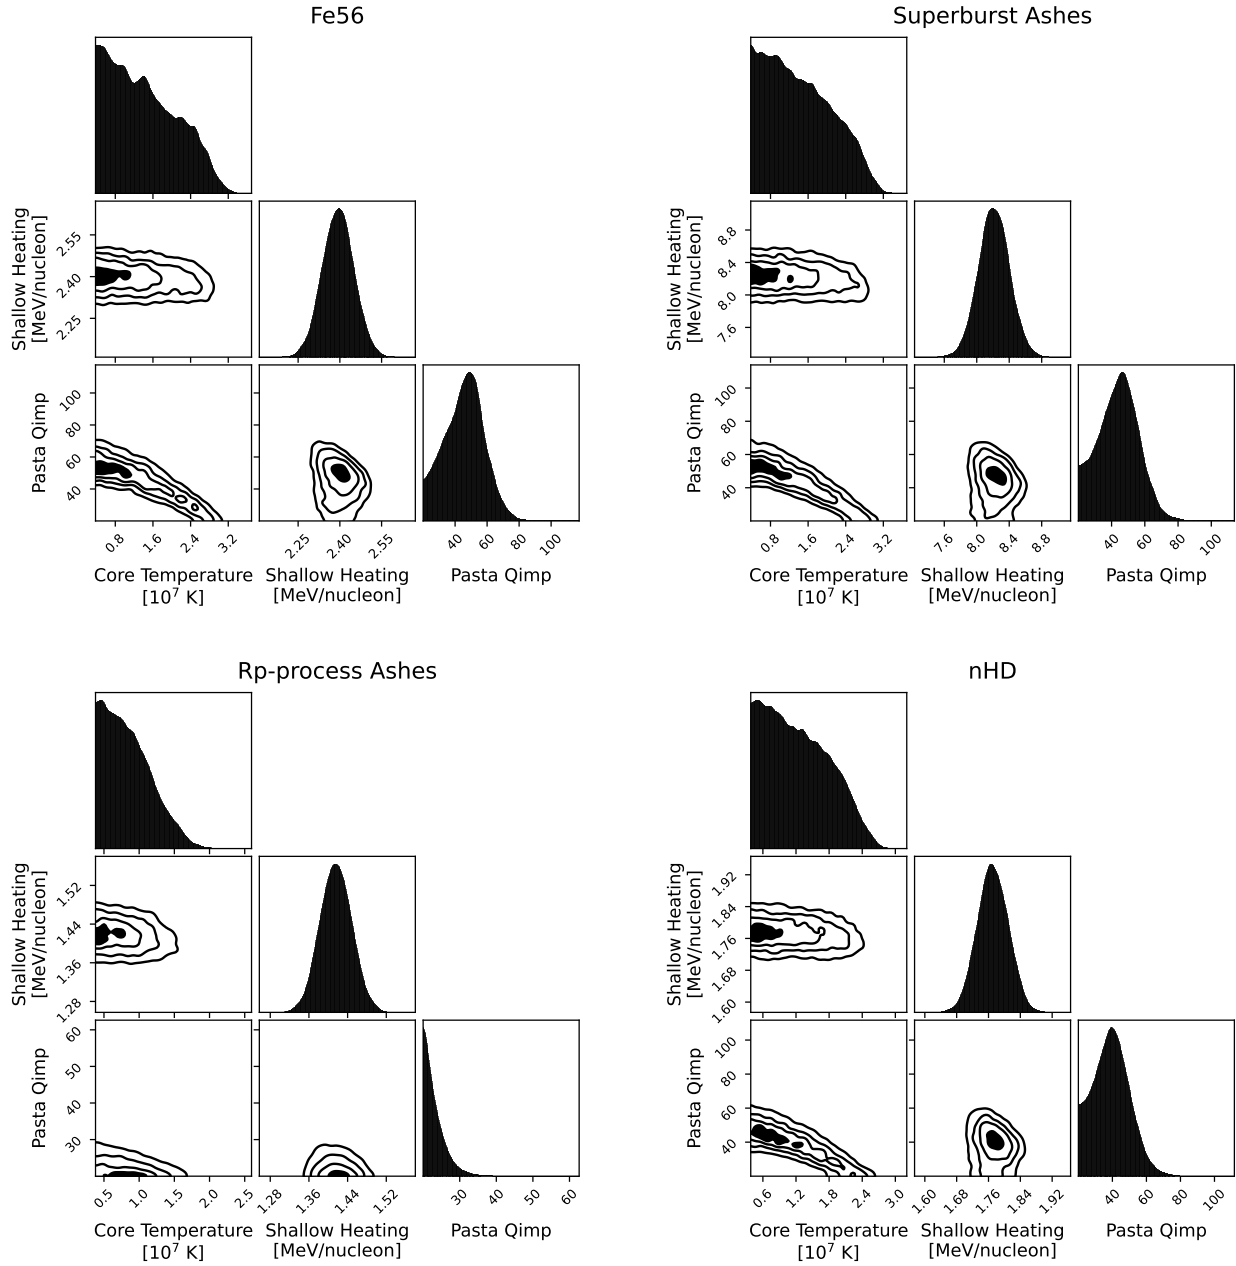


Figure A.4 Posterior distributions of  $T_c$ ,  $Q_{sh}$ , and  $Q_{imp,pasta}$  for MXB 1659–29. The plots display the non-diffusive and nHD models with a pasta layer.

### XTE J1701-462 (No Pasta)

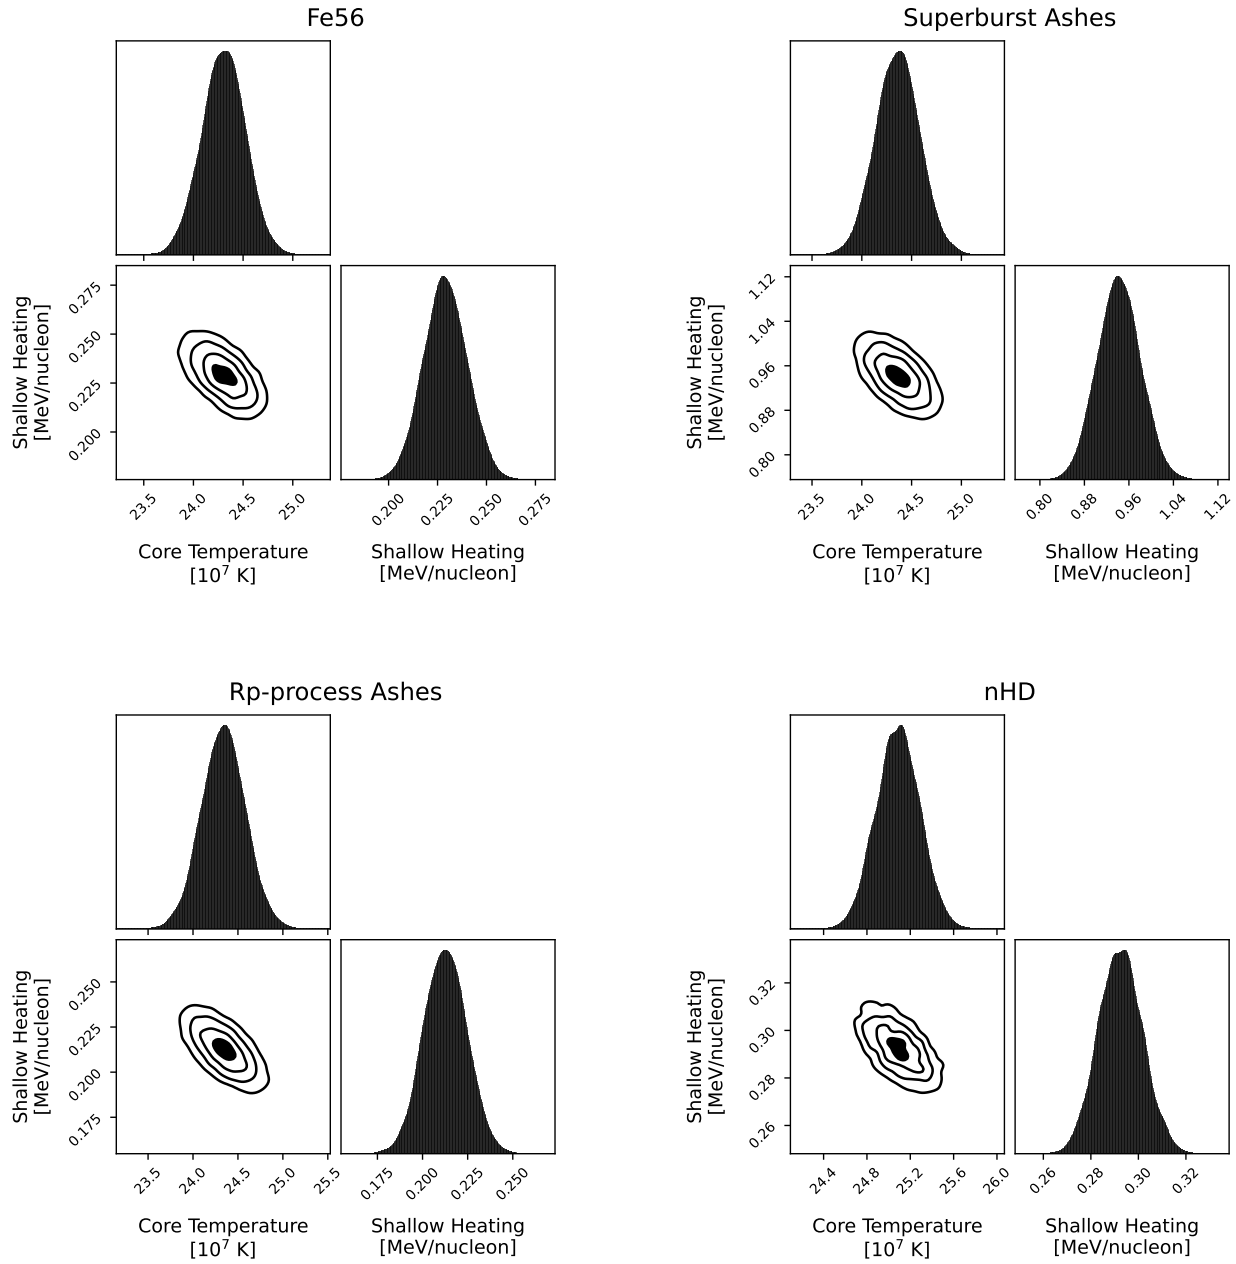


Figure A.5 Posterior distributions of  $T_c$  and  $Q_{sh}$  for XTE J1701-462. The non-diffusive and nHD models without a pasta layer are shown.

### XTE J1701-462 (Pasta)

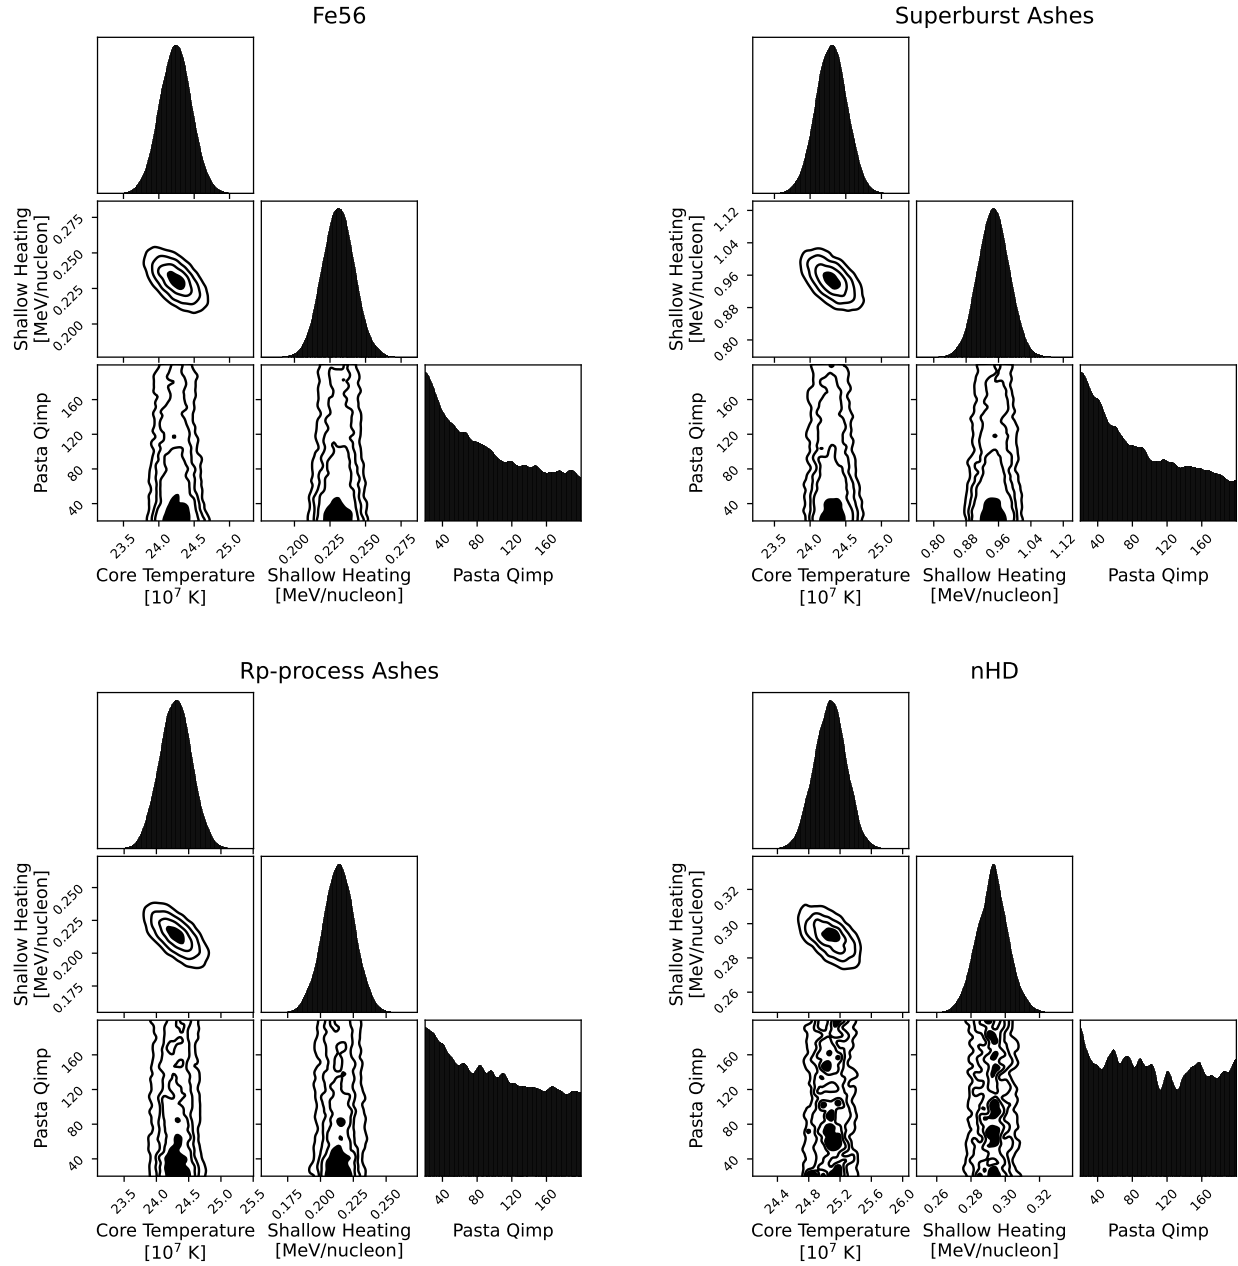


Figure A.6 Posterior distributions of  $T_c$ ,  $Q_{sh}$ , and  $Q_{imp,pasta}$  for XTE J1701-462. Each non-diffusive model and the nHD model, both with a pasta layer, are shown.

### IGR J17480-2446 (No Pasta)

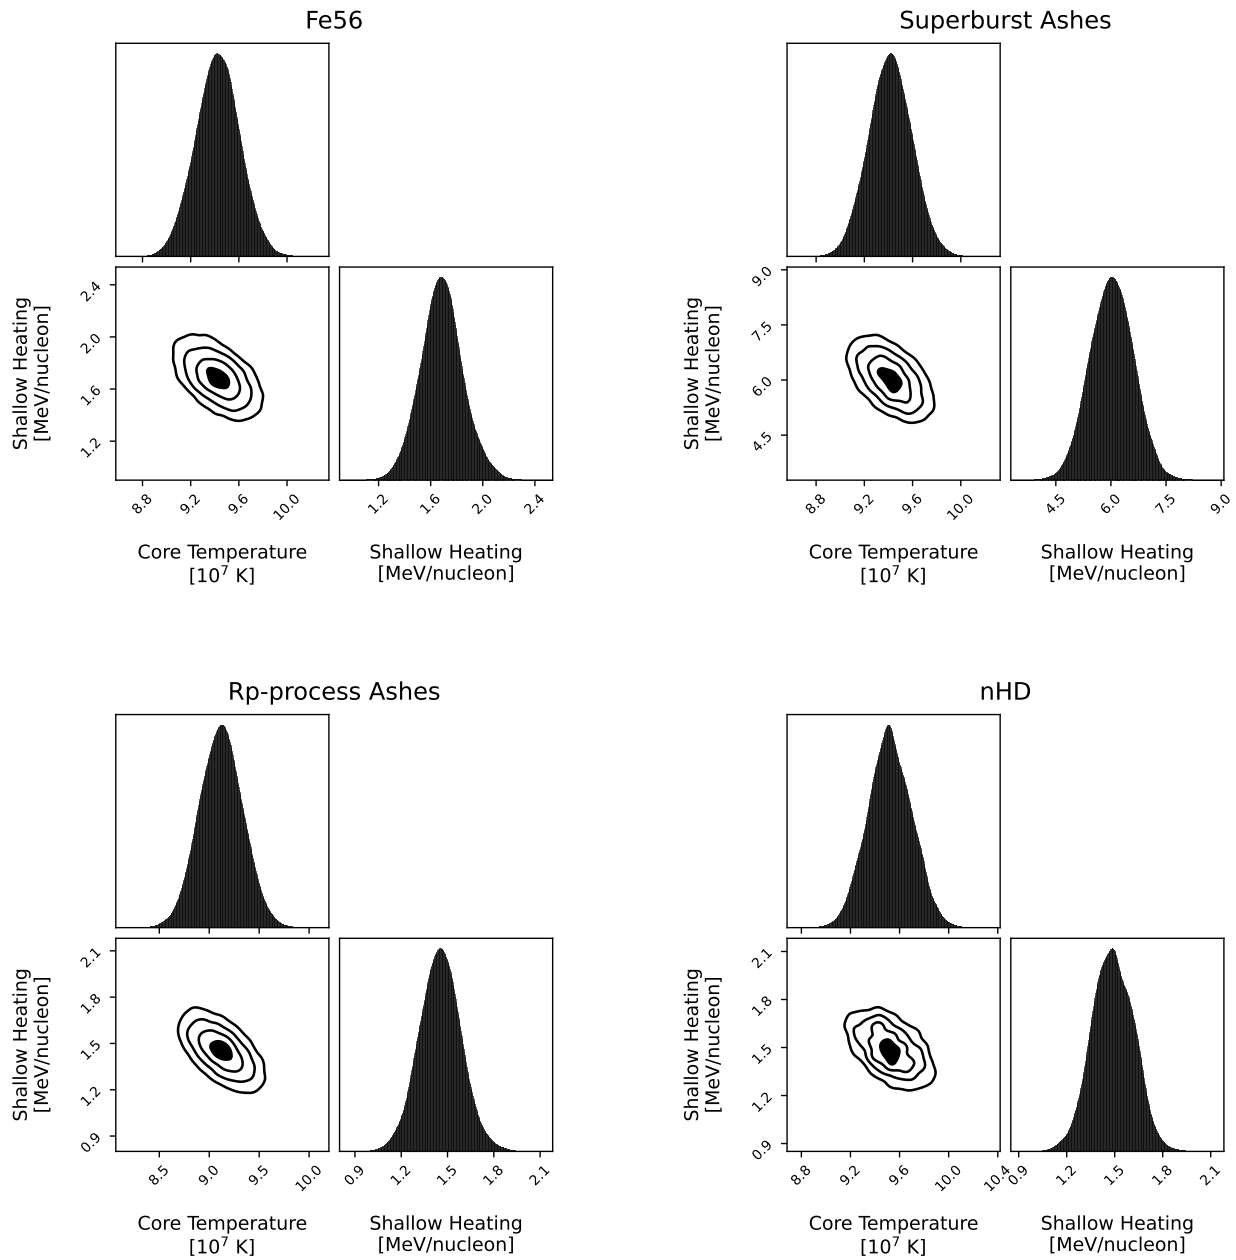


Figure A.7 Posterior distributions of  $T_c$  and  $Q_{sh}$  for IGR J17480–2446. The plots show the non-diffusive and the nHD models without a pasta layer.



### IGR J17480-2446 (Pasta)

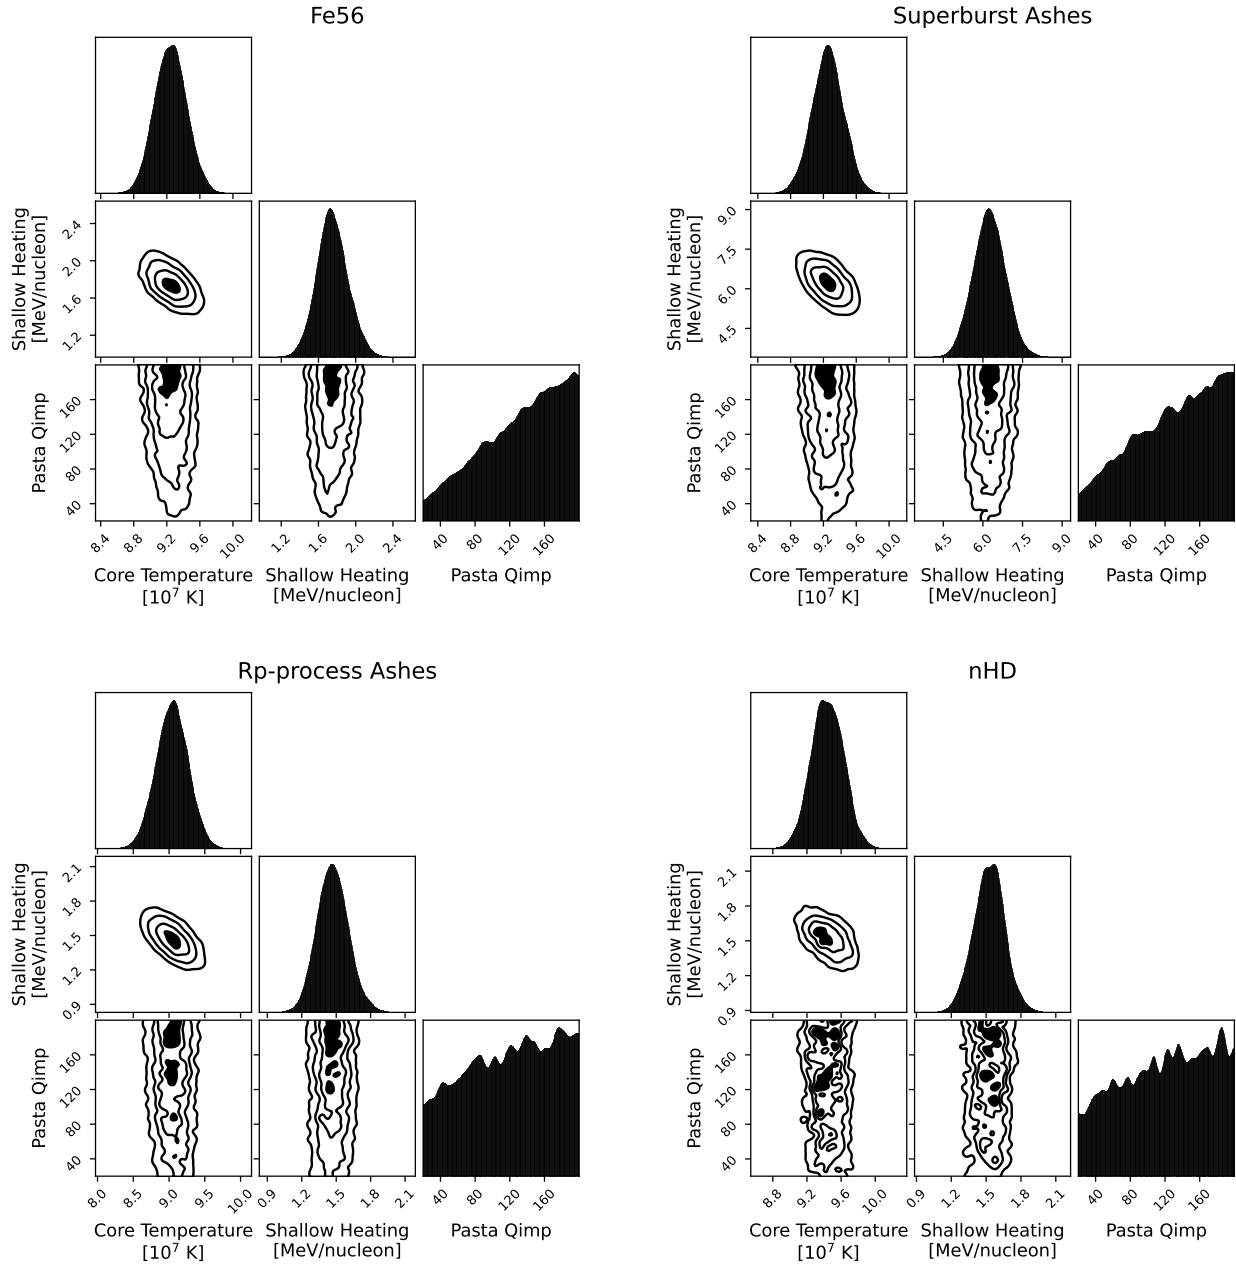


Figure A.8 Posterior distributions of  $T_c$ ,  $Q_{sh}$ , and  $Q_{imp,pasta}$  for IGR J17480–2446. Each non-diffusive model and the nHD model with a pasta layer are displayed.

### Swift J174805.3-244637 (No Pasta)

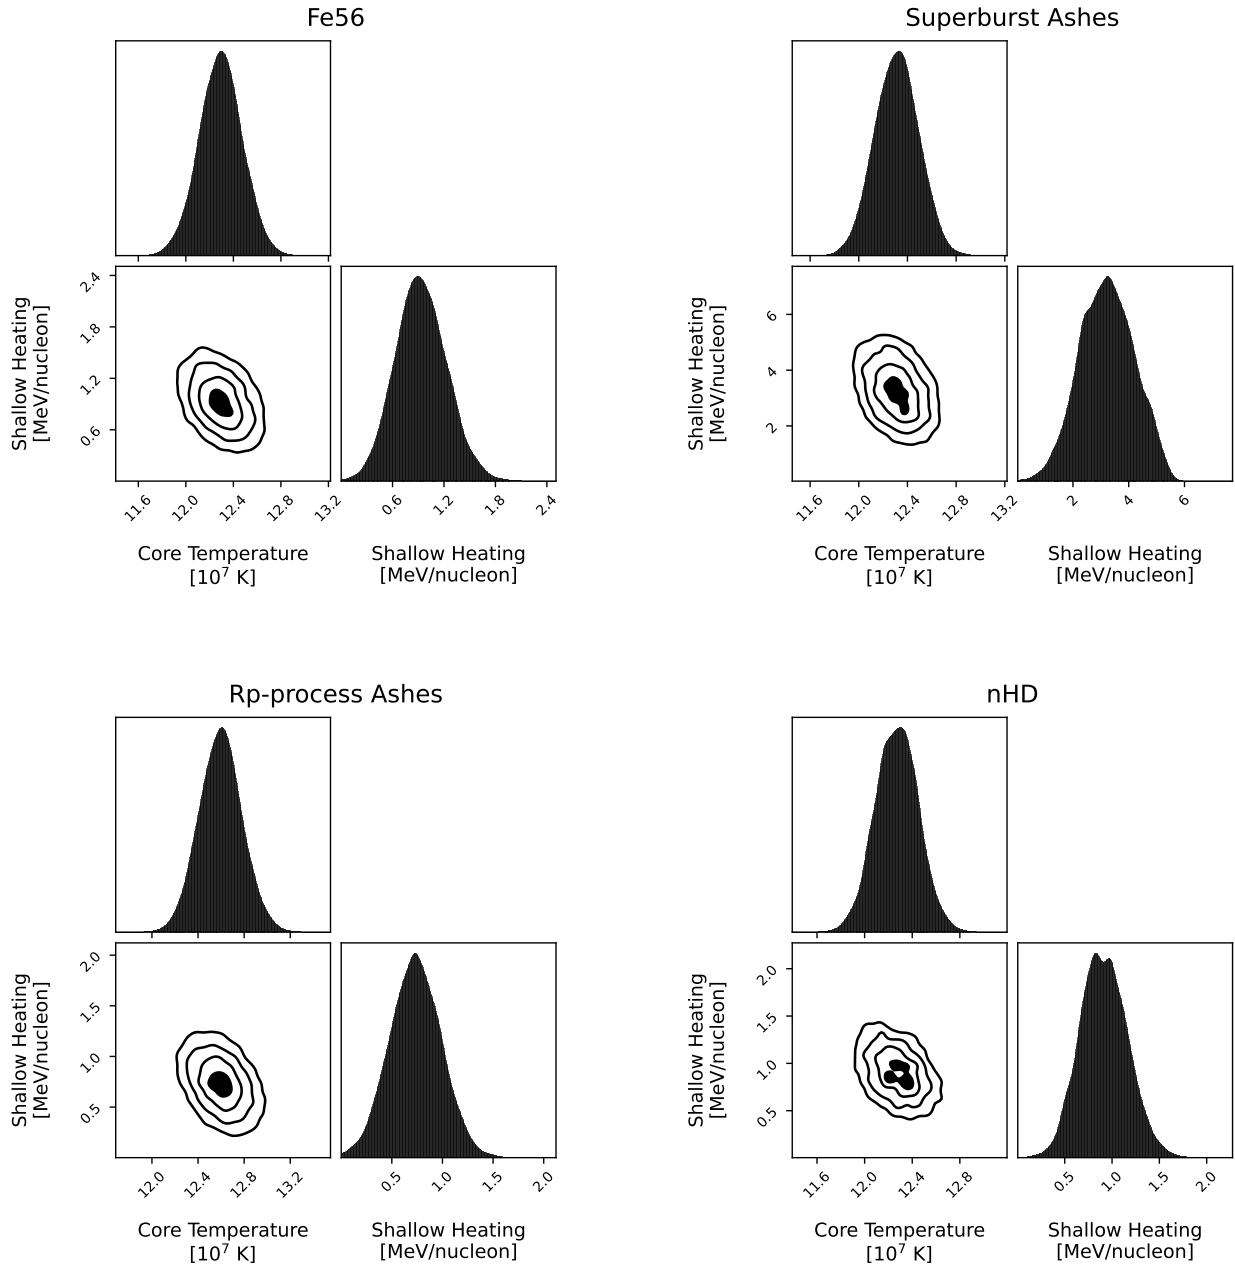


Figure A.9 Posterior distributions of  $T_c$  and  $Q_{sh}$  for Swift 174805.3–244637. The plots display the non-diffusive and the nHD models without a pasta layer.

### Swift J174805.3-244637 (Pasta)

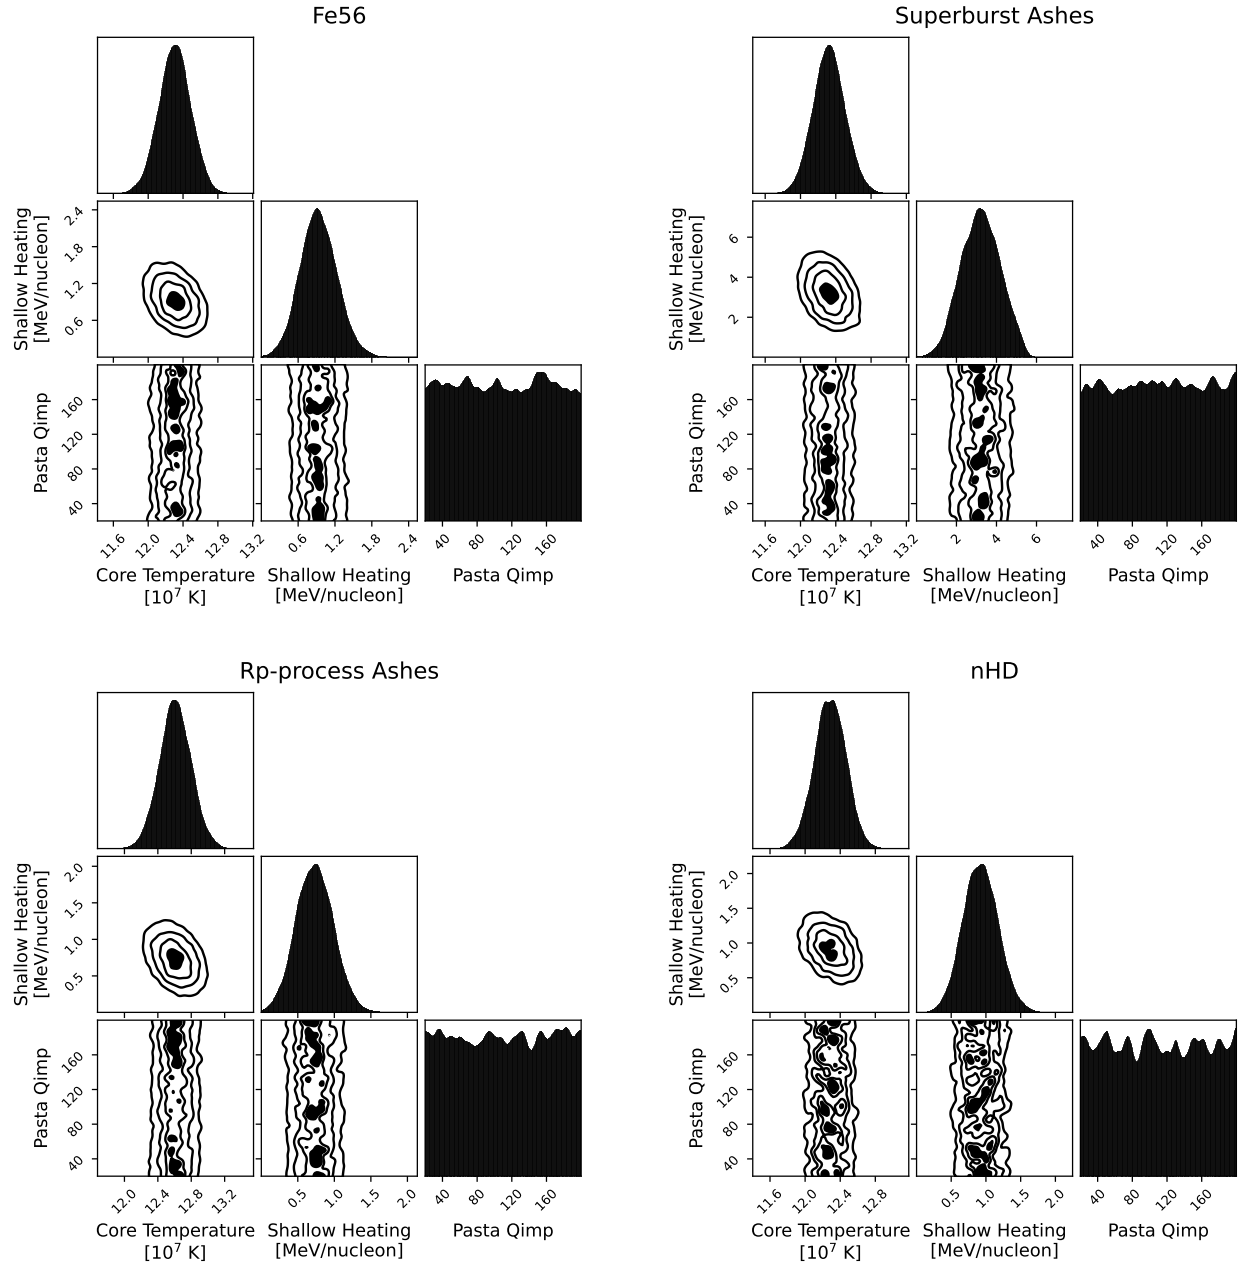


Figure A.10 Posterior distributions of  $T_c$ ,  $Q_{sh}$ , and  $Q_{imp,pasta}$  for Swift 174805.3–244637. The non-diffusive and nHD models with a pasta layer are shown.