



Investigating the Lower Mass Gap with Low-mass X-Ray Binary Population Synthesis

Jared C. Siegel^{1,2,3}, Ilija Kiato^{2,4}, Vicky Kalogera^{2,5}, Christopher P. L. Berry^{2,6}, Thomas J. Maccarone⁷, Katelyn Breivik⁸, Jeff J. Andrews^{2,9}, Simone S. Bavera¹⁰, Aaron Dotter², Tassos Fragos¹⁰, Konstantinos Kovelakas¹⁰, Devina Misra¹⁰, Kyle A. Rocha^{2,5}, Philipp M. Srivastava^{2,11}, Meng Sun², Zepei Xing¹⁰, and Emmanouil Zapartas¹²

¹Department of Astronomy and Astrophysics, University of Chicago, Chicago, IL 60637, USA; siegeljc@princeton.edu

²Center for Interdisciplinary Exploration and Research in Astrophysics (CIERA), Northwestern University, 1800 Sherman Ave, Evanston, IL 60201, USA

³Department of Astrophysical Sciences, Princeton University, 4 Ivy Lane, Princeton, NJ 08544, USA

⁴Department of Physics, University of Illinois at Urbana-Champaign, 1110 W Green St, Urbana, IL 61801, USA

⁵Department of Physics and Astronomy, Northwestern University, 2145 Sheridan Ave, Evanston, IL 60208, USA

⁶SUPA, School of Physics and Astronomy, University of Glasgow, Kelvin Building, University Ave, Glasgow G12 8QQ, UK

⁷Department of Physics & Astronomy, Texas Tech University, Box 41051, Lubbock, TX 79409, USA

⁸Center for Computational Astrophysics, Flatiron Institute, 162 Fifth Ave, New York, NY 10010, USA

⁹Department of Physics, University of Florida, 2001 Museum Rd., Gainesville, FL 32611, USA

¹⁰Departement d'Astronomie, Université de Genève, Chemin Pegasi 51, CH-1290 Versoix, Switzerland

¹¹Electrical and Computer Engineering, Northwestern University, 2145 Sheridan Road, Evanston, IL 60208, USA

¹²IAASARS, National Observatory of Athens, Penteli, 15236, Greece

Received 2022 September 14; revised 2023 July 18; accepted 2023 July 19; published 2023 September 8

Abstract

Mass measurements from low-mass black hole X-ray binaries (LMXBs) and radio pulsars have been used to identify a gap between the most massive neutron stars (NSs) and the least massive black holes (BHs). BH mass measurements in LMXBs are typically only possible for transient systems: outburst periods enable detection via all-sky X-ray monitors, while quiescent periods enable radial velocity measurements of the low-mass donor. We quantitatively study selection biases due to the requirement of transient behavior for BH mass measurements. Using rapid population synthesis simulations (COSMIC), detailed binary stellar-evolution models (MESA), and the disk instability model of transient behavior, we demonstrate that transient LMXB selection effects introduce observational biases, and can suppress mass-gap BHs in the observed sample. However, we find a population of transient LMXBs with mass-gap BHs form through accretion-induced collapse of an NS during the LMXB phase, which is inconsistent with observations. These results are robust against variations of binary evolution prescriptions. The significance of this accretion-induced collapse population depends upon the maximum NS birth mass $M_{\text{NS,birth-max}}$. To reflect the observed dearth of low-mass BHs, COSMIC and MESA models favor $M_{\text{NS,birth-max}} \lesssim 2M_{\odot}$. In the absence of further observational biases against LMXBs with mass-gap BHs, our results indicate the need for additional physics connected to the modeling of LMXB formation and evolution.

Unified Astronomy Thesaurus concepts: X-ray transient sources (1852); Low-mass x-ray binary stars (939); Stellar-mass black holes (1611); Stellar evolutionary models (2046)

1. Introduction

Black hole X-ray binaries (XRBs) are composed of a black hole (BH) accreting mass from a nondegenerate donor star: either through Roche-lobe overflow (RLO) or captured winds. There are currently >50 candidate BH XRB systems (Remillard & McClintock 2006; Corral-Santana et al. 2016), broadly divided by donor mass into high-mass XRBs (HMXB) and low-mass XRBs (LMXB). HMXBs host donor stars of masses $M_{\text{donor}} \gtrsim 5M_{\odot}$ and predominantly transfer mass to the compact object via strong winds, while LMXBs host donor stars of masses $M_{\text{donor}} \lesssim 2M_{\odot}$ and transfer mass via RLO. XRBs offer powerful insights into binary evolution, compact objects, and accretion disks. The XRB Cygnus X-1 provided the first evidence for the existence of a BH (Bolton 1972), and since then X-ray and optical-infrared observations have allowed the identification of tens of transient BH systems (Corral-Santana et al. 2016).

In cases of RLO (most commonly in LMXBs), mass transferred from the donor star forms an accretion disk around the BH. Depending on the structure and composition of the system, a thermal instability may form within the disk and generate transient X-ray emission, cycling through periods of bright outburst and dim quiescence (Cannizzo et al. 1982; King et al. 1996; Lasota et al. 2008). Transient systems are prime targets for follow-up radial velocity measurements of the donor star: periods of bright outburst make the system discoverable in X-rays, while periods of quiescence allow for inference of the accretor's mass via uncontaminated radial velocity measurements of the donor star.

Mass measurements provide key insights into the formation and evolution of BHs, and the majority of Galactic BHs with mass estimates are transient XRBs. Exceptions include the X-ray-faint binary VFTS 243, which Shenar et al. (2022) proposed is composed of a $25M_{\odot}$ O-type star and a $>9M_{\odot}$ BH, as well as the proposed noninteracting low-mass binary companions to 2MASS J05215658+4359220 (Thompson et al. 2019) and V723 Mon (Jayasinghe et al. 2021); however, V723 Mon is consistent with a stripped-star companion instead of a compact object (El-Badry et al. 2022). These systems are outnumbered by the ~ 20 transient LMXBs currently with dynamical mass measurements (Corral-Santana et al. 2016).

Using the mass measurements of the observed LMXB population, prior studies have attempted to constrain the underlying stellar-mass BH distribution (Bailyn et al. 1998; Özel et al. 2010; Farr et al. 2011). Each of these investigations identified a gap between the lowest mass BHs ($\gtrsim 5M_{\odot}$) and the maximum mass ($2\text{--}3M_{\odot}$) for nonrotating neutron stars (NSs; Rhoades & Ruffini 1974; Kalogera & Baym 1996; Müller & Serot 1996; Özel & Freire 2016; Margalit & Metzger 2017; Ai et al. 2020; Shao et al. 2020; Raaijmakers et al. 2021); this feature is commonly referred to as the *lower mass gap*.

Understanding the nature of the purported lower mass gap offers key insights into the properties of degenerate matter and supernova explosion mechanisms (Fryer et al. 2012; Mandel & Müller 2020; Zevin et al. 2020; Liu et al. 2021; Patton et al. 2022). If instability growth and launch of the supernova proceeds on rapid timescales (~ 10 ms), accretion onto proto-NSs is suppressed, and simulations predict a paucity of compact objects born with masses between 2 and $5M_{\odot}$ (Fryer et al. 2012; Belczynski et al. 2012; Fryer et al. 2022); however, if instability growth is delayed (timescales greater than ~ 200 ms), simulations predict a continuous compact-object birth-mass distribution. Motivated by the observed Galactic LMXB sample, prior studies have typically favored the rapid instability timescale, but an underlying assumption of these studies is that BH mass measurements are unbiased.

Recent detections of BH masses in systems other than LMXBs complicate the picture of the lower mass gap. The binary sources of gravitational waves GW190814 and GW200210_092254 are each inferred to have a source containing a compact object lying in the mass gap: $2.59^{+0.08}_{-0.09}M_{\odot}$ and $2.83^{+0.47}_{-0.42}M_{\odot}$, respectively (Abbott et al. 2020, 2021a). The mass estimates for the compact-object companions to 2MASS J05215658+4359220 and V723 Mon (assuming a compact-object interpretation) also lie within the mass gap: $3.3^{+2.8}_{-0.7}M_{\odot}$ and $(3.04 \pm 0.06)M_{\odot}$, respectively (Thompson et al. 2019; Jayasinghe et al. 2021). These measurements are in apparent contradiction with the existence of a lower mass gap, and they must be reconciled with the Galactic LMXB BH mass distribution.

Understanding the nature of detection biases in the Galactic LMXB population is critical to interpreting the observed BH mass distribution. Through approximate scaling of the observed XRB sample, Özel et al. (2010) argued detection biases due to outburst flux cannot account for the mass gap. Jonker et al. (2021) noted that one potential selection effect comes from the fact that objects with high extinction are more difficult to obtain follow-up radial velocity measurements for. Assuming an anticorrelation between supernova kick magnitude and compact-object mass (e.g., Fryer et al. 2012; Atri et al. 2019), this preference for LMXBs off the Galactic plane may bias against mass inference of higher-mass BHs. While this effect alone would not produce a mass gap as an observational selection effect, it may still shape the observed BH mass distribution. Kreidberg et al. (2012) also demonstrated that assuming zero or constant emission from the accretion flow can lead to systematic overestimates of compact-object masses. The role of detection biases in the Galactic LMXB population, in particular the requirement of transient behavior for BH mass measurement, remains unclear.

Here, we adopt a forward-modeling approach to assess selection biases: generating samples of synthetic binary systems, evolving the systems forward in time, and inferring

the detectable population using the disk instability model of LMXB transient behavior. With these methods, we investigate the impact of observational biases on the lower mass gap and the underlying BH mass distribution. Through our analysis, we uncover that the observed gap has implications for the maximum NS mass at birth ($M_{\text{NS,birth-max}}$).

In Section 2.1 and Section 2.3, we describe our population synthesis (COSMIC; Breivik et al. 2020) and stellar evolution (MESA; Paxton et al. 2011, 2013, 2015, 2018, 2019) methods, respectively; the population synthesis Milky Way model is outlined in Section 2.2. In Section 2.4, we outline our treatment of LMXB transient behavior, and in Section 2.5, we present the formulation of detection probabilities. In Section 2.6, we detail how we draw together a synthetic population. We describe our results in Section 3, discuss their implications in Section 4, and summarize our conclusions in Section 5.

2. LMXB Simulations

Here, we outline the methods for simulating LMXB samples with rapid binary population synthesis (COSMIC; Section 2.1) and binary stellar-evolution models (MESA; Section 2.3). While MESA provides detailed modeling, the computational efficiency of COSMIC allows us to explore uncertain aspects of binary evolution physics. Taken together, the suite of proof-of-concept COSMIC simulations contextualize the results of improved mass-transfer modeling with MESA, by marginalizing over uncertain aspects of binary evolution physics.

2.1. Population Models

Binary population synthesis offers valuable insights into binary evolution channels and the expected distributions of target populations. We employ the rapid binary population synthesis code COSMIC (Breivik et al. 2020) to generate Milky Way LMXB populations. Originating from BSE (Hurley et al. 2002), COSMIC approximates single-star evolution using fitting formulae to detailed models and applies prescriptions for binary evolution to these formulae. To complement these simulations, we also consider a three-dimensional grid (in compact-object mass M_{CO} , donor mass M_{donor} , and orbital period P space) of LMXBs using MESA (Section 2.3). The computational efficiency of COSMIC allows us to explore uncertain aspects of binary evolution physics that would not be computationally feasible with MESA.

Independent of contextualizing MESA simulations, verifying the extent to which binary population synthesis methods can replicate the observed LMXB sample is of interest. For example, Ivanova (2006) proposed that some LMXBs could be fed by pre-main-sequence donor stars, given the observed relation between LMXBs' orbital periods and the donors' effective temperatures, as well as the fact that only stars more massive than $\sim 2M_{\odot}$ have a pre-main-sequence lifetime shorter than the main-sequence lifetime of a star that explodes and forms a BH ($\lesssim 10^7$ yr). However, no rapid population synthesis codes simulate pre-main-sequence interactions. These constraints on the binary population suppress the formation of LMXBs with massive BHs and shorter-period binaries. Any rapid population synthesis investigation of LMXBs that does not consider pre-main-sequence interactions is impacted by these limitations. While population synthesis methods likely cannot reproduce all aspects of the observed LMXB sample, quantifying these discrepancies is vital to developing improved

population synthesis engines. Moreover, we are employing population synthesis as a proof-of-concept survey of binary evolution prescriptions: rather than replicating all aspects of the observed LMXB distribution, we are interested in uncovering any significant dependencies between the synthetic observed LMXB sample and uncertain aspects of binary evolution physics.

To generate Galactic LMXB population synthesis samples, we simulate Milky Way stellar populations for a range of binary evolution prescriptions. For each COSMIC population:

1. The maximum NS mass is set to $3M_{\odot}$ (Section 1; Rhoades & Ruffini 1974; Kalogera & Baym 1996).
2. Compact-object birth masses from core-collapse supernovae are assigned following the *Delayed* prescription of Fryer et al. (2012); this differs from the alternative *Rapid* prescription in the assumed timescale of instability growth and launch of the supernova. The Delayed remnant prescription produces a continuous mass distribution, while the Rapid prescription produces a gap between the most massive NSs and the least massive BHs. The Delayed prescription is adopted to explore whether LMXB selection biases can hide an underlying population of low-mass BHs.
3. Compact-object natal kicks are drawn from a bimodal distribution: standard iron core-collapse supernova kicks are drawn from a Maxwellian distribution with a dispersion of 265 km s^{-1} (Hobbs et al. 2005), while kicks for electron-capture supernovae and ultra-stripped supernovae are drawn according to a Maxwellian distribution with a dispersion of 20 km s^{-1} (Giacobbo & Mapelli 2019). Natal kicks are reduced by a factor of $1 - f_{\text{fb}}$, where f_{fb} is the fraction of the ejected supernova mass that will fall back onto the proto-compact object (Fryer et al. 2012).
4. Initial conditions (e.g., eccentricity, orbital period, primary mass, and mass ratio) are independently drawn following Sana et al. (2012).

We investigate model dependence on six uncertain factors of binary evolution:

1. Stars without a distinct core–envelope boundary (e.g., stars on the Hertzsprung gap) that instigate a common envelope (CE) event are either assumed to survive the CE (*Optimistic*) or are assumed to merge (*Pessimistic*; Belczynski et al. 2008).
2. The CE efficiency α parameterizes the transfer of orbital energy to the envelope, i.e., how easily the CE is unbound from the system (Webbink 1984; de Kool 1990); higher values of α lead to longer-period binaries post-CE. Prior studies using one-dimensional simulations (Fragos et al. 2019) and comparison to gravitation-wave observations (Giacobbo & Mapelli 2018; Santoliquido et al. 2021; Zevin et al. 2021) have favored higher CE efficiencies ($\alpha \sim 5$). To explore these higher values, α is set to $\alpha = 1.0$ or $\alpha = 5.0$.
3. The minimum ZAMS mass ratio between binary members is set to $q > 0.01$, where $q \equiv M_{\text{donor}}/M_{\text{primary}}$, or is defined as a function of primary mass (*Lifetime-limited*). For the Lifetime-limited case, q is restricted such that the pre-main-sequence lifetime of the secondary is shorter than the main-sequence lifetime of the primary.

4. In addition to gravitational radiation, close-in binaries (e.g., cataclysmic variables and LMXBs) can efficiently lose orbital angular momentum via magnetic braking of a tidally coupled magnetic wind. Magnetic braking is implemented following Hurley et al. (2002, hereafter H02) or Ivanova & Taam (2003, hereafter IT03).
5. The amount of mass accreted during Roche-lobe overflow is either fixed to 50% efficiency (Belczynski et al. 2008) or is a function of the accretor’s type (*State-dependent*). For the State-dependent prescription, the accretion efficiency during Roche-lobe overflow is either (i) limited to ten times the thermal rate, if the accretor is a main-sequence, Hertzsprung gap, or core helium burning star; or (ii) unlimited, if the accretor is a giant branch, early asymptotic giant branch, or asymptotic giant branch star. Compact objects also experience conservative mass transfer, provided the mass-transfer rate is sub-Eddington. Accretion onto an NS is treated independently of whether the system is transient. While outburst events during accretion onto a white dwarf can lead to a loss of mass, only a small fraction of accreting NSs show evidence of mass loss during bursts (Degenaar et al. 2018); this difference between white dwarf and NS accretors is likely attributable to the deeper potential wells of NSs.
6. The maximum NS birth mass from core-collapse supernova $M_{\text{NS,birth-max}}$ is of particular interest to the formation and evolution of LMXBs; lowering $M_{\text{NS,birth-max}}$ suppresses the formation of LMXBs through accretion-induced collapse (AIC) of an NS. To vary $M_{\text{NS,birth-max}}$ while preserving computational feasibility, $M_{\text{NS,birth-max}}$ is fixed to $3M_{\odot}$ within COSMIC and is lowered in post-processing by removing systems with an NS birth mass above a given threshold. This method of imposing a maximum NS birth mass is a nonphysical toy model. The goal is to demonstrate how a truncated NS birth-mass distribution impacts LMXB demographics.

To simulate stellar populations analogous to the Milky Way, we assign birth times, spatial positions, and metallicities following a three-component Milky Way model (Section 2.2).

2.2. Milky Way Model

For a given set of binary evolution prescriptions, we generate a synthetic Galactic sample of LMXBs. To do so, we resample a given COSMIC population and assign system properties in accordance with the Milky Way (e.g., star formation history and spatial distribution). The Milky Way is approximated as a composite of three subpopulations: thin disk, thick disk, and bulge.

The star formation history and metallicity of each component are adopted from Robin et al. (2003). For the thick disk and bulge, star formation is modeled as a 1 Gyr continuous burst, beginning at the birth of the Milky Way and 1 Gyr after the birth of the Milky Way, respectively. Star formation for the thin disk is continuous from 1 Gyr after the birth of the Milky Way to the present. Stars in the thin disk and bulge have solar metallicity, while stars in the thick disk have 15% of solar metallicity. System birth times are relative to a Milky Way age of $\mathcal{A}_{\text{MW}} = 11 \text{ Gyr}$. LMXBs’ positions within the galaxy are

Table 1
Milky Way Component Parameters

Component	Age ^a (Gyr)	Metallicity [Fe/H]	Star Formation History	Stellar Mass ($10^{10}M_{\odot}$)	Spatial Distribution Parameters
Thin disk	$\mathcal{A}_{\text{MW}}-1$ Gyr	0.0	Continuous	4.32	$z_{\text{disk}} = 0.3$ kpc, $R_{\text{disk}} = 2.90$ kpc
Thick disk	\mathcal{A}_{MW}	-0.8	1 Gyr burst	1.44	$z_{\text{disk}} = 0.9$ kpc, $R_{\text{disk}} = 3.31$ kpc
Bulge	$\mathcal{A}_{\text{MW}}-1$ Gyr	0.0	1 Gyr burst	0.89	$a = 1.8$, $r_0 = 0.075$ kpc, $r_{\text{cut}} = 2.1$ kpc, $q = 0.5$

Notes. Component masses and spatial distributions follow McMillan (2011). Star formation histories and metallicities are adopted from Robin et al. (2003).

^a The adopted Milky Way age is $\mathcal{A}_{\text{MW}} = 11$ Gyr.

fixed to their birthplaces; we do not include the effects of SN kicks on systemic velocity.

The component masses and spatial distributions for the Milky Way model follow McMillan (2011). The thin and thick disk spatial distributions are treated as an exponential,

$$\rho_{\text{disk}}(R, z) \propto \exp\left(-\frac{|z|}{z_{\text{disk}}} - \frac{R}{R_{\text{disk}}}\right), \quad (1)$$

where z_{disk} is the scale height and R_{disk} is the scale length; for the thin disk, $z_{\text{disk}} = 0.3$ kpc and $R_{\text{disk}} = 2.9$ kpc; for the thick disk, $z_{\text{disk}} = 0.9$ kpc and $R_{\text{disk}} = 3.31$ kpc (McMillan 2011). The bulge spatial distribution is

$$\rho_{\text{bulge}}(R, z) \propto \frac{\exp[-(r'/r_{\text{cut}})^2]}{(1 + r'/r_0)^a}, \quad (2)$$

$$r' = \sqrt{R^2 + (z/q)^2}, \quad (3)$$

where $a = 1.8$, $r_0 = 0.075$ kpc, $r_{\text{cut}} = 2.1$ kpc, and the axial ratio $q = 0.5$. Each component is treated as axisymmetric.

The parameters of the Milky Way model are summarized in Table 1.

2.3. MESA Models

Rapid binary population synthesis enables the efficient modeling of statistically significant stellar samples, but such methods have drawbacks. Rapid binary population synthesis codes, including COSMIC, assume binary members have the same properties as a single star with matching mass and metallicity at thermal equilibrium; this can lead to systematic biases in the binaries' evolution (Gallegos-Garcia et al. 2021; Fragos et al. 2023). As outlined above, the COSMIC simulations serve as a proof-of-concept survey over uncertain aspects of binary evolution physics and contextualize a smaller grid of binaries modeled in greater detail with MESA. Here, we outline the LMXB grid evolved using MESA.

We consider a three-dimensional grid of binaries consisting of a compact object and a hydrogen-rich main-sequence star. Adopted from Fragos et al. (2023), the grid consists of 1375 binaries, initialized uniformly in $\log_{10}(M_{\text{CO}}/M_{\odot})$, $\log_{10}(M_{\text{donor}}/M_{\odot})$, and $\log_{10}(P/\text{day})$. Compact-object masses span the range $1.1-10M_{\odot}$, donor star masses span $0.5-3M_{\odot}$, and orbital periods span $1.26-3162$ days. The grid includes LMXBs with mass-gap BH, which allows us to investigate the observability of such sources. The maximum NS mass is set to $3M_{\odot}$ (as in the COSMIC simulations). All systems are initialized at solar metallicity, and the donor stars are seeded with a rotational period equal to their orbital period (Fragos et al. 2023).

To facilitate comparison between the population synthesis and MESA methods, we initialize and evolve grids of binaries

with COSMIC at the same initial conditions as the MESA grid. We consider the same combinations of binary evolution prescriptions as in Section 2.1, with the exception of CE efficiency and minimum ZAMS mass ratio (which do not apply to this phase of evolution).

2.4. Transient LMXB Behavior

2.4.1. Instability Cycle

LMXB accretion disks are complex rotating, turbulent, irradiated systems. Despite this complexity, the outburst cycles of LMXBs are reasonably well approximated by the disk instability model (King et al. 1996; Dubus et al. 1999; Lasota 2001, 2016). This simplified model of accretion physics assumes: (i) a constant mass-transfer rate, determined by the secular mass-transfer rate from the binary evolution calculations; (ii) that each outburst is identical, a consequence of the assumed accretion of the entire disk during an outburst; and (iii) that the mass transfer is conservative when sub-Eddington. Observational work shows these assumptions are imperfect (e.g., the observed bimodality of outburst durations; Lasota 2001, Section 4.1). The disk instability model is adopted for this work because it is quantitatively tractable, and efforts to improve the model are beyond the scope of this paper.

Under the disk instability model, a disk is thermally stable provided radiative cooling varies faster with temperature than viscous heating (Frank et al. 1992; Dubus et al. 1999). As a result, a rapid change in opacity instigates a thermal instability in the disk (e.g., if a disk is composed of ionized hydrogen and the temperature falls low enough for hydrogen to recombine). If thermally unstable, a disk will cycle through hot outburst states and cold quiescent states.

For a given LMXB, the conditions for thermal equilibrium of the accretion disk (typically presented in the disk surface density Σ and effective temperature T_{eff} plane) can be numerically calculated as a function of disk radius. Equilibrium solutions in the $\Sigma-T_{\text{eff}}$ plane (*S-curves*) are composed of a stable hot branch and a stable cold branch connected by an instability strip (Lasota 2001, Section 3). LMXBs where the mass-transfer rate results in a T_{eff} on the instability strip will be thermally unstable, leading the disk to cycle between the hot and cold branches. The inflection points of the S-curve correspond to critical accretion rates that bound the instability strip. By calculating S-curves for a range of disk radii and binary system parameters (e.g., orbital period, donor mass, and accretor mass), the cold branch and hot branch bounding critical accretion rates can be interpolated for a given LMXB.

In principle, stable LMXB accretion disks can occupy the cold branch or the hot branch of the S-curve. However, binary evolution favors mass-transfer rates above the cold branch critical mass-transfer rate (Lasota 2001). Systems may occupy

the cold branch late in their evolution, when their donor stars become degenerate around the point that orbital periods begin to increase (Knigge et al. 2011); in such cases, the outbursts would not contribute significantly to the detectable population of sources, because the outbursts would be short, faint, and infrequent. In this study, an LMXB is considered a transient source if its mass-transfer rate falls below the hot branch critical accretion rate $\dot{M}_{\text{donor}} < \dot{M}_{\text{crit}}$.

For systems with hydrogen-dominated donor stars, the hot branch critical mass-transfer rate for X-ray irradiated disks is (Dubus et al. 1999)

$$\dot{M}_{\text{crit}} = 1.5 \times 10^{15} \left(\frac{M_{\text{BH}}}{M_{\odot}} \right)^{-0.4} \left(\frac{R_{\text{disk}}}{10^{10} \text{ cm}} \right)^{2.1} \times \left(\frac{C}{5 \times 10^{-4}} \right) \text{ g s}^{-1}, \quad (4)$$

where M_{BH} is the BH mass, R_{disk} is the outer radius of the accretion disk, and C is a disk structure constant (which we fix to $C = 5 \times 10^{-4}$; Dubus et al. 1999). For helium-dominated donors, where the opacity of the inflow and accretion disk are markedly different from the hydrogen-dominated case, the hot-branch irradiated-disk critical mass-transfer rate is (Lasota et al. 2008)

$$\dot{M}_{\text{crit}} = 2.1 \times 10^{16} \left(\frac{C_{\text{He}}}{10^{-3}} \right)^{-0.22} \times \left(\frac{\alpha}{0.1} \right)^{-0.03-0.01 \log(C_{\text{He}}/10^{-3})} \times \left(\frac{R_{\text{disk}}}{10^{10} \text{ cm}} \right)^{2.51-0.05 \log(C_{\text{He}}/10^{-3})} \times \left(\frac{M_{\text{BH}}}{M_{\odot}} \right)^{-0.74+0.05 \log(C_{\text{He}}/10^{-3})} \text{ g s}^{-1}, \quad (5)$$

where C_{He} is a disk structure constant for helium-dominated donors and α is a viscosity parameter (we fix $C_{\text{He}} = 10^{-3}$ and $\alpha = 0.1$; Lasota et al. 2008). We adopt the standard formulation of disk size: $R_{\text{disk}} = 70\%$ of the BH's Roche-lobe radius (King et al. 1997).

If the mass-transfer rate is below the critical rate, the system will undergo a cycle of outburst and quiescent phases of duration T_{outburst} and $T_{\text{quiescence}}$, respectively. If the mass-transfer rate is above the critical rate, the system will be a persistent X-ray source. For transient LMXBs, the period of the instability cycle is $T_{\text{cycle}} \equiv T_{\text{quiescence}} + T_{\text{outburst}}$. Assuming the entire disk is accreted onto the BH during an outburst event, the duration of the quiescent phase is

$$T_{\text{quiescence}} = \frac{M_{\text{disk-max}}}{\dot{M}_{\text{donor}}}, \quad (6)$$

where \dot{M}_{donor} is the mass-transfer rate from the donor star. Adopting the critical surface density Σ_{crit} of Cannizzo et al. (1988), and assuming the surface density of the quiescent disk prior to infall is comparable to Σ_{crit} (Dubus et al. 2001), the maximum accretion disk mass is

$$M_{\text{disk-max}} \approx \int_{R_{\text{ISCO}}}^{R_{\text{disk}}} 2\pi \Sigma_{\text{crit}}(R) R dR, \quad (7)$$

where R_{ISCO} is the radius of the innermost stable circular orbit.

The duration of the outburst T_{outburst} is the total time necessary for the disk to accrete onto the BH. Because the mass infall rate during the instability $\dot{M}_{\text{ins}}(t)$ varies during the outburst, T_{outburst} cannot be approximated analogously to $T_{\text{quiescence}}$ in Equation (6). Following King & Ritter (1998), the mass infall rate $\dot{M}_{\text{ins}}(t)$ is treated as an exponential decay followed by a linear decay.

The bolometric outburst light curve is related to the infall rate by

$$L_{\text{bol}}(t) = \eta(t) c^2 \dot{M}_{\text{ins}}(t), \quad (8)$$

where $\eta(t)$ is the radiative efficiency and c is the speed of light. Adopting the methods of Kneivt et al. (2014), we consider two potential formulations of $\eta(t)$, corresponding to two different transitions to radiatively inefficient accretion (RIA): a *sharp* transition,

$$\eta = \begin{cases} 0, & L_{\text{bol}} < \nu L_{\text{Edd}} \\ 0.1, & L_{\text{bol}} \geq \nu L_{\text{Edd}} \end{cases}, \quad (9)$$

or a *smooth* transition,

$$\eta = \begin{cases} 0.1 \left(\frac{\dot{M}_{\text{ins}}}{\nu \dot{M}_{\text{Edd}}} \right), & L_{\text{bol}} < \nu L_{\text{Edd}} \\ 0.1, & L_{\text{bol}} \geq \nu L_{\text{Edd}} \end{cases}, \quad (10)$$

where we either fix ν to 0 (no transition to RIA) or 0.05; here, \dot{M}_{Edd} is the Eddington mass-transfer rate, and L_{Edd} is the corresponding Eddington luminosity.

In this approximate outburst model, including a transition to RIA shortens the X-ray outburst lifetime without changing the peak luminosity of the outburst. To highlight this dependence, outburst light curves for a typical LMXB system under varying BH masses and transitions to RIA are presented in Figure 1.

Because the outburst luminosity may drop below a detectable level before the full T_{outburst} has elapsed, we introduce the effective outburst time T'_{outburst} , the total amount of time that the X-ray outburst flux is above an observable threshold. For a limiting X-ray flux of F_{fiducial} , the outburst is observable if

$$F_{\text{fiducial}} < \frac{L_{\text{bol}}(t)}{4\pi d^2 f_{\text{corr}}}, \quad (11)$$

where f_{corr} converts the bolometric flux into the observable range and d is the distance to the system. We adopt $f_{\text{corr}} = 4$ and $F_{\text{fiducial}} = 10 \text{ mCrab}$ (Section 2.5). For COSMIC, d is drawn for each system from the spatial distributions described in Section 2.2; for the MESA grid, d is fixed to 8 kpc (Section 2.6).

In Section 2.5, LMXBs' effective outburst times T'_{outburst} and cycle lifetimes T_{cycle} are mapped to probabilities of detection with an all-sky X-ray monitor.

2.5. Detection Probabilities

Transient X-ray binaries have predominantly been discovered by all-sky monitors (ASMs). Although the observed sample of transient LMXBs has been built up by a myriad of instruments, we treat the ASM on board the Rossi X-ray Timing Explorer (RXTE; Bradt et al. 1993) as our benchmark detector. Active from 1996 to 2012, the RXTE ASM observed in the 1.5–12 keV band with a highly stochastic pointing pattern that typically scanned a given source 5–10 times per




Figure 1. X-ray outburst light curves for a typical LMXB system for a range of BH masses and transitions to RIA. On the left, we assume no transition to RIA and vary the BH mass with fixed orbital period and companion mass. On the right, we present the sharp and smooth transitions to RIA for fixed BH mass. Both panels assume $M_{\text{donor}} = 0.7M_{\odot}$, $P = 10$ hr, and $d = 8$ kpc. The shaded gray region highlights where the X-ray luminosity drops below the detectable limit (10 mCrab). All luminosities have been converted to the 2–10 keV range by dividing the bolometric luminosity by $f_{\text{corr}} = 4$.

day (Levine et al. 1996). RXTE discovered a significant portion of the transient LMXB sample (Wen et al. 2006). The detection requirements of the RXTE ASM can be approximated as a flux limit of $F_{\text{fiducial}} = 10$ mCrab for a minimum emission time of $T_{\text{fiducial}} = 1$ day.

For a transient LMXB source to be identified by an all-sky survey and be eligible for mass inference, it must (i) be bright enough during outburst to be discovered and (ii) transition from an outburst to a quiescent state during the survey. The latter requirement enables mass inference via radial velocity follow-up of the low-mass donor. We neglect inclination effects and potential biases induced by the position of the Sun.

If an outburst occurs during an ASM survey, the probability the outburst is detected by the ASM is (Knevitt et al. 2014)

$$p_{\text{discover}} = \min \left\{ 1, \frac{T'_{\text{outburst}}}{T_{\text{fiducial}}} \right\}, \quad (12)$$

where T'_{outburst} is the effective outburst time: the total amount of time that the X-ray outburst is brighter than F_{fiducial} (Section 2.4).

In addition to the requirement on the system’s outburst flux, an LMXB must also transition from outburst to quiescence during the survey in order to be classified as a transient. The probability of such a transition is

$$p_{\text{transition}} = \min \left\{ 1, \frac{T_{\text{survey}}}{T_{\text{cycle}}} \right\}, \quad (13)$$

where T_{cycle} is the period of the instability cycle, $T_{\text{cycle}} \equiv T_{\text{quiescence}} + T_{\text{outburst}}$, and T_{survey} is the duration of the X-ray monitor’s survey.

Considering both the probability of being seen in outburst and the probability of observing a transition, it follows that the probability of a transient LMXB source being identified is

$$p_{\text{detect}} = p_{\text{discover}}(T'_{\text{outburst}}) \times p_{\text{transition}}(T_{\text{cycle}}). \quad (14)$$

Motivated by RXTE, we adopt $T_{\text{survey}} = 15$ yr. For persistent LMXBs, we set $p_{\text{detect}} = 0$.

Using Equation (14), the probability of identifying an LMXB as a transient source can be inferred from the outburst time T'_{outburst} and cycle lifetime T_{cycle} ; T'_{outburst} and T_{cycle} can in turn be inferred from the evolutionary state of the LMXB, e.g., mass-transfer rate, BH mass, donor mass, and orbital period. In Section 2.6, we leverage Equation (14) to generate synthetic samples of observed LMXBs for both COSMIC and MESA.

2.6. Synthetic Observed Samples

To facilitate comparison between the observed LMXB sample and synthetic populations of LMXBs, a Monte Carlo observing procedure is employed. Samples are drawn from the evolution tracks of COSMIC and MESA, based on the assumed age of the Milky Way, and are then compared with the Galactic LMXB population to investigate the nature of the lower mass gap.

While similar in philosophy, the synthetic observing process differs between the COSMIC Milky Way populations and the MESA grid; the differences are motivated by the far lower number of binaries simulated in the MESA grid. For a given COSMIC Milky Way population, we collect 5000 snapshots of the population within a time window of 10.5–11.5 Gyr, informed by our adopted Milky Way age of $\mathcal{A}_{\text{MW}} = (11 \pm 0.5)$ Gyr. The time of each snapshot is independently and uniformly drawn within the window. At every snapshot, we infer the properties of the population via linear interpolation of the COSMIC time steps. Statistical properties of the synthetic observed sample can then be inferred by considering all 5000 snapshots.

As described in Section 2.3, we also consider a three-dimensional grid of compact-object–main-sequence binaries with MESA, and initialize COSMIC populations with matching initial conditions for a range of binary evolution prescriptions. Given the discrete distribution of initial conditions and the lower number of binaries, synthetic observed samples are drawn from the grids using a modified version of the Monte Carlo procedure applied to the COSMIC populations. For a given trial, we draw a random Milky Way age \mathcal{A}_{MW} from $\mathcal{U}(10.5, 11.5)$ Gyr and draw each binary a birth time independently from $\mathcal{U}(0, \mathcal{A}_{\text{MW}})$; this assignment of binary

birth times is analogous to continuous star formation and is based on the thin disk component of the Milky Way model (Table 1). Linear interpolation is then applied to infer each binary’s properties at \mathcal{A}_{MW} . While distances are assigned via spatial distributions for the COSMIC Milky Way populations, a nominal distance of 8 kpc is adopted for each binary in the MESA grid, due to the low number of binaries. In total, 5000 trails are drawn for each grid.

For a single snapshot, the number of observed transient LMXB systems is

$$N_{\text{transient}} = \sum_{i=1}^{N_{\text{LMXB}}} p_{\text{detect},i}, \quad (15)$$

where N_{LMXB} is the total number of LMXB systems in the population during the snapshot. Systems composed of a BH accreting mass from a nondegenerate donor star via RLO are considered LMXBs, which excludes degenerate donors, in order to mirror the observed sample of LMXBs with BH mass estimates.

To quantify the density of transient LMXB systems within the lower mass gap, we define the fraction

$$f_{\text{low}}(p_{\text{detect}}) = \frac{1}{N_{\text{transient}}} \sum_{i=1}^{N_{\text{LMXB}}} \begin{cases} p_{\text{detect},i}, & M_{\text{BH}} < M_{\text{bound}} \\ 0, & M_{\text{BH}} \geq M_{\text{bound}} \end{cases}. \quad (16)$$

Here, M_{bound} is nominally fixed to $4.5M_{\odot}$.

The fraction $f_{\text{low-intrinsic}}$ is then defined as the unweighted fraction of LMXB systems with a BH mass less than M_{bound} ,

$$f_{\text{low-intrinsic}} = \frac{1}{N_{\text{LMXB}}} \sum_{i=1}^{N_{\text{LMXB}}} \begin{cases} 1, & M_{\text{BH}} < M_{\text{bound}} \\ 0, & M_{\text{BH}} \geq M_{\text{bound}} \end{cases}. \quad (17)$$

This $f_{\text{low-intrinsic}}$ fraction characterizes the *intrinsic* density of LMXBs with BHs within the lower mass gap, while $f_{\text{low}}(p_{\text{detect}})$ characterizes the *observed* density of LMXBs with a mass-gap BH. Considering $f_{\text{low}}(p_{\text{detect}})$ and $f_{\text{low-intrinsic}}$ together highlights the influence of transient detection effects on the observed Galactic LMXB sample.

Following the procedures above, we generate synthetic LMXB samples for each COSMIC population and MESA grid. We discuss these samples and their implications for the lower mass gap in Section 3.

3. Results

Using rapid binary population synthesis (COSMIC) coupled with detailed stellar-structure and mass-transfer models (MESA), we investigate the impact of observational biases on the Galactic LMXB sample. We demonstrate that transient LMXB selection effects do introduce some biases into the observed LMXB sample and that the observed lower mass gap has implications for the maximum NS birth-mass $M_{\text{NS, birth-max}}$.

To survey how well population synthesis methods match the observed LMXB sample, and to explore dependencies between the LMXB sample and uncertain aspects of binary evolution physics, our investigation begins with a suite of COSMIC rapid binary population synthesis simulations. The COSMIC results are discussed in Section 3.1. For 32 combinations of binary evolution prescriptions, Galactic LMXB samples from COSMIC host mass-gap BHs, with a significant fraction of mass-gap BH forming through AIC of an NS. Regardless of binary evolution prescriptions, the COSMIC populations show

minimal to moderate dependence on the adopted X-ray outburst selection effects and favor $M_{\text{NS, birth-max}} \lesssim 2M_{\odot}$ (to suppress the formation of AIC LMXB).

Given the consistency of the COSMIC results, we next consider a three-dimensional grid of LMXBs evolved using MESA. These results are discussed in Section 3.2. While the MESA grid samples a narrower range of binary evolution prescriptions than the COSMIC populations, MESA employs a more detailed treatment of XRB mass transfer. For $M_{\text{NS, birth-max}} = 3M_{\odot}$, MESA and COSMIC yield similar BH mass distributions, but if $M_{\text{NS, birth-max}} \lesssim 2M_{\odot}$, MESA produces a dearth of LMXBs with mass-gap BHs and reveals a greater dependence on the adopted transition to RIA in the disk instability model.

3.1. COSMIC Results

Using COSMIC, we consider proof-of-concept Milky Way populations over a range of binary evolution prescriptions. In addition to evaluating the ability of population synthesis methods to replicate the observed LMXB sample, we leverage the simulation suite to investigate dependencies between the LMXB sample and uncertain aspects of binary evolution physics. Given the substantial level of uncertainty in several binary evolution prescriptions and the order-of-magnitude nature of our study, treating the COSMIC simulations in aggregate is preferred. Here, we first outline the principal results of the simulation suite, before investigating dependencies between BH mass-gap occupancy and different prescriptions of binary evolution physics.

For a given Milky Way COSMIC population, a synthetic observed sample of LMXBs is extracted from the evolution tracks through a Monte Carlo process (Section 2.6). In Table 2, the p_{detect} -weighted fraction of LMXBs in the synthetic observed sample hosting a BH in the mass gap $f_{\text{low}}(p_{\text{detect}})$ is reported for each population. Detection probabilities are assigned under the assumption that there is no transition to RIA. For a selection of representative populations, the mass and period distributions of the synthetic observed LMXBs are presented in Figure 2. Every model predicts $f_{\text{low}}(p_{\text{detect}})$ is greater than 40% (using a benchmark $M_{\text{NS, birth-max}} = 3M_{\odot}$ and $M_{\text{bound}} = 4.5M_{\odot}$).

For LMXBs hosting a BH within the mass gap, the BHs formed through one of two channels: AIC of an NS or intrinsic formation of a BH (directly post core collapse of a star). While both channels are dependent on the maximum NS mass allowed by the equation of state (because it sets the transition between NS and BH in the compact-object birth-mass distribution), the AIC LMXBs also rely on $M_{\text{NS, birth-max}}$: if $M_{\text{NS, birth-max}}$ lies significantly below the maximum NS mass, the formation of AIC LMXBs will be suppressed relative to a population where $M_{\text{NS, birth-max}}$ lies near the maximum NS mass.

AIC LMXBs account for a significant fraction (>15% and reaching as high as 90%, for $M_{\text{NS, birth-max}} = 3M_{\odot}$ and assuming no transition to RIA) of the detection-weighted synthetic observed populations. The AIC LMXBs typically host BHs with $M_{\text{BH}} < 5M_{\odot}$. The populations where the fraction of LMXBs formed via AIC is below 30% all adopted a mass-accretion efficiency of 50%, while the populations where the fraction of LMXBs formed via AIC is above 90% all adopted the State-dependent mass-accretion prescription. The fraction of LMXBs formed via AIC is lowest for the population with

Table 2
Population Models

CE Survival	α	q Limit	Magnetic Braking	Accretion Limit	$f_{\text{low}}(M_{\text{NS, birth}} < 1.5M_{\odot})^{\text{a}}$	$f_{\text{low}}(M_{\text{NS, birth}} < 3.0M_{\odot})$
Optimistic	1.0	Lifetime-limited	IT03	State-dependent	$0.39^{+0.04}_{-0.04}$	$0.89^{+0.02}_{-0.03}$
Optimistic	1.0	Lifetime-limited	IT03	50% efficiency	$0.59^{+0.02}_{-0.04}$	$0.67^{+0.02}_{-0.04}$
Optimistic	1.0	Lifetime-limited	H02	State-dependent	$0.79^{+0.02}_{-0.02}$	$0.955^{+0.005}_{-0.006}$
Optimistic	1.0	Lifetime-limited	H02	50% efficiency	$0.58^{+0.03}_{-0.03}$	$0.72^{+0.02}_{-0.02}$
Optimistic	1.0	0.01	IT03	State-dependent	$0.37^{+0.01}_{-0.02}$	$0.80^{+0.01}_{-0.02}$
Optimistic	1.0	0.01	IT03	50% efficiency	$0.58^{+0.03}_{-0.02}$	$0.65^{+0.04}_{-0.02}$
Optimistic	1.0	0.01	H02	State-dependent	$0.603^{+0.013}_{-0.009}$	$0.880^{+0.006}_{-0.007}$
Optimistic	1.0	0.01	H02	50% efficiency	$0.56^{+0.01}_{-0.02}$	$0.67^{+0.01}_{-0.01}$
Optimistic	5.0	Lifetime-limited	IT03	State-dependent	$0.44^{+0.02}_{-0.02}$	$0.943^{+0.004}_{-0.004}$
Optimistic	5.0	Lifetime-limited	IT03	50% efficiency	$0.65^{+0.02}_{-0.03}$	$0.80^{+0.01}_{-0.01}$
Optimistic	5.0	Lifetime-limited	H02	State-dependent	$0.56^{+0.03}_{-0.03}$	$0.965^{+0.002}_{-0.005}$
Optimistic	5.0	Lifetime-limited	H02	50% efficiency	$0.60^{+0.03}_{-0.04}$	$0.84^{+0.01}_{-0.01}$
Optimistic	5.0	0.01	IT03	State-dependent	$0.39^{+0.02}_{-0.01}$	$0.893^{+0.012}_{-0.008}$
Optimistic	5.0	0.01	IT03	50% efficiency	$0.53^{+0.01}_{-0.02}$	$0.70^{+0.02}_{-0.01}$
Optimistic	5.0	0.01	H02	State-dependent	$0.43^{+0.01}_{-0.02}$	$0.914^{+0.005}_{-0.004}$
Optimistic	5.0	0.01	H02	50% efficiency	$0.50^{+0.01}_{-0.02}$	$0.741^{+0.009}_{-0.012}$
Pessimistic	1.0	Lifetime-limited	IT03	State-dependent	$0.36^{+0.04}_{-0.05}$	$0.88^{+0.01}_{-0.01}$
Pessimistic	1.0	Lifetime-limited	IT03	50% efficiency	$0.57^{+0.04}_{-0.05}$	$0.64^{+0.03}_{-0.04}$
Pessimistic	1.0	Lifetime-limited	H02	State-dependent	$0.63^{+0.03}_{-0.03}$	$0.948^{+0.005}_{-0.004}$
Pessimistic	1.0	Lifetime-limited	H02	50% efficiency	$0.56^{+0.02}_{-0.03}$	$0.71^{+0.04}_{-0.02}$
Pessimistic	1.0	0.01	IT03	State-dependent	$0.47^{+0.01}_{-0.04}$	$0.837^{+0.008}_{-0.034}$
Pessimistic	1.0	0.01	IT03	50% efficiency	$0.56^{+0.02}_{-0.02}$	$0.64^{+0.04}_{-0.02}$
Pessimistic	1.0	0.01	H02	State-dependent	$0.48^{+0.01}_{-0.01}$	$0.856^{+0.003}_{-0.003}$
Pessimistic	1.0	0.01	H02	50% efficiency	$0.597^{+0.014}_{-0.009}$	$0.69^{+0.02}_{-0.01}$
Pessimistic	5.0	Lifetime-limited	IT03	State-dependent	$0.35^{+0.06}_{-0.05}$	$0.988^{+0.002}_{-0.001}$
Pessimistic	5.0	Lifetime-limited	IT03	50% efficiency	$0.63^{+0.04}_{-0.05}$	$0.83^{+0.02}_{-0.01}$
Pessimistic	5.0	Lifetime-limited	H02	State-dependent	$0.70^{+0.02}_{-0.02}$	$0.984^{+0.001}_{-0.002}$
Pessimistic	5.0	Lifetime-limited	H02	50% efficiency	$0.59^{+0.03}_{-0.02}$	$0.930^{+0.006}_{-0.008}$
Pessimistic	5.0	0.01	IT03	State-dependent	$0.28^{+0.03}_{-0.03}$	$0.928^{+0.014}_{-0.003}$
Pessimistic	5.0	0.01	IT03	50% efficiency	$0.455^{+0.01}_{-0.0090}$	$0.643^{+0.009}_{-0.013}$
Pessimistic	5.0	0.01	H02	State-dependent	$0.351^{+0.012}_{-0.007}$	$0.911^{+0.002}_{-0.002}$
Pessimistic	5.0	0.01	H02	50% efficiency	$0.40^{+0.03}_{-0.03}$	$0.72^{+0.02}_{-0.01}$

Notes. Using COSMIC, a set of population synthesis simulations spanning uncertain aspects of binary evolution physics reveals that rapid population synthesis consistently fills the mass gap. The detection-weighted fraction of LMXBs with a BH in the mass gap $f_{\text{low}}(p_{\text{detect}})$ is presented for 32 exploratory COSMIC Milky Way populations (assuming no transition to RIA and either $M_{\text{NS, birth-max}} = 1.5M_{\odot}$ or $3.0M_{\odot}$). For a given population, the average and the 68% uncertainty are reported from 5000 snapshots of the evolution tracks between 10.5 and 11.5 Gyr, to reflect the uncertainty in the age of the Milky Way. Each population adopts a different combination of binary evolution prescriptions, and each row of the table corresponds to a different COSMIC population. As outlined in Section 2.1, the prescriptions for uncertain aspects of binary evolution physics are varied between the COSMIC populations, including the conditions for merging during a CE event (Optimistic versus Pessimistic), the efficiency of CE ejection ($\alpha = 1$ or 5), the minimum mass ratio q at ZAMS ($q > 0.01$ or Lifetime-limited, where q is restricted such that the pre-main-sequence lifetime of the secondary is shorter than the main-sequence lifetime of the primary), the magnetic braking implementation (following either H02 or IT03), the efficiency of accretion during RLO (either 50% efficiency or State-dependent efficiency, where compact objects have conservative mass accretion up to the Eddington limit). For each population, the compact-object birth-mass distribution is fixed to the Delayed prescription of Fryer et al. (2012), the maximum NS mass is fixed to $3M_{\odot}$, and the binaries' initial conditions are independently drawn from Sana et al. (2012). The data behind Table 2 are openly available from the Zenodo repository at DOI:10.5281/zenodo.8155601.

^a The fraction $f_{\text{low}}(p_{\text{detect}})$ is defined relative to $M_{\text{bound}} = 4.5M_{\odot}$.

$\alpha = 1.0$, Lifetime-limited q , Pessimistic CE survival, magnetic braking following IT03, and 50% accretion efficiency; the fraction of LMXBs formed via AIC is maximized for the population with $\alpha = 5.0$, Lifetime-limited minimum q , Pessimistic CE survival, magnetic braking following IT03, and State-dependent accretion efficiency.

AIC LMXBs also dominate (>20% and reaching as high as 99%) the subpopulation of systems with orbital periods shorter than 10 hr. Consistent with the total fraction of LMXBs formed via AIC, the fraction of AIC LMXBs among systems with periods shorter than 10 hr is reduced by adopting 50% accretion

efficiency and maximized by adopting State-dependent accretion efficiency.

We next consider how the principal results of the COSMIC simulation suite (i.e., population synthesis fails to replicate the observed LMXB population and consistently produces a population of AIC LMXBs) depends on the adopted binary evolution prescriptions, beginning with the details of the disk instability physics. From population synthesis, the detection of LMXBs hosting mass-gap BHs is largely insensitive to the adopted transition to RIA. In Figure 3, $f_{\text{low}}(p_{\text{detect}})$ is presented against $f_{\text{low-intrinsic}}$ for every combination of population

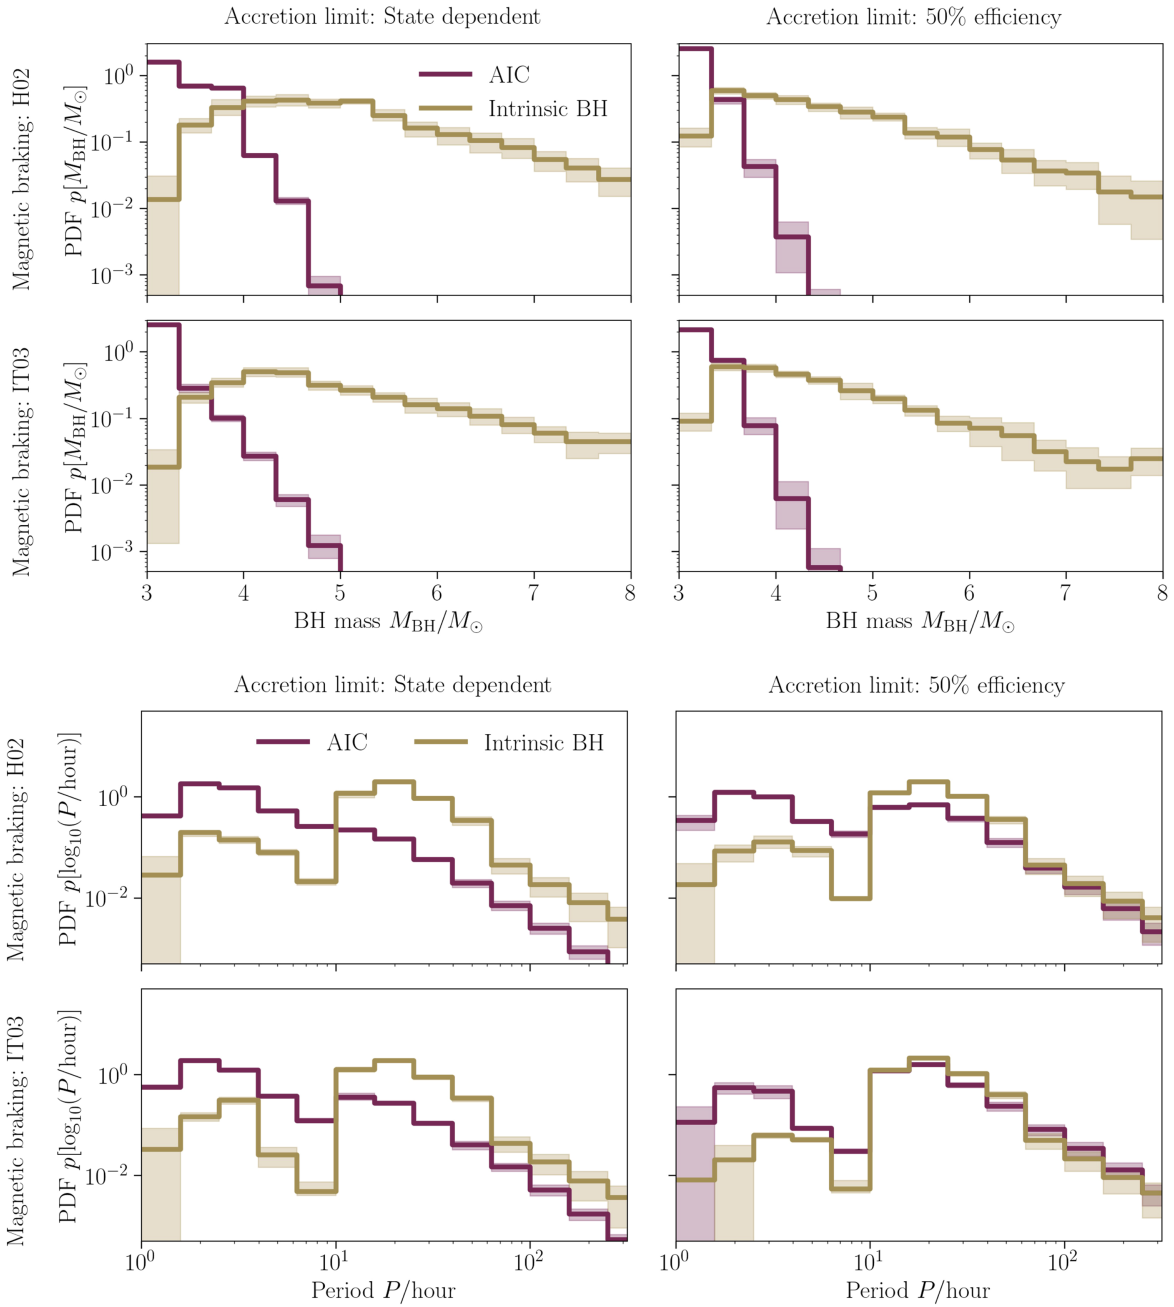


Figure 2. The BH mass (upper panels) and orbital period (lower panels) distributions for a selection of synthetic Galactic LMXB samples. Detection probabilities are calculated with no transition to RIA. Each panel corresponds to a different set of binary evolution prescriptions. For each presented population $\alpha = 5$, the minimum q at ZAMS follows the Lifetime-limited prescription, and $M_{\text{NS,birth-max}} = 3.0M_{\odot}$. The solid lines represent the average distribution across all 5000 snapshots of the evolution tracks, and the shaded regions demarcate the 16th–84th percentiles. AIC LMXBs are highlighted in purple. The detection-weighted BH mass and orbital period distributions of LMXBs from population synthesis are broadly insensitive to changes in the adopted binary evolution prescriptions.

synthesis prescription and X-ray outburst prescription. Since $f_{\text{low-intrinsic}}$ characterizes the *intrinsic* density of BHs within the lower mass gap, and $f_{\text{low}(p_{\text{detect}})}$ characterizes the *observed* density of LMXBs with a mass-gap BH, considering them together highlights the influence of transient detection effects on the observed Galactic LMXB sample.

As shown in Figure 3, modeling the X-ray outburst with no RIA transition yields $f_{\text{low}(p_{\text{detect}})}$ consistent with modeling the X-ray outburst with a smooth RIA transition. For both of these prescriptions, the detection-weighted mass distribution is only marginally different from the unweighted distribution: $f_{\text{low}(p_{\text{detect}})}$ is $\sim 10\%$ higher than $f_{\text{low-intrinsic}}$. Adopting a sharp

transition to RIA lowers $f_{\text{low}(p_{\text{detect}})}$ relative to the other RIA transitions; however, each population still yields an $f_{\text{low}(p_{\text{detect}})}$ greater than 40% (using a benchmark $M_{\text{NS,birth-max}} = 3M_{\odot}$ and $M_{\text{bound}} = 4.5M_{\odot}$). Because the outburst duration T_{outburst} is positively correlated with orbital period, and the AIC LMXBs typically have orbital periods shorter than 10 hr, the sensitivity of $f_{\text{low}(p_{\text{detect}})}$ to the RIA transition prescription is expected. The sharp transition dramatically decreases the outburst duration (Figure 1), making the short-period LMXBs, which tend to be AIC LMXBs, no longer detectable.

Discovering a transient LMXB source with an ASM relies on two independent probabilities: (i) the probability that the

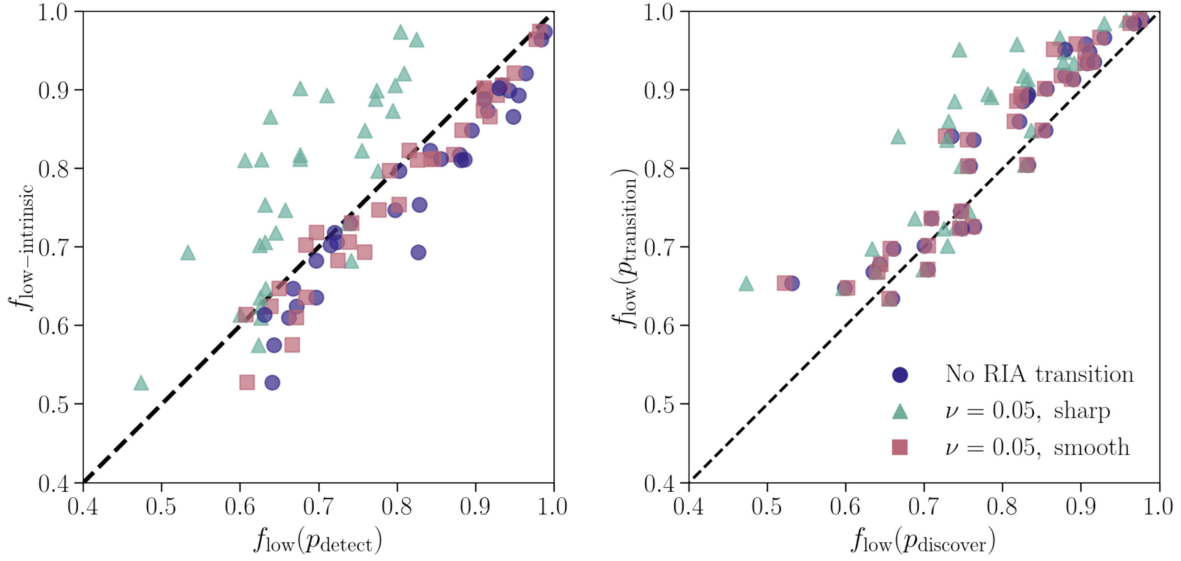


Figure 3. From population synthesis, each combination of binary evolution prescription and disk instability model yields an observed LMXB mass distribution with a substantial fraction of mass-gap BHs (for $M_{\text{NS,birth-max}} = 3M_{\odot}$). Left: the detection probability-weighted fraction of LMXBs with mass gap BHs, $f_{\text{low}}(p_{\text{detect}})$, is presented against the unweighted fraction, $f_{\text{low-intrinsic}}$, for each model (using either no transition to RIA, a sharp transition, or a smooth transition). For points that fall above the dashed black line, detection effects lead to a dearth of mass-gap BH detections relative to the intrinsic population. Right: the p_{discover} detection-weighted fraction of LMXBs with a mass-gap BH is presented against the $p_{\text{transition}}$ -weighted fraction. This panel breaks down the two components of the LMXB detection probability: (i) being bright enough during outburst for discovery p_{discover} and (ii) transitioning from outburst to quiescence during the survey $p_{\text{transition}}$. Relative to $p_{\text{transition}}$, weighting by p_{discover} suppresses observation of mass-gap BHs.

source is bright enough in outburst to be discovered (p_{discover}) and (ii) the probability that the source transitions from outburst to quiescence during the ASM’s observing window ($p_{\text{transition}}$). In Figure 3, the p_{discover} -weighted fraction of LMXB with $M_{\text{BH}} < 4.5M_{\odot}$ is presented against the $p_{\text{transition}}$ -weighted fraction. Similarly to $f_{\text{low-intrinsic}}$, these fractions are not directly comparable to the observed LMXB sample; however, they outline the relative effects of p_{discover} and $p_{\text{transition}}$ on the mass gap. As shown in Figure 3, weighting by $p_{\text{transition}}$ fills the mass gap relative to p_{discover} weighting. Compared to short-period LMXBs, long-period LMXBs are disfavored by $p_{\text{transition}}$, because the quiescence time is positively related with orbital period. Because AIC LMXBs tend to have shorter periods than LMXBs with intrinsic BHs, $p_{\text{transition}}$ favors AIC LMXBs and raises the density of BHs in the mass gap relative to p_{discover} weighting.

The prevalence of LMXBs hosting a mass-gap BH in the COSMIC populations is discrepant with the Galactic LMXB sample. While AIC LMXBs comprise a substantial fraction of the COSMIC populations for a variety of disk instability models and binary evolution prescriptions, these systems are dependent on $M_{\text{NS,birth-max}}$. Lowering $M_{\text{NS,birth-max}}$ suppresses the number of AIC LMXBs, so tuning $M_{\text{NS,birth-max}}$ could bring the population synthesis models into better agreement with observations. Below, we investigate the dependence between the AIC LMXB population and $M_{\text{NS,birth-max}}$.

In Figure 4, $f_{\text{low}}(p_{\text{detect}})$ is presented as a function of $M_{\text{NS,birth-max}}$ for a selection of COSMIC populations. While $f_{\text{low}}(p_{\text{detect}})$ decreases for each population as $M_{\text{NS,birth-max}}$ is lowered, the magnitude of the effect depends strongly on the accretion-limit prescription. This dependence reflects the higher occurrence rate of AIC LMXBs in populations with State-dependent accretion limits. For the State-dependent prescription, compact objects have conservative mass transfer, provided the mass-transfer rate is sub-Eddington; accretors therefore retain a greater fraction of the transferred mass, which

corresponds to more NSs undergoing AIC. The increased accretion rates from the State-dependent prescription also translate to greater mass gain for the intrinsic BH LMXBs, which in turn reduces $f_{\text{low}}(p_{\text{detect}})$. Populations with 50% accretion efficiency have both higher $f_{\text{low}}(p_{\text{detect}})$ for intrinsic BHs and weaker dependence on $M_{\text{NS,birth-max}}$.

Over the range of uncertain binary evolution physics that we considered, LMXB populations modeled with COSMIC yield a substantial fraction of LMXB BHs with $M_{\text{BH}} < 4.5M_{\odot}$ (over 40%, assuming no transition to RIA) for $M_{\text{NS,birth-max}} = 3M_{\odot}$. The simulations are in tension with the observed sample of LMXBs, where there is a dearth of BHs with masses below $\approx 4.5M_{\odot}$ (Bailyn et al. 1998; Özel et al. 2010; Farr et al. 2011). The fraction of mass-gap BHs in the synthetic Galactic sample is minimized by adopting a CE efficiency of $\alpha = 5$, $q > 0.01$, magnetic braking following IT03, State-dependent accretion efficiency, pessimistic CE survival, no transition to RIA, and $M_{\text{NS,birth-max}} = 1.5M_{\odot}$; for this population, 28% of LMXBs host a BH in the mass gap. While lowering $M_{\text{NS,birth-max}}$ lowers the density of LMXB systems in the mass gap, $f_{\text{low}}(p_{\text{detect}})$ remains greater than 25% for every COSMIC population and is substantially higher for most, assuming no transition to RIA.

3.2. MESA Results

Our proof-of-concept rapid population synthesis Milky Way populations consistently populate the lower mass gap. These simulations also reveal a dependence between $M_{\text{NS,birth-max}}$ and the existence of the mass gap. Informed by the COSMIC simulations, here we investigate a grid of binary MESA models. We adopt a three-dimensional grid of binaries (in compact-object mass, donor mass, and orbital period space), consisting of a compact object and a hydrogen-rich main-sequence star from Fragos et al. (2023). A series of COSMIC populations are initialized with the same initial conditions as the MESA grid, to facilitate comparison.

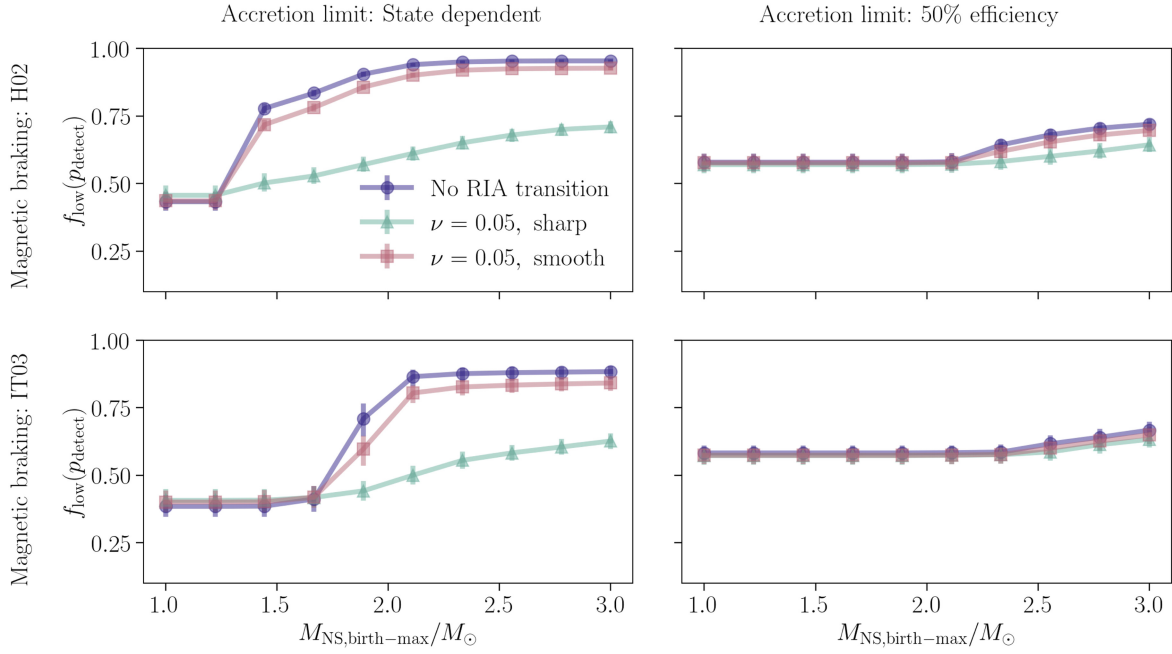


Figure 4. The detection-weighted fraction of LMXBs with a mass-gap BH $f_{\text{low}}(p_{\text{detect}})$ as a function of $M_{\text{NS, birth-max}}$ for a selection of populations. Fractions for each population are calculated by averaging over all 5000 snapshots, and the uncertainties are taken from the 68% uncertainty across the snapshots. Lowering the maximum NS birth mass $M_{\text{NS, birth-max}}$ reduces the occurrence of LMXB systems with $M_{\text{BH}} < 4.5M_{\odot}$. The magnitude of this trend depends on the adopted accretion prescription; changing the magnetic braking prescription moderately affects the fraction of mass-gap BH detections, while modifying the adopted accretion physics significantly affects $f_{\text{low}}(p_{\text{detect}})$.

Table 3
COSMIC and MESA Grid Comparison

Magnetic Braking		Accretion Limit	RIA Transition	$f_{\text{low}}(M_{\text{NS, birth}} < 1.5M_{\odot})^a$	$f_{\text{low}}(M_{\text{NS, birth}} < 3.0M_{\odot})$
COSMIC	H02	50% efficiency	No transition	0.345	0.406
	H02	50% efficiency	$\nu = 0.05$, sharp	0.422	0.492
	H02	50% efficiency	$\nu = 0.05$, smooth	0.347	0.407
	IT03	50% efficiency	No transition	0.345	0.406
	IT03	50% efficiency	$\nu = 0.05$, sharp	0.418	0.490
	IT03	50% efficiency	$\nu = 0.05$, smooth	0.344	0.405
	H02	State-dependent	No transition	0.354	0.437
	H02	State-dependent	$\nu = 0.05$, sharp	0.439	0.516
	H02	State-dependent	$\nu = 0.05$, smooth	0.351	0.435
	IT03	State-dependent	No transition	0.352	0.435
	IT03	State-dependent	$\nu = 0.05$, sharp	0.439	0.515
	IT03	State-dependent	$\nu = 0.05$, smooth	0.353	0.437
MESA			No transition	0.105	0.428
			$\nu = 0.05$, sharp	0.242	0.507
			$\nu = 0.05$, smooth	0.105	0.429

Notes. Relative to rapid population synthesis, LMXBs evolved using MESA underpopulate the mass gap if $M_{\text{NS, birth-max}} \lesssim 2M_{\odot}$, and show greater dependence on the adopted RIA transition. The detection-weighted fraction of LMXBs with a BH in the mass gap $f_{\text{low}}(p_{\text{detect}})$ is presented for a three-dimensional grid of binaries evolved with MESA, as well as for grids evolved using COSMIC; the grids evolved with COSMIC are initialized at the same initial conditions as the grid evolved using MESA. Values of the detection-weighted fraction $f_{\text{low}}(p_{\text{detect}})$ are reported for each combination of $M_{\text{NS, birth-max}} = 1.5M_{\odot}$ or $3.0M_{\odot}$ and smooth, sharp, or nonexistent transition to RIA. To marginalize over uncertain aspects of binary evolution physics, different combinations of binary evolution prescriptions are adopted for the COSMIC grids, including the magnetic braking implementation (following either H02 or IT03) and the efficiency of accretion during RLO (either 50% efficiency or State-dependent efficiency, where compact objects have conservative mass accretion up to the Eddington limit).

^a The fraction $f_{\text{low}}(p_{\text{detect}})$ is defined relative to $M_{\text{bound}} = 4.5M_{\odot}$.

Using the Monte Carlo sampling procedure described in Section 2.6, $f_{\text{low}}(p_{\text{detect}})$ is reported for the MESA grid and each COSMIC grid in Table 3, for a range of accretion disk treatments and $M_{\text{NS, birth-max}}$. For $M_{\text{NS, birth-max}} = 3M_{\odot}$, the MESA and COSMIC grids yield generally consistent $f_{\text{low}}(p_{\text{detect}})$. However, when $M_{\text{NS, birth-max}}$ is lowered, MESA predicts a consistently lower $f_{\text{low}}(p_{\text{detect}})$ than COSMIC, and the

dependence on the assumed transition to RIA is exacerbated. When $M_{\text{NS, birth-max}} \lesssim 2M_{\odot}$, $f_{\text{low}}(p_{\text{detect}})$ for an LMXB population modeled with MESA is approximately half that of an LMXB population modeled with COSMIC.

In Figure 5, $f_{\text{low}}(p_{\text{detect}})$ is presented as a function of $M_{\text{NS, birth-max}}$ for the MESA grid and a representative COSMIC grid; for a more comprehensive view, the average number of




Figure 5. Top: the detection-weighted fraction of LMXBs with a mass-gap BH $f_{\text{low}}(p_{\text{detect}})$ as a function of $M_{\text{NS,birth-max}}$. Each color corresponds to a different RIA transition. For each panel, the results for no transition to RIA overlap with the results for a smooth transition to RIA. Bottom: the average number of detections of AIC LMXBs (over $N_{\text{trials}} = 5000$ snapshots of the evolutionary tracks) with $M_{\text{BH}} < 4.5 M_{\odot}$ (as a solid line); the dashed horizontal line is the average number of detections of LMXB that formed with a BH and have $M_{\text{BH}} < 4.5 M_{\odot}$; the dotted horizontal line is the average number of detections of LMXB (from any formation channel) with $M_{\text{BH}} > 4.5 M_{\odot}$. The left column presents the results from the MESA grid, while the right column presents the results from the COSMIC grid. While MESA and COSMIC predict similar LMXB BH mass distributions when $M_{\text{NS,birth-max}} = 3 M_{\odot}$, MESA produces a dearth of LMXBs with mass gap BHs relative to COSMIC if $M_{\text{NS,birth-max}} \lesssim 2 M_{\odot}$.

LMXB detections (separated by formation channel and BH mass) is presented as well. For $M_{\text{NS,birth-max}} \lesssim 2 M_{\odot}$, the MESA grid yields $f_{\text{low}}(p_{\text{detect}})$ a factor of two lower than the COSMIC grids. This disparity is enhanced when either a smooth or nonexistent transition to RIA is adopted. The sharp RIA transition actually reduces the number of detected AIC LMXBs for the MESA grid. However, the behavior of $f_{\text{low}}(p_{\text{detect}})$ is dominated by the increased detection of LMXBs with $M_{\text{BH}} > 4.5 M_{\odot}$ when either a smooth or nonexistent transition to RIA is adopted for the MESA grid and the less dramatic increase for the COSMIC grid (i.e., a smooth or nonexistent transition to RIA leads to slightly more mass-gap BH detections but a far greater increase in the number of higher-mass LMXB detections).

Similarly to the BH mass distributions, the orbital period distributions from MESA and COSMIC also diverge. In Figure 6, orbital period distributions are presented for three subpopulations of LMXBs (AIC LMXBs, LMXBs with intrinsic BH in the mass gap, and LMXBs with intrinsic BH above the mass gap) from the MESA grid and each COSMIC grid. To highlight the influence of p_{detect} , Figure 6 includes both the underlying orbital period distributions (i.e., only Monte Carlo sampled) and the synthetic observed orbital period distributions (i.e., Monte Carlo sampled and detection weighted).

While the observed orbital period distributions of the COSMIC grids are generally consistent with MESA, the underlying distributions show significant differences. Relative to the COSMIC grids, the underlying orbital period distribution of AIC LMXBs is skewed toward shorter orbital periods in the

MESA grid; without p_{detect} weighting, 80% of AIC LMXBs in the MESA grid have $P < 10$ hr, but $< 35\%$ of AIC LMXBs in the COSMIC grids have $P < 10$ hr. For LMXBs with intrinsic BHs and $P < 10$ hr, the density of LMXBs with mass-gap BHs is lower in the MESA grid, relative to the COSMIC grids; in the MESA grid, $f_{\text{low-intrinsic}} = 7\%$ for LMXBs with intrinsic BHs and $P < 10$ hr, while the fraction is $\sim 30\%$ in the COSMIC grids. For $P > 10$ hr, this trend continues, but the magnitude is lessened; in the MESA grid, $f_{\text{low-intrinsic}} = 15\%$ for LMXBs with intrinsic BH and $P > 10$ hr, while the fraction is $\sim 20\%$ for the COSMIC grids.

The differences in the orbital period and mass distributions between COSMIC and MESA are both consequences of how RLO mass transfer is treated in each code. COSMIC treats RLO mass transfer following Hurley et al. (2002), in which stars are (nonphysically) approximated as in thermal equilibrium throughout mass loss. For main-sequence donors in the MESA models, mass is removed for stars overflowing their Roche lobes at the beginning of each time step, such that the star's radius remains within its Roche lobe, while for giant stars, the methods of Kolb & Ritter (1990) are adopted; stars' radii are allowed to extend beyond their Roche lobes, and mass transfer is self-consistently calculated based on the local fluid conditions (Fragos et al. 2023). Differences in mass-transfer rate naturally affect an LMXB's BH mass and orbital evolution.

Within the first 100 Myr of evolution, every compact object in the grid accretes more mass when MESA is used than when COSMIC is used. Similarly, every donor star in the grid loses more mass when MESA is used than when COSMIC is used. Of

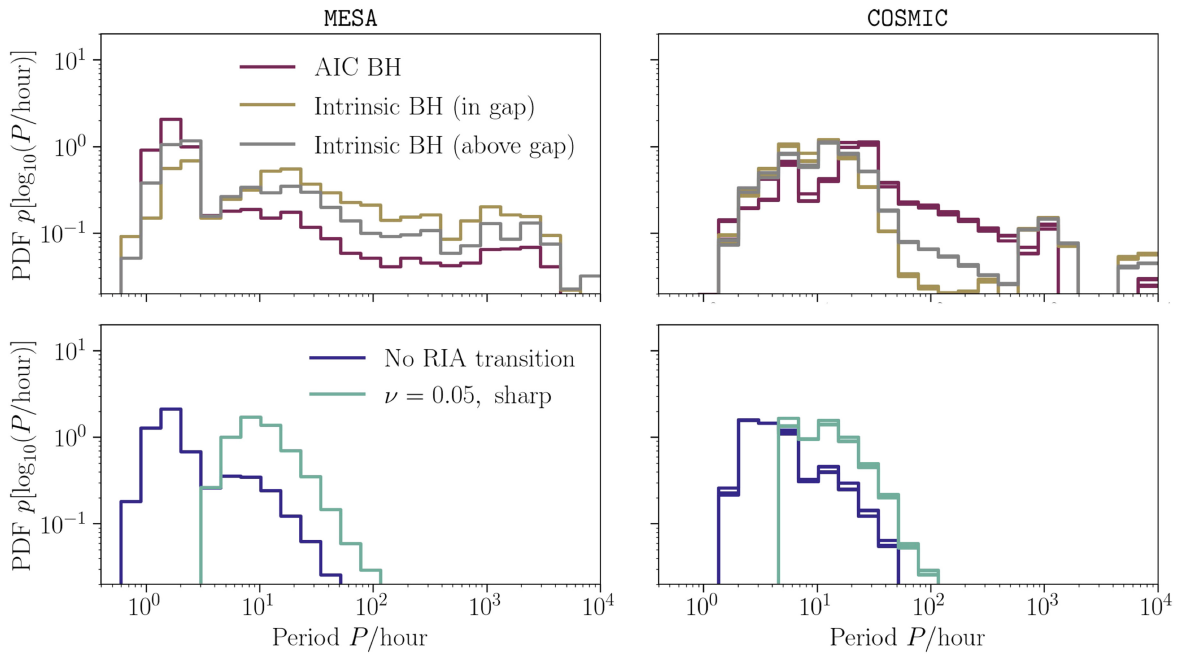


Figure 6. The LMXB orbital period distribution for the MESA grid is presented in the left column, while the orbital period distributions of the COSMIC grids (which are generated with matching initial conditions) are presented on the right. Top: Monte Carlo sampled orbital period distributions are presented for three subpopulations of LMXBs: AIC LMXBs (purple), LMXBs with intrinsic BHs in the mass gap (gold), and LMXBs with intrinsic BHs above the mass gap (gray). Bottom: Monte Carlo sampled and detection-weighted orbital period distributions for the entire population of LMXBs are presented for either no transition to RIA (blue) or a sharp transition (teal); here, applying a smooth transition to RIA is nearly indistinguishable from applying no transition. LMXBs evolved using MESA typically occupy shorter orbital periods than LMXBs evolved with COSMIC.

LMXBs in the grid that do not merge by 100 Myr, the compact objects of LMXBs evolved using MESA accrete on average $0.7M_{\odot}$ more mass by 100 Myr than compact objects in LMXBs evolved using COSMIC; weighting by detection probability (assuming no transition to RIA), the compact objects of LMXBs evolved using MESA accrete on average $0.4M_{\odot}$ more mass by 100 Myr than the compact objects of LMXBs evolved using COSMIC. The greater mass accretion by compact objects in MESA results in both more massive compact objects and shorter orbital periods.

To highlight these differences, the evolution tracks of a single LMXB evolved using MESA are presented against the evolution tracks from COSMIC (for an LMXB with the same initial conditions) in Figure 7. The majority of both the orbital period evolution and mass gain for the compact object occurs in the first few hundred million years in the MESA grid. During this period, the mass-loss rate from the donor star is larger in MESA than in COSMIC by at least two orders of magnitude. Because the donor mass-loss rate is still sub-Eddington, the mass exchange is conservative for both codes. In contrast, the timescales for orbital period evolution and mass gain in the COSMIC grid extend to several billion years. This leads to more dramatic orbital period shrinking and more accretion onto the compact object, and thus results in the overall higher-mass BHs and shorter periods in the MESA grid when compared to COSMIC.

In summary, a COSMIC population study of uncertain aspects of binary evolution and disk instability models consistently found a substantial fraction of LMXBs hosting BHs with $M_{\text{BH}} < 4.5M_{\odot}$ (over 40%) for $M_{\text{NS, birth-max}} = 3M_{\odot}$. After lowering $M_{\text{NS, birth-max}}$ to $1.5M_{\odot}$, $f_{\text{low}}(p_{\text{detect}}) > 25\%$ for every COSMIC population. For $M_{\text{NS, birth-max}} = 3M_{\odot}$, MESA and COSMIC yield similar BH mass distributions and both fill

the lower mass gap. However, if $M_{\text{NS, birth-max}}$ is lowered to below $2M_{\odot}$, the LMXBs produced with MESA have at least a factor of two fewer BHs with $M_{\text{BH}} < 4.5M_{\odot}$. This difference is exacerbated when a smooth or nonexistent transition to RIA is adopted. These results highlight the importance of accurately modeling mass transfer in understanding the formation of compact objects.

4. Discussion

While Galactic LMXBs have previously been used to constrain the underlying stellar-mass BH distribution (Bailyn et al. 1998; Özel et al. 2010; Farr et al. 2011), the impact of detection biases on the observed LMXB sample has remained unclear (Özel et al. 2010; Kreidberg et al. 2012; Jonker et al. 2021). Here, we presented a quantitative study of selection biases, motivated by the fact that BH mass measurement is possible only for *transient* XRB systems. We examined whether the requirement for transient behavior can hide an underlying population of low-mass BHs and create an *apparent* mass gap in the observed sample.

We followed a forward-modeling approach. Using both rapid binary population synthesis (COSMIC) and detailed stellar evolution models (MESA), we generated initial synthetic binary systems, evolved the systems forward in time, and inferred the detectable populations using the disk instability model of LMXB transient behavior. To reflect the criteria for LMXB mass inference, synthetic samples were weighted by the probability an RXTE-like ASM (Levine et al. 1996) would detect the source as transient. From these results, we demonstrated that transient LMXB selection effects do introduce some biases into the observed sample, but this is not sufficient to completely hide mass-gap BHs. Additionally,




Figure 7. The differences in the orbital period and mass distributions between COSMIC and MESA both originate from the treatment of RLO mass accretion. Relative to COSMIC, evolving LMXBs with MESA leads to greater mass accretion by the BHs, and by extension, shorter orbital periods. Here, the evolution tracks of a single LMXB evolved with MESA (dashed lines) are presented against the evolution tracks when COSMIC is used (solid lines). The binary is initialized with $M_{\text{BH}} = 2.5M_{\odot}$, $M_{\text{donor}} = 1.16M_{\odot}$, and $P = 0.413$ days. For the COSMIC evolution tracks, the magnetic braking formulation of H02 and the State-dependent accretion efficiency were adopted.

observations of the lower mass gap have implications for $M_{\text{NS,birth-max}}$, as this impacts the number of BHs (or lack thereof) formed through AIC. We conclude that, while selection effects are important in understanding XRB observations, the lack of LMXB BH seen in the lower mass gap places constraints upon viable models of binary stellar evolution.

Using COSMIC, we generated Milky Way binary populations for an array of binary evolution prescriptions and initial conditions. While current rapid population synthesis methods likely cannot reproduce all properties of the observed LMXB sample, they enable exploratory studies to uncover any significant dependencies between the synthetic observed LMXB sample and the binary evolution prescriptions. In concert with the suite of population synthesis simulations, we considered a grid of MESA binary models. We adopted a three-dimensional grid of binaries consisting of a compact object and a hydrogen-rich main-sequence star from Fragos et al. (2023). While the MESA grid does not consider the full range of binary evolution prescriptions investigated with COSMIC, here we explored an improved physical treatment of XRB mass transfer. Our key results are:

1. Our rapid population synthesis results cannot replicate the observed mass distribution of Galactic LMXBs. The COSMIC results indicate that detection of LMXBs containing mass-gap BHs is a robust prediction of the Delayed explosion mechanism. For 32 combinations of binary evolution prescriptions and three disk instability models (smooth, sharp, or nonexistent transition to RIA) over 40% of the LMXBs in the synthetic Galactic samples have $M_{\text{BH}} < 4.5M_{\odot}$ (with $M_{\text{NS,birth-max}} = 3M_{\odot}$).
2. LMXBs with a mass-gap BH form through one of two channels: AIC of an NS or directly post core collapse of a star. Unlike LMXBs with intrinsic BHs, the formation of

AIC LMXBs relies upon $M_{\text{NS,birth-max}}$ (i.e., lowering $M_{\text{NS,birth-max}}$ suppresses the formation of AIC LMXBs). Because $M_{\text{BH}} \lesssim 5M_{\odot}$ for the AIC LMXBs, the existence of a lower mass gap in the Galactic LMXB sample can potentially constrain $M_{\text{NS,birth-max}}$.

3. While rapid population synthesis and MESA models are consistent for high values of the maximum NS birth mass (e.g., $M_{\text{NS,birth-max}} = 3M_{\odot}$), LMXBs evolved with MESA have a lower fraction of BHs in the mass gap relative to the COSMIC models, if $M_{\text{NS,birth-max}} \lesssim 2M_{\odot}$. Among the synthetic LMXB samples, the MESA models produce a factor of two fewer LMXBs with a mass-gap BH than the analogous COSMIC systems. This trend is exacerbated when a smooth or nonexistent transition to RIA is adopted.
4. Relative to our rapid population synthesis models, BHs in LMXBs evolved using MESA generally undergo greater mass accretion. This results in both more massive BHs and shorter-period LMXBs in the MESA models.

Drawing population-level conclusions from the MESA grid alone presents considerable challenges. Unlike the COSMIC Milky Way models, the initial conditions for the MESA grid are not drawn from informed distributions, and the MESA grid only considers the compact-object–main-sequence phase of evolution. While we find that the COSMIC simulated LMXB samples host significantly more BHs in the mass gap than samples evolved using MESA, we cannot predict the occurrence rate of mass-gap BH for the Milky Way using the MESA grid.

To suppress the formation of AIC LMXBs, both the COSMIC Milky Way populations and the MESA grid favor $M_{\text{NS,birth-max}} \lesssim 2M_{\odot}$. This constraint on $M_{\text{NS,birth-max}}$ is consistent with the observed NS mass distribution. Masses have been inferred for >70 NSs via pulsar timing, eclipsing

X-ray binaries, and optical radial velocity measurements of an NS's companion (Alsing et al. 2018; Lattimer 2019). From these measurements, the NS mass distribution displays a strong peak near $1.4M_{\odot}$ with a few objects heavier than $\sim 1.6M_{\odot}$ and a maximum around $\sim 2.3M_{\odot}$ (Shao et al. 2020). However, the most massive NSs are in binaries with low-mass companions (typically white dwarfs) and likely underwent accretion growth. With the exceptions of GW190814 and GW200210_092254, which are consistent with containing a compact object lying in the mass gap (Abbott et al. 2020, 2021a, 2021b), there are few gravitational-wave sources with significant support for a mass-gap component, and observations are consistent with a dip in the NS population above $\sim 2.1M_{\odot}$ (Abbott et al. 2023). The preference for $M_{\text{NS,birth-max}} \lesssim 2M_{\odot}$ in both the COSMIC and MESA simulations is approximately consistent with these observed samples.

Over a wide range of binary evolution prescriptions, $M_{\text{NS,birth-max}}$, and disk instability models, COSMIC Milky Way populations fill the mass gap. In isolation, the robustness of this result (and its contradiction with observation) would favor a supernova explosion mechanism that inherently creates a mass gap (e.g., the Rapid prescription of Fryer et al. 2012). However, gravitational-wave observations still favor mechanisms that (at least partially) populate the gap (e.g., the Delayed prescription of Fryer et al. 2012). Therefore, there is still physics regarding the formation of mass-gap BHs (either from AIC of an NS or core collapse) that remains to be understood. Given the overproduction of mass-gap BHs with COSMIC relative to MESA, greater study of LMXB formation and observably using detailed stellar evolution models (e.g., POSYDON; Fragos et al. 2023) including up-to-date supernova and stellar-collapse prescriptions (e.g., Ertl et al. 2020; Fryer et al. 2022) is vital to untangling the mystery of the lower mass gap. Additionally, the presence of observational selection biases should motivate the building of more sensitive ASMs than RXTE, such as those proposed for missions like eXTP and STROBE-X (Zhang et al. 2016; Ray et al. 2019), as well as the use of sufficient optical and infrared resources to make mass estimates for the fainter X-ray transients, which will tend to have fainter optical counterparts.

5. Conclusions

For a supernova engine that fills the lower BH mass gap, transient LMXB selection effects do introduce significant biases into the observed LMXB sample. However, unless there are further (unaccounted for) observational biases against finding LMXBs with mass-gap BHs, population synthesis models fail to reproduce this aspect of the observed LMXB population. This result is robust against variations of uncertain aspects of binary evolution physics, e.g., CE efficiency, CE survivability, minimum ZAMS mass ratio, magnetic braking, and accretion efficiency. This points to the need for additional physics not currently included in our COSMIC and MESA simulations, such as a supernova mechanism that suppresses formation of mass-gap objects.

Regardless of whether the low-mass BHs form from core collapse, the results of our COSMIC and MESA models lead to the robust implication that the NS birth masses must be suppressed above $\sim 2M_{\odot}$. Otherwise, the mass gap would be filled by low-mass BHs formed through NS AIC. This constraint on the maximum NS birth mass is independent of

whether the supernova engine forms mass-gap BHs. This result alone motivates a re-examination of the physics included in the COSMIC and MESA simulations, such as a supernova engine that limits the maximum birth mass of NS.

Acknowledgments

The authors thank Michael Zevin and Simon Stevenson for useful discussions, and the referee for useful suggestions. J.S. and I.K. were supported as CIERA REU students by the National Science Foundation (NSF) under grant No. 1757792. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the NSF. J.S. was also partially supported by CIFAR (through V.K.'s Senior Fellowship). V.K. was partially supported through a CIFAR Senior Fellowship, a Guggenheim Fellowship, and the Gordon and Betty Moore Foundation (grant award GBMF8477). C.P.L.B. was supported by the CIERA Board of Visitors Research Professorship. J.J.A. was supported by Northwestern University through a CIERA Postdoctoral Fellowship. K.R. was supported by the Gordon and Betty Moore Foundation (PI Kalogera, grant award GBMF8477) and the Riedel Family Graduate Fellowship in CIERA. A.D., P.S., and M.S. were supported by the Gordon and Betty Moore Foundation (grant award GBMF8477). S.B. B., T.F., K.K., D.M., and Z.X. were supported by a Swiss National Science Foundation Professorship grant (PP00P2_176868; PI Fragos). K.K. acknowledges support from the Federal Commission for Scholarships for Foreign Students for the Swiss Government Excellence Scholarship (ESKAS No. 2021.0277). Z.X. acknowledges support from the Chinese Scholarship Council (CSC). E.Z. acknowledges funding support from the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation program (grant agreement No. 772086). The Flatiron Institute is funded by the Simons Foundation. This research was supported in part through the computational resources and staff contributions provided for the Quest High-Performance Computing Cluster at Northwestern University, which is jointly supported by the Office of the Provost, the Office for Research, and Northwestern University Information Technology.

Data supporting the findings reported in this paper are openly available from the Zenodo repository at DOI:10.5281/zenodo.8155601.

Software: astropy (Robitaille et al. 2013; Price-Whelan et al. 2018); COSMIC (Breivik et al. 2020); matplotlib (Hunter 2007); MESA (Paxton et al. 2011, 2013, 2015, 2018, 2019); numpy (van der Walt et al. 2011); pandas (McKinney 2010; The pandas development team 2020); POSYDON (Fragos et al. 2023); scipy (Virtanen et al. 2020).

ORCID iDs

Jared C. Siegel  <https://orcid.org/0000-0002-9337-0902>


Vicky Kalogera  <https://orcid.org/0000-0001-9236-5469>

Christopher P. L. Berry  <https://orcid.org/0000-0003-3870-7215>

Katelyn Breivik  <https://orcid.org/0000-0001-5228-6598>

Jeff J. Andrews  <https://orcid.org/0000-0001-5261-3923>

Simone S. Bavera  <https://orcid.org/0000-0002-3439-0321>

Aaron Dotter  <https://orcid.org/0000-0002-4442-5700>

Tassos Fragos  <https://orcid.org/0000-0003-1474-1523>
 Konstantinos Kovelakas  <https://orcid.org/0000-0003-3684-964X>
 Devina Misra  <https://orcid.org/0000-0003-4260-960X>
 Kyle A. Rocha  <https://orcid.org/0000-0003-4474-6528>
 Meng Sun  <https://orcid.org/0000-0001-9037-6180>
 Zepei Xing  <https://orcid.org/0000-0002-0031-3029>

References

- Abbott, R., Abbott, T. D., Abraham, S., et al. 2020, *ApJL*, **896**, L44
 Abbott, R., Abbott, T. D., Acernese, F., et al. 2023, *PhRvX*, **13**, 011048
 Abbott, R., Abbott, T. D., Acernese, F., et al. 2021a, arXiv:2111.03606
 Abbott, R., Abbott, T. D., Acernese, F., et al. 2021b, arXiv:2108.01045
 Ai, S., Gao, H., & Zhang, B. 2020, *ApJ*, **893**, 146
 Alsing, J., Silva, H. O., & Berti, E. 2018, *MNRAS*, **478**, 1377
 Atri, P., Miller-Jones, J. C. A., Bahramian, A., et al. 2019, *MNRAS*, **489**, 3116
 Bailyn, C. D., Jain, R. K., Coppi, P., & Orosz, J. A. 1998, *ApJ*, **499**, 367
 Belczynski, K., Kalogera, V., Rasio, F. A., et al. 2008, *ApJS*, **174**, 223
 Belczynski, K., Wiktorowicz, G., Fryer, C. L., Holz, D. E., & Kalogera, V. 2012, *ApJ*, **757**, 91
 Bolton, C. T. 1972, *Natur*, **235**, 271
 Bradt, H. V., Rothschild, R. E., & Swank, J. H. 1993, *A&AS*, **97**, 355
 Breivik, K., Coughlin, S., Zevin, M., et al. 2020, *ApJ*, **898**, 71
 Cannizzo, J. K., Ghosh, P., & Wheeler, J. C. 1982, *ApJL*, **260**, L83
 Cannizzo, J. K., Shafter, A. W., & Wheeler, J. C. 1988, *ApJ*, **333**, 227
 Corral-Santana, J. M., Casares, J., Muñoz-Darias, T., et al. 2016, *A&A*, **587**, A61
 de Kool, M. 1990, *ApJ*, **358**, 189
 Degenaar, N., Ballantyne, D. R., Belloni, T., et al. 2018, *SSRv*, **214**, 15
 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
 El-Badry, K., Seeburger, R., Jayasinghe, T., et al. 2022, *MNRAS*, **512**, 5620
 Ertl, T., Woosley, S. E., Sukhbold, T., & Janka, H. T. 2020, *ApJ*, **890**, 51
 Farr, W. M., Sravan, N., Cantrell, A., et al. 2011, *ApJ*, **741**, 103
 Fragos, T., Andrews, J. J., Bavera, S. S., et al. 2023, *ApJS*, **264**, 45
 Fragos, T., Andrews, J. J., Ramirez-Ruiz, E., et al. 2019, *ApJL*, **883**, L45
 Frank, J., King, A., & Raine, D. 1992, *Accretion Power in Astrophysics*, Vol. 21 (Cambridge: Cambridge Univ. Press)
 Fryer, C. L., Belczynski, K., Wiktorowicz, G., et al. 2012, *ApJ*, **749**, 91
 Fryer, C. L., Olejak, A., & Belczynski, K. 2022, *ApJ*, **931**, 94
 Gallegos-Garcia, M., Berry, C. P. L., Marchant, P., & Kalogera, V. 2021, *ApJ*, **922**, 110
 Giacobbo, N., & Mapelli, M. 2018, *MNRAS*, **480**, 2011
 Giacobbo, N., & Mapelli, M. 2019, *MNRAS*, **482**, 2234
 Hobbs, G., Lorimer, D. R., Lyne, A. G., & Kramer, M. 2005, *MNRAS*, **360**, 974
 Hunter, J. D. 2007, *CSE*, **9**, 90
 Hurley, J. R., Tout, C. A., & Pols, O. R. 2002, *MNRAS*, **329**, 897
 Ivanova, N. 2006, *ApJL*, **653**, L137
 Ivanova, N., & Taam, R. E. 2003, *ApJ*, **599**, 516
 Jayasinghe, T., Stanek, K. Z., Thompson, T. A., et al. 2021, *MNRAS*, **504**, 2577
 Jonker, P. G., Kaur, K., Stone, N., & Torres, M. A. P. 2021, *ApJ*, **921**, 131
 Kalogera, V., & Baym, G. 1996, *ApJL*, **470**, L61
 King, A. R., Kolb, U., & Burderi, L. 1996, *ApJL*, **464**, L127
 King, A. R., Kolb, U., & Szuszkiewicz, E. 1997, *ApJ*, **488**, 89
 King, A. R., & Ritter, H. 1998, *MNRAS*, **293**, L42
 Kneivitt, G., Wynn, G. A., Vaughan, S., & Watson, M. G. 2014, *MNRAS*, **437**, 3087
 Knigge, C., Baraffe, I., & Patterson, J. 2011, *ApJS*, **194**, 28
 Kolb, U., & Ritter, H. 1990, *A&A*, **236**, 385
 Kreidberg, L., Bailyn, C. D., Farr, W. M., & Kalogera, V. 2012, *ApJ*, **757**, 36
 Lasota, J.-P. 2001, *NewAR*, **45**, 449
 Lasota, J. P., Dubus, G., & Kruk, K. 2008, *A&A*, **486**, 523
 Lasota, J.-P. 2016, in *Astrophysics of Black Holes*, Astrophysics and Space Science Library, ed. C. Bambi, Vol. 440 (Berlin: Springer), 1
 Lattimer, J. M. 2019, *Univ*, **5**, 159
 Levine, A. M., Bradt, H., Cui, W., et al. 1996, *ApJL*, **469**, L33
 Liu, T., Wei, Y.-F., Xue, L., & Sun, M.-Y. 2021, *ApJ*, **908**, 106
 Mandel, I., & Müller, B. 2020, *MNRAS*, **499**, 3214
 Margalit, B., & Metzger, B. D. 2017, *ApJL*, **850**, L19
 McKinney, W. 2010, in *Proc. of the 9th Python in Sci. Conf.*, ed. S. van der Walt & J. Millman (Austin, TX: Scipy) 56
 McMillan, P. J. 2011, *MNRAS*, **414**, 2446
 Müller, H., & Serot, B. D. 1996, *NuPhA*, **606**, 508
 Özel, F., & Freire, P. 2016, *ARA&A*, **54**, 401
 Özel, F., Psaltis, D., Narayan, R., & McClintock, J. E. 2010, *ApJ*, **725**, 1918
 Patton, R. A., Sukhbold, T., & Eldridge, J. J. 2022, *MNRAS*, **511**, 903
 Paxton, B., Bildsten, L., Dotter, A., et al. 2011, *ApJS*, **192**, 3
 Paxton, B., Cantiello, M., Arras, P., et al. 2013, *ApJS*, **208**, 4
 Paxton, B., Marchant, P., Schwab, J., et al. 2015, *ApJS*, **220**, 15
 Paxton, B., Schwab, J., Bauer, E. B., et al. 2018, *ApJS*, **234**, 34
 Paxton, B., Smolec, R., Schwab, J., et al. 2019, *ApJS*, **243**, 10
 Price-Whelan, A. M., Sipőcz, B. M., Günther, H. M., et al. 2018, *AJ*, **156**, 123
 Raaijmakers, G., Greif, S. K., Hebel, K., et al. 2021, *ApJL*, **918**, L29
 Ray, P., Arzoumanian, Z., Ballantyne, D., et al. 2019, *BAAS*, **51**, 231
 Remillard, R. A., & McClintock, J. E. 2006, *ARA&A*, **44**, 49
 Rhoades, C. E., & Ruffini, R. 1974, *PhRvL*, **32**, 324
 Robin, A. C., Reylé, C., Derrière, S., & Picaud, S. 2003, *A&A*, **409**, 523
 Robitaille, T. P., Tollerud, E. J., Greenfield, P., et al. 2013, *A&A*, **558**, A33
 Sana, H., de Mink, S. E., de Koter, A., et al. 2012, *Sci*, **337**, 444
 Santoliquido, F., Mapelli, M., Giacobbo, N., Bouffanais, Y., & Artale, M. C. 2021, *MNRAS*, **502**, 4877
 Shao, D.-S., Tang, S.-P., Jiang, J.-L., & Fan, Y.-Z. 2020, *PhRvD*, **102**, 063006
 Shenar, T., Sana, H., Mahy, L., et al. 2022, *NatAs*, **6**, 1085
 The pandas development team 2020, pandas-dev/pandas: Pandas, v1.1.1, Zenodo, doi:10.5281/zenodo.3509134
 Thompson, T. A., Kochanek, C. S., Stanek, K. Z., et al. 2019, *Sci*, **366**, 637
 van der Walt, S., Colbert, S. C., & Varoquaux, G. 2011, *CSE*, **13**, 22
 Virtanen, P., Gommers, R., Oliphant, T. E., et al. 2020, *NatMe*, **17**, 261
 Webbink, R. F. 1984, *ApJ*, **277**, 355
 Wen, L., Levine, A. M., Corbet, R. H. D., & Bradt, H. V. 2006, *ApJS*, **163**, 372
 Zevin, M., Bavera, S. S., Berry, C. P. L., et al. 2021, *ApJ*, **910**, 152
 Zevin, M., Spera, M., Berry, C. P. L., & Kalogera, V. 2020, *ApJL*, **899**, L1
 Zhang, S. N., Feroci, M., Santangelo, A., et al. 2016, *Proc. SPIE*, **9905**, 99051Q