



Monte Carlo simulations of black hole mergers in AGN discs: Low χ_{eff} mergers and predictions for LIGO

B. McKernan^{1,2,3*}, K. E. S. Ford^{1,2,3}, R. O’Shaughnessy⁴ and D. Wysocki⁴

¹Department of Astrophysics, American Museum of Natural History, New York, NY 10024, USA

²Graduate Center, City University of New York, 365 5th Avenue, New York, NY 10016, USA

³Department of Science, BMCC, City University of New York, New York, NY 10007, USA

⁴Center for Computational Relativity and Gravitation, Rochester Institute of Technology, Rochester, NY 14623, USA

Accepted 2020 March 10. Received 2020 February 13; in original form 2019 July 5

ABSTRACT

Accretion discs around supermassive black holes are promising sites for stellar mass black hole mergers detectable with LIGO. Here we present the results of Monte Carlo simulations of black hole mergers within 1-d AGN disc models. For the spin distribution in the disc bulk, key findings are: (1) The distribution of χ_{eff} is naturally centred around $\bar{\chi}_{\text{eff}} \approx 0.0$, (2) the width of the χ_{eff} distribution is narrow for low natal spins. For the mass distribution in the disc bulk, key findings are: (3) mass ratios $\tilde{q} \sim 0.5\text{--}0.7$, (4) the maximum merger mass in the bulk is $\sim 100\text{--}200 M_{\odot}$, (5) ~ 1 per cent of bulk mergers involve $\text{BH} > 50 M_{\odot}$ with (6) $\simeq 80$ per cent of bulk mergers are pairs of first generation BH. Additionally, mergers at a migration trap grow an IMBH with typical merger mass ratios $\tilde{q} \sim 0.1$. Ongoing LIGO non-detections of black holes $> 10^2 M_{\odot}$ puts strong limits on the presence of migration traps in AGN discs (and therefore AGN disc density and structure) as well as median AGN disc lifetime. The highest merger rate occurs for this channel if AGN discs are relatively short-lived (≤ 1 Myr) so multiple AGN episodes can happen per Galactic nucleus in a Hubble time.

Key words: accretion, accretion discs – black hole physics – gravitational waves – galaxies: active.

1 INTRODUCTION

Advanced LIGO (Aasi et al. 2015) and Advanced Virgo (Acernese et al. 2015) have revealed a population of merging black holes (BHs) that is surprisingly numerous at the upper end of the rate estimate, and significantly more massive than those observed in our own Galaxy (LIGO 2019). Most tantalizingly, the first few mergers are associated with a low observed χ_{eff} , which could imply a precursor population of BHs biased towards low spin, or a precursor population of BHs biased towards anti-alignment or a wide range of spin alignments. Important channels for BH merger are likely to include field binary mergers (e.g. Belczynski et al. 2010; deMink & Mandel 2016), dynamical mergers in globular clusters or galactic nuclei (e.g. Antonini 2016; Rodriguez et al. 2016) and active galactic nucleus (AGN) discs (McKernan et al. 2014; Yang et al. 2019). The next few years of LIGO operation will increase the population of black hole mergers and allow us to construct underlying prior distributions of mass and spin in the precursor BHs (e.g. Fishbach, Holz & Farr

2017; Gerosa & Berti 2017; Wysocki, Lange & O’Shaughnessy 2018).

In the local Universe, black hole density seems greatest in our own Galactic nucleus. The observed rate of occurrence of black hole X-ray binaries implies a cusp of BHs in the central parsec (Generozov et al. 2018; Hailey et al. 2018) consistent with previous conjectures (Morris 1993; Miralda-Escudé & Gould 2000). As a result one of the most potentially promising sites to generate a high rate of black hole mergers yielding overmassive BHs detectable with LIGO are AGN discs (McKernan et al. 2012, 2014; Bartos et al. 2017; Stone et al. 2017; McKernan et al. 2018; Secunda et al. 2019). The hints from LIGO/Virgo’s O1 and O2 of low χ_{eff} are intriguing and hopefully strongly constraining – we require not only massive BH’s, but also binary spin configurations that produce low χ_{eff} often enough to explain the mergers in O1 and O2.

2 LOW χ_{EFF} DOES NOT NECESSARILY MEAN LOW SPIN

To provide context for our simulation results, let us consider how low χ_{eff} BHs can arise in the context of AGN disc mergers. First, it

* E-mail: bmckernan@amnh.org

is important to understand precisely what LIGO measures when it comes to black hole spin. Gravitational wave measurements reliably constrain the combination (Abbott et al. 2016a)

$$\chi_{\text{eff}} = \frac{M_1 \vec{a}_1 + M_2 \vec{a}_2}{M_1 + M_2} \cdot \vec{L} \quad (1)$$

when a merger occurs between two BHs of masses M_1, M_2 , and spins \vec{a}_1, \vec{a}_2 respectively. Notably, this combination is nearly conserved in GR and only involves the projection of the spins on to the angular momentum (\vec{L}) of the merging binary. For our purposes we can rewrite χ_{eff} as

$$\chi_{\text{eff}} = \frac{M_1}{M_b} |a_1| \cos \phi_1 + \frac{M_2}{M_b} |a_2| \cos \phi_2, \quad (2)$$

where (ϕ_1, ϕ_2) are the angles between the spins \vec{a}_1, \vec{a}_2 and the orbital angular momentum of the binary \vec{L}_{bin} , which we assume is aligned or anti-aligned with the angular momentum of the disc gas, and $M_b = M_1 + M_2$ is the binary mass.

From equation (2), low χ_{eff} could occur in a given merger if both spin magnitudes $|a_1, a_2|$ are small. However, χ_{eff} could also be small if spins \vec{a}_1, \vec{a}_2 are anti-aligned so that $\cos(\phi_1, \phi_2)$ are opposite sign. Or, χ_{eff} could be small if a_2 is large, a_1 is small, and $m_1 \gg m_2$. Or, if both $\cos(\phi_1, \phi_2)$ are small, which can happen if both spins are oriented at large angles to \vec{L}_{bin} , then low χ_{eff} can occur.

2.1 How we might get a χ_{eff} black hole binary?

In this section we describe our prior assumptions for BH in AGN discs, and how these assumptions produce binaries with preferentially low χ_{eff} . For BH spin magnitudes, we assume an initial uniform spin magnitude distribution between $a = [-1, +1]$. For spin alignment relative to the AGN gas disc, we assume the initial angle ϕ between a given BHs' spin and the angular momentum of the disc gas (\vec{L}_{disk}) is drawn from a uniform distribution over $\phi \in [0, \pi]$, with $\phi = 0(\pi)$ corresponding to perfect alignment(anti-alignment) with the disc gas angular momentum. [This assumption modestly favours alignment or anti-alignment relative to the gas disc, compared to isotropic spin misalignment on the sphere.] Thirdly, we draw masses randomly from our population, which for the purposes of this section is uniform in m_1, m_2 , though the result does not change appreciably for a power-law mass distribution. Based on these assumptions, Fig. 1 shows the inferred distribution of χ_{eff} for two choices of the maximum value of a (1 and 0.5).

To better understand why even BH spin distributions which individually weakly favour alignment produce χ_{eff} preferentially near zero, we consider a simplifying limit, where the more massive BH (labelled '1' below) dominates χ_{eff} , as usually holds for our power-law mass distributions. Fig. 2 shows the resulting allowed range of $\chi_{\text{eff}} \simeq |a_1| \cos(\phi_1)$ for our initial population of BHs as a function of ϕ_1 , bounded by the black and dashed red lines. For example a black hole in the AGN disc with initial angle $\phi_1 \sim 0$ would produce, if it dominates a binary, an χ_{eff} randomly drawn from a uniform distribution between $\chi_{\text{eff}} \sim [0, 1]$. By contrast, a black hole with initial $\phi \sim 1.0(2.0)$ rad would have χ_{eff} randomly drawn from a uniform distribution between $\chi_{\text{eff}} \sim [0, 0.5]([-0.5, 0])$. The solid red line in Fig. 2 indicates median χ_{eff} for the initial black hole population as a function of ϕ . From Fig. 2 any BHs with ϕ close to $\pi/2$ must produce small $|\chi_{\text{eff}}|$. Conversely, if ϕ is significantly different than $\pi/2$, roughly half of BHs will have spins and thus $|\chi_{\text{eff}}|$ less than 1/2. Thus, we expect that among the first generation of mergers in this channel, $|\chi_{\text{eff}}|$ must be biased towards low values.




Figure 1. Inferred distribution of χ_{eff} for individual BHs allowed envelope for χ_{eff} for individual BHs embedded in an AGN disc drawn from a population of uniform flat spin $a \in [-1, +1]$ where the spin vectors are oriented at angle ϕ compared to the disc gas. The first generation of mergers in this channel will have χ_{eff} drawn from this distribution. The angle ϕ is drawn from a uniform distribution $\phi \in [0, \pi/2]$ rad for $a > 0$ and $\phi \in [\pi/2, \pi]$ rad for $a < 0$. The blue and red solid lines correspond to the distribution of χ_{eff} under the assumptions that $a_i < 1$ and $a_i < 0.5$, respectively. The dotted lines are calculated assuming $a_2 = 0$ (i.e. that the first spin dominates χ_{eff}).




Figure 2. Allowed envelope for χ_{eff} for binaries dominated by an individual BH embedded in an AGN disc drawn from a population of uniform flat spin $a \in [-1, +1]$ where the spin vectors are oriented at angle ϕ compared to the disc gas. The first generation of mergers in this channel will have χ_{eff} drawn from this distribution. The angle ϕ is drawn from a uniform distribution $\phi \in [0, \pi/2]$ rad for $a > 0$ and $\phi \in [\pi/2, \pi]$ rad for $a < 0$. The black solid line corresponds to the envelope of the distribution of χ_{eff} as a function of ϕ . The red solid line corresponds to median χ_{eff} as a function of ϕ but would also correspond to the envelope of χ_{eff} for a low natal spin magnitude distribution concentrated between $a \in [-0.5, +0.5]$.

2.2 How do we get low χ_{eff} mergers in this channel?

In AGN discs, we expect multiple generations of mergers (McKernan et al. 2012, 2014). In the first generation of mergers in this model, the spin magnitudes and orientations are random, and so from Fig. 2, we are drawing from a population that will be biased towards low values of χ_{eff} . Once BHs encounter each other within their mutual Hill sphere, the binaries have $\mathcal{O}(50:50)$ odds that their

orbital angular momentum about its centre of mass is prograde or retrograde. Therefore around 1/4 of first generation mergers will have both black hole spins anti-aligned with the orbital angular momentum, potentially creating a genuinely low-spin black hole (see Section 3.3 below). The net spin that results can be low magnitude (see equation 7 below for dependencies) but will also tend to be aligned or anti-aligned with the disc gas. Additionally, $\sim 1/2$ of first generation mergers will consist of BHs with opposite sign spins, also potentially yielding low spin magnitude BHs.

Second generation mergers will mainly consist of mergers between first and second generation BHs. The second generation of BH will be aligned or anti-aligned with the disc gas, but the first generation remains randomly aligned. Thus, the 1g–2g mergers will have a higher fraction of larger χ_{eff} mergers, but will still favour low χ_{eff} . As the generations evolve, collisions between second or third generations will tend to consist of fully aligned or anti-aligned mergers. Around half of the later generation of 2g–2g or higher mergers will consist of anti-aligned spin BH, and therefore a low χ_{eff} . An additional $\sim 1/4$ of later generation mergers will have spins anti-aligned with binary angular momentum. Thus, low χ_{eff} mergers should still occur, but the distribution should broaden with increasing generation.

3 ASSUMPTIONS BEHIND THE SIMULATIONS

We ran a series of simulations of embedded black hole migrators in 1-d AGN disc models. In this section we discuss the assumptions behind the different components of our simulations, divided according to useful sub-sections.

3.1 The AGN disc and the BHs in the disc

We started with two different models for the gas discs, given by the Sirko & Goodman (2003) and Thompson, Quataert & Murray (2005) disc models. The disc models are calculated for a central supermassive black of mass $M = 10^8 M_{\odot}$ (Sirko & Goodman 2003) and $\sim 10^9 M_{\odot}$ (Thompson et al. 2005) respectively. We assumed the disc was populated by BHs drawn from an initial mass function $M^{-\gamma}$, with lower and upper bounds given by $[5, 50] M_{\odot}$. We note the process of grind-down hardens the power-law index γ relative to the physical IMF appropriate to the AGN disc environment (Yang et al. 2019). We assumed that the black hole spherical component in the galactic nucleus consisted of $\sim 2 \times 10^4$ BHs (Generozov et al. 2018), and that a fraction N_{SBH} between ~ 0.5 per cent and 5 per cent of this component lived in the disc, corresponding to average aspect ratios for the disc models in Sirko & Goodman (2003) and Thompson et al. (2005) respectively. We assumed that the BHs were all settled in the equatorial plane of the disc, so that any binaries that form would have orbital angular momentum either aligned or anti-aligned with the disc gas orbital angular momentum ($\vec{L}_{\text{bin}} = \pm \vec{L}_{\text{disc}}$). The BHs are assumed to lie distributed along the 1-d radial density profile of each disc model.

3.2 Black hole orbits and migration

We assumed that half the initial black hole population in the disc (N_{SBH}) lay on prograde circular Keplerian orbits and half the initial population lay on retrograde circular orbits. We also assumed that a component of the spherical population of BHs would end up with orbits ground-down into the AGN disc during the run, adding an additional N_{grind} population, all with prograde orbits but with random spin magnitudes and orientations. We assumed that gas

torques cause embedded black hole orbits to change over time as they undergo migration (a change of semimajor axis) within the disc. The so-called Type I migration time assuming a Sirko & Goodman (2003) model disc is given by McKernan et al. (2018) as

$$\tau_{\text{mig}} \sim 0.03 \text{ Myr} \left(\frac{N}{3} \right)^{-1} \left(\frac{R}{10^3 r_g} \right)^{1/2} \left(\frac{M}{5 M_{\odot}} \right)^{-1} \left(\frac{h/R}{0.01} \right)^2 \times \left(\frac{\Sigma}{10^7 \text{ kg m}^{-2}} \right)^{-1} \left(\frac{M}{10^8 M_{\odot}} \right)^{3/2} \quad (3)$$

and this is the time for a massive migrator to migrate onto the SMBH through the disc. The equivalent Type I migration time for a typically thinner, but less dense Thompson et al. (2005) model disc can be parametrized as

$$\tau_{\text{mig}} \sim 0.3 \text{ Myr} \left(\frac{N}{3} \right)^{-1} \left(\frac{R}{10^3 r_g} \right)^{1/2} \left(\frac{M}{5 M_{\odot}} \right)^{-1} \left(\frac{h/R}{10^{-3}} \right)^2 \times \left(\frac{\Sigma}{10^4 \text{ kg m}^{-2}} \right)^{-1} \left(\frac{M}{10^9 M_{\odot}} \right)^{3/2} \quad (4)$$

or approximately $\times 10$ slower migration than (Sirko & Goodman 2003) at the thinnest part of the Thompson et al. (2005) disc model. However, as we have shown in previous work, there can be regions of AGN discs where the net torque on a migrator is zero, leading to the formation of a migration trap within the disc, where masses can encounter each other and merge quickly (Bellovary et al. 2016; Secunda et al. 2019). Here we followed (Bellovary et al. 2016) and assumed that a migration trap lay at $\sim 700 r_g$ in the Sirko & Goodman (2003) model disc and at $\sim 500 r_g$ in the Thompson et al. (2005) model disc¹. All BHs on prograde orbits are assumed to migrate inwards at a rate given by equation (3) from orbits lying outside the migration trap. BHs interior to the migration trap (few in number since this is a small fraction of the area of the disc) are assumed to migrate outwards through the disc to the migration trap. Retrograde BHs are assumed to not migrate since the spiral density perturbations generated by retrograde orbiters are ~ 1 per cent that of the prograde migrators (McKernan et al. 2014), and so we assume migration torques on retrograde orbiters are negligible for the few Myr an AGN disc might persist.

3.3 Black hole spins

The BHs were given initial dimensionless spin parameter magnitudes (a) drawn from a flat, uniform distribution between $a = [-1, +1]$. The BHs have an angle ϕ between their spin vector (\vec{a}) and the angular momentum vector of the disc (\vec{L}_{disc}). For BHs with prograde spin ($a > 0$), the initial value of ϕ was drawn from a uniform distribution $\phi = [0, 90^\circ]$ where $\phi = 0^\circ$ corresponds to perfect alignment between the prograde spin and \vec{L}_{disc} . For retrograde spin BHs ($a < 0$), the initial value of ϕ was drawn from a uniform distribution $\phi = [90.01^\circ, 180^\circ]$ where $\phi = 180^\circ$ corresponds to perfect anti-alignment with \vec{L}_{disc} . By assuming a uniform distribution in ϕ rather than $\cos(\phi)$ we are assuming that the addition of the AGN disc to the nuclear cluster (which dominates the masses of individual embedded objects) has had a minor effect on the initial distribution of spins. Over time, the spins of the initial black hole population are allowed to evolve as gas accretion reduces $|a|$ for the $a < 0$ population and increases $|a|$ for the

¹Note our migration trap locations are a factor of two different from Bellovary et al. (2016) since in that paper the migration traps should have been written in units of the Schwarzschild radius $r_{\text{Sch}} = 2r_g$

$a > 0$ population. Accreting gas mass will also tend to torque the spin into alignment with the disc gas angular momentum, effectively decreasing ϕ towards zero (see the next section).

3.4 Black hole accretion

BHs on prograde orbits in the disc were assumed to accrete at a fraction f_{Edd} of the Eddington rate. For most simulations we used $f_{\text{Edd}} \sim 1$. As gas accretion occurs we also expect spin alignment with the disc in the limit of 1 per cent – 10 per cent mass growth from the disc (Bogdanovic, Reynolds & Miller 2007). The mass-doubling time is ~ 40 Myr at the Eddington accretion rate, so several Myr of gas accretion should be required in order to align all the BHs spins with the AGN disc. Spin up or down via accretion takes a relatively long time compared to torquing of angle ϕ , and if AGN discs are short lived, gas accretion will have little impact on the spin magnitude. More important in the context of an AGN disc, it turns out, is the net spin of a merged black hole (see Section 3.6 below). We expect most mergers to occur early on in the disc's lifetime. Early in the AGN lifetime, gas damping is circularizing BH orbits and the number of collision targets is a maximum (McKernan et al. 2012, 2014). We assumed that spinning BHs which were not aligned with the angular momentum of the gas disc (i.e. $\cos\phi > 0$) would experience a torque due to accreting gas, thus driving the black hole spin towards alignment with the gas disc. Using the maximum estimate from Bogdanovic et al. (2007) we assumed that ~ 1 per cent mass accretion at $f_{\text{Edd}} \sim 1$ was sufficient to drive a prograde orbiting black hole spin into alignment with the gas disc. We assume the rate of torquing of alignment is uniform over time. While we may be overestimating the rate of spin alignment by up to an order of magnitude, super-Eddington gas accretion ($f_{\text{Edd}} > 1$) is also possible for BHs embedded in the disc, so we believe our choice of alignment rate is a reasonable starting point. Throughout, we ignore the effect of black hole accretion on the surrounding gas and so we assume that the migration and gas hardening torques are unaffected by the backreaction of the radiation and winds. We will build a feedback prescription into future simulations based on the results of numerical experiments.

3.5 Binary formation and hardening

We assume that the binary initial fraction is zero. However, by distributing the BHs randomly across the AGN disc, a fraction randomly end up close to each other. We assume a binary forms when two prograde orbiters lie within their mutual Hill sphere

$$R_{\text{Hill}} = R \left(\frac{q_{\text{bin}}}{3} \right)^{1/3}, \quad (5)$$

where R is the semimajor axis of the centre of mass of both BHs and $q_{\text{bin}} = M_b/M_{\text{SMBH}}$ is the ratio of the binary mass ($M_b = M_1 + M_2$) to the SMBH mass (assumed to be $10^8 M_\odot$). The disc is assumed to extend to a distance of $2 \times 10^4 r_g$, so two prograde BHs $M_1 = 50 M_\odot$, $M_2 = 50 M_\odot$ both with semimajor axes $R \sim 10^4 r_g$ form a binary if they lie within $R_H \sim 70 r_g (R/10^4 r_g)(M_{\text{bin}}/100 M_\odot)^{1/3}$ of their centre of mass. In principle a binary can form between a retrograde black hole and a prograde black hole if the separation between the two is $\sim q_{\text{bin}} R$, but the probability of this occurring is very small.

Once a binary forms we assume that about half the binaries randomly form with prograde angular momentum and half form randomly with retrograde angular momentum. In numerical experiments, we have actually found that out of 30 binary formation

encounters that we studied, approximately 17/30 were retrograde, 8/30 were prograde, and 5/30 were indeterminate (see Secunda et al., in preparation). So, the ratio of retrograde: prograde binaries may actually be closer to 2:1, but for our simulations we assumed 1:1 throughout. We will explore differences in this ratio in future work.

The rate of hardening is assumed to follow the Baruteau, Cuadra & Lin (2011) results where retrograde binaries harden $\sim \times 5$ faster than prograde binaries. In particular, it takes ~ 200 orbits around the centre of mass for a retrograde binary to halve its semimajor axis. By assuming that this hardening prescription applies throughout the binary lifetime, binaries experience a runaway hardening effect from gas torques and rapidly get driven to the regime of GW-emission domination, whereupon the binary promptly merges. This torque prescription fails after several halvings of the semimajor axis, the binary will be 'hard' compared to the encounter velocity of prograde migrators, so tertiary encounters will become important in hardening the binary to the point where GW emission dominates. Approximately 10 such encounters would be required to harden a binary with semimajor axis ($R_H/2$) to GW-dominated merger (Leigh et al. 2018). We will investigate stalled binary hardening and the importance of tertiary encounters in future work.

3.6 Mergers

Merging binary black holes produce a remnant, whose final mass and spin reflects the input BH properties and the overall radiative mass and angular momentum losses during merger. For clarity, we will not employ very accurate but black-box approximations like Varma et al (2019), instead adopting simplified but transparent approximations. We shall approximate the merged mass as (Tichy & Marronetti 2008)

$$M_{\text{final}} = M_b (1 - 0.2\nu - 0.208\nu^2(a_1 + a_2)), \quad (6)$$

where $\nu = \mu/M_b$, the symmetric mass ratio, or $\nu = q_b/((1 + q_b)^2)$ where $q_b = M_2/M_1$ is the binary mass ratio. We assume that the merged spin magnitude is given by (Tichy & Marronetti 2008)

$$a_{\text{merger}} \approx 0.686 (5.04\nu - 4.16\nu^2) + 0.4 \left(\frac{a_1}{(0.632 + 1/q)^2} + \frac{a_2}{(0.632 + q)^2} \right). \quad (7)$$

In this expression, when calculating the magnitude of the post-merger spin, we assume ϕ_i are small. We assume that black hole binaries always form with orbital angular momentum oriented parallel or antiparallel to the angular momentum of the gas in the disc so $L_\pm = +1(-1)$ is aligned (anti-aligned) with the angular momentum of the gas. Thus, ϕ_{merger} is always fully aligned with the disc (0rad for $L = +1$) or anti-aligned with the disc (π rad for $L = -1$).

When gravitational radiation starts to dominate the BBH inspiral, the BH is still in a sufficiently wide orbit that its orbital angular momentum dominates over the component spin angular momenta. Moreover, to a good first approximation, the total angular momentum direction is conserved during binary inspiral, and the remnant BH's spin direction points along this same axis. We therefore assume the remnant BH has perfect alignment with its initial orbital angular momentum direction. For comparable-mass binaries, the orbital angular momentum at merger dominates, so the remnant angular momentum direction will point in the same direction as the orbit, independent of BH spins. Conversely, for highly asymmetric binaries, the orbital angular momentum at merger is much smaller

than the binary's spin angular momentum, and the spin angular momentum direction is unchanged. This effect is captured in equation (7), above.

The merging BH experiences a recoil kick which only very slightly perturbs its orbit, given the BH's large Keplerian orbital velocity around the central SMBH. Because our BH spins are largely aligned or antialigned with the disc, the kick magnitude never reaches the most extreme options possible, nor does it misalign the remnant's orbit from the disc. For the SMBH and BH of interest, recoil kicks would perturb orbital eccentricity by a small amount (i.e. $v_{\text{recoil}}/v_{\text{Kep}} \ll 1$). Since orbital eccentricity damping is very fast in AGN discs (McKernan et al. 2012) we expect recoil to not significantly impact the merger or evolutionary history of BHs in the disc (barring exceptionally high numbers of BHs in the disc, so BHs are often separated by only a few Hill radii). We therefore ignore BH recoil kicks in all subsequent discussion.

3.7 Grind-down population

When BHs lie on orbits inclined to the disc they are expected to experience a head-wind and drag force as they pass through the disc. The net result is to extract energy from the black hole orbit and grind its inclination down into the plane of the disc over time (e.g. Syer, Clarke & Rees 1991; Subr & Karas 2005; McKernan et al. 2014). Grind down of orbits is most efficient at small radii, so we assumed that some of the spherical component of BHs is ground down into the disc at radii $< 5000 r_g$ over time. We typically assumed $N_{\text{grind}} \sim 100 \text{ Myr}^{-1}$, corresponding to $\mathcal{O}(10 \text{ per cent})$ of a typical initial disc population.

The grind-down rate contributes directly to the rate of first-generation mergers drawn from this population, as well as providing the seeds for all future generations of mergers.

3.8 Caveats

One of the biggest caveats for our simulations is that we ignore all tertiary encounters. Tertiary encounters can both help and hinder binary formation. For example, if binaries can only be hardened by gas torques to a modest separation and stall thereafter, the resulting population of hard binaries can become hardened to merger via $\mathcal{O}(10)$ tertiary encounters (Leigh et al. 2018). This is because the relative velocity of a prograde encounter is less than the binary velocity around its own centre of mass (Sigurdsson & Phinney 1993). Conversely, encounters between a binary and either the spherical component of BHs plunging through the disc or a black hole on a retrograde orbit can ionize wide binaries, since the relative velocity of the encounter is large (Leigh et al. 2018).

We ignore scattering between retrograde and prograde BHs upon close encounters, although we have started numerical studies of this effect (Secunda et al., in preparation). We also assume that the retrograde BHs have zero orbital eccentricity. This is unlikely. BHs on prograde orbits have initial eccentricities efficiently damped in a short time ($< 0.1 \text{ Myr}$) (McKernan et al. 2012). However, BHs on retrograde orbits are unlikely to have their initial eccentricities damped by much, if at all. As a result, the rate of encounter between retrograde and prograde orbiters could be higher than presented here, with obvious consequences for the rate of three-body encounters and binary ionization as well as scattering events. A future version of this simulation will build in scattering events based on numerical studies of prograde and retrograde encounters (Secunda et al., in preparation).

We also ignore complicated multiple encounters within a Hill sphere. Particularly early on, several BHs may lie within a mutual Hill sphere. For now our primary assumption is that the BH pair separated by the smallest 1-d radial distance form a binary and ignore the complexity of interactions with other nearby BHs, although we also later test an alternative scenario where the most massive BH within a mutual Hill sphere form a binary. Furthermore, since our disc model is 1-d, we ignore 2-d and 3-d complications like resonances or slightly inclined orbits. We also ignore precession effects. These are problems we will tackle in future extensions of this work.

Finally, we ignore objects other than BHs. There should be a large population of other stellar remnants and stars embedded in the disc (McKernan et al. 2012) also subject to similar gas torques. This will lead to binary formation between neutron stars, white dwarfs, stars, and BHs. If gas torques are efficient at merging such binaries, electromagnetic signatures may result (McKernan et al. 2018). However, if gas torques lead to stalled hard binaries, BHs in binaries with non-BHs may end up swapping partners in tertiary encounters. Such encounters will be the subject of future numerical work.

4 SIMULATION RESULTS

In this section we show results from individual runs to illustrate some of the possible variations within our model. We shall then outline the Monte Carlo results of a large number of runs, characterizing the parameter distributions. In Table 1 we list a range of individual runs to illustrate our choice of particular input parameters. In Table 2 we list key results from each of the individual runs.

4.1 Initial results from individual disc runs

Table 1 lists twelve fiducial runs (R1–R12) that illustrate a range of possible scenarios for merging stellar mass black holes in AGN discs. In Table 2 we list some of the key results from the fiducial runs. We divide the results in Table 2 into black holes merging in the bulk of the disc and the black holes merging at a migration trap (if present). For context, the average mass ratio (\bar{q}) and $\bar{\chi}_{\text{eff}}$ of the first ten black hole mergers detected with LIGO are $\bar{q} \sim 0.7$ and $\bar{\chi}_{\text{eff}} \sim 0.06$ (LIGO 2019).

We illustrate our results by discussing outcomes of representative runs. In R1 from Table 1 we had $N_{\text{SBH}} = 869$ BHs initially, drawn from an IMF with $\gamma = -1$, and there are a further $N_{\text{grind}} = 100$ added during the 1 Myr disc lifetime. Fig. 3 shows the black hole masses involved in mergers over time. The black crosses correspond to mergers occurring in the disc away from the migration trap. The red filled-in circles correspond to mergers at the migration trap and reveal the rapid and steady growth of the IMBH at the migration trap. A majority of mergers in the bulk ($> 50 \text{ per cent}$) occur early on ($< 0.1 \text{ Myr}$) as BHs that happen to lie close to each other in the disc merge quickly. We estimate that our initial random radial distribution was equivalent to an effective initial binary fraction of ~ 0.15 in R1. Following this burst of mergers, the number of possible targets for subsequent merger then drops and BHs have to migrate in order to find new targets for merger.

The evolution of the mass spectrum for R1 is shown in Fig. 4. The black solid line corresponds to the input initial mass function (IMF) and the red solid line corresponds to the mass function after 1 Myr. The IMF has evolved towards a broken power-law distribution, with the break lying near the high mass end of the IMF (McKernan et al. 2018), along with a single IMBH. Note that the low mass end of

Table 1. Sample runs to illustrate the variations that arise with changing input parameters. Column 1 is the run name. Column 2 is the initial number of BH embedded in the disc. Column 3 is the number of BH ground down by the disc per Myr. Column 4 is the IMF index $M^{-\gamma}$. Columns 5(6) are the lower (upper) bounds to the IMF. Column 7 is the AGN lifetime in Myr. Column 8 is the spin distribution (u=uniform, flat between $a = [-1, +1]$). Column 9 is the position of the migration trap if present ($700r_g$ in the (Sirko & Goodman 2003) disc model and $250r_g$ in the (Thompson et al. 2005) disc model. Column 10 shows the choice of disc model (SG=(Sirko & Goodman 2003); TQM=(Thompson et al. 2005)). Column 11 shows the ratio of binary hardening time-scale for retrograde (t_-) versus prograde (t_+) binaries, with $t_-/t_+ = 5$ the result from (Baruteau et al. 2011).

Run	N_{BH}	N_{gr} (/Myr)	γ	M_{Lower} (M_{\odot})	M_{Upper} (M_{\odot})	τ_{AGN} (Myr)	a	trap (r_g)	disc	t_-/t_+
R1	869	10^2	1	5	50	1	u	$700r_g$	SG	5
R2	869	10^2	1	5	50	1	u	$700r_g$	SG	1
R3	100	10^2	1	5	50	1	u	$700r_g$	SG	5
R4	851	10^2	2	5	50	1	u	$700r_g$	SG	5
R5	851	10^2	2	5	50	5	u	$700r_g$	SG	5
R6	851	10^2	2	5	15	1	u	$700r_g$	SG	5
R7	851	10^2	2	5	50	1	(1-a)	$700r_g$	SG	5
R8	851	0	2	5	50	1	u	$700r_g$	SG	5
R9	851	0	2	5	50	5	u	$700r_g$	SG	5
R10	851	0	2	5	50	5	u	$700r_g$	SG	1
R11	851	10^2	2	5	50	1	u	none	SG	5
R12	851	10^2	2	5	50	1	u	$500r_g$	TQM	5

Table 2. Initial results from sample runs. Column 1 is the run name. Column 2 is the number of mergers in the bulk of the disc during the run. Column 3 is the number of mergers at the trap (if present). Column 4 is the ratio of bulk mergers involving at least 1 BH of $2g$ or higher, or $ngmg/1g1g$ where $n > 1$, $m \geq 1$. Column 4 is the median mass ratio (\tilde{q}) per bulk merger with associated standard deviation. Column 5 shows the range of q for bulk mergers during the run. Column 6 is the median $\tilde{\chi}_{\text{eff}}$ for mergers in the bulk, with associated standard deviation. Column 7 shows the range of χ_{eff} for the bulk mergers during the run. Column 8 shows the largest mass BH in the run (mostly at trap). Column 9 shows the time of 1st merger at the trap and Column 10 the mass of the 1st merger at the trap.

Run	N_{bulk}	N_{trap}	$ngmg/1g1g$	\tilde{q}	$[q_{\text{min}}, q_{\text{max}}]$	$\tilde{\chi}_{\text{eff}}$	$[\chi_{\text{min}}, \chi_{\text{max}}]$	M_{IMBH} (M_{\odot})	t_{trap} (Myr)	M_{trap} (M_{\odot})
R1	227	25	0.22	0.46 ± 0.24	[0.08,1.0]	0.07 ± 0.37	$[-0.88, 0.83]$	552.7	0.12	63.0
R2	189	26	0.16	0.44 ± 0.25	[0.06,1.0]	0.08 ± 0.40	$[-0.94, 0.87]$	745.0	0.12	73.0
R3	16	21	0.19	0.47 ± 0.20	[0.18,1.0]	-0.09 ± 0.37	$[-0.83, 0.56]$	486.0	0.38	46.6
R4	172	17	0.15	0.63 ± 0.27	[0.10,1.0]	0.03 ± 0.33	$[-0.79, 0.88]$	272.2	0.24	42.1
R5	329	159	0.29	0.50 ± 0.25	[0.07,1.0]	0.03 ± 0.44	$[-0.96, 0.94]$	2409.7	0.14	30.2
R6	175	14	0.09	0.70 ± 0.19	[0.26,1.0]	-0.03 ± 0.34	$[-0.79, 0.81]$	126.6	0.19	12.1
R7	178	17	0.18	0.57 ± 0.26	[0.09,1.0]	0.01 ± 0.23	$[-0.59, 0.82]$	270.2	0.24	42.7
R8	157	1	0.08	0.60 ± 0.26	[0.11,1.0]	0.07 ± 0.37	$[-0.86, 0.88]$	16.3	0.24	16.3
R9	221	27	0.27	0.50 ± 0.26	[0.07,1.0]	0.00 ± 0.44	$[-0.94, 0.94]$	709.8	0.24	16.3
R10	221	20	0.25	0.49 ± 0.26	[0.05,1.0]	0.14 ± 0.46	$[-0.93, 0.93]$	503.4	0.24	24.1
R11	181	0	0.18	0.54 ± 0.27	[0.05,1.0]	0.00 ± 0.38	$[-0.89, 0.89]$	139.4	none	none
R12	53	149	0.08	0.64 ± 0.24	[0.14,1.0]	0.00 ± 0.31	$[-0.68, 0.64]$	1625.5	0.06	36.2

the distribution is supported by the addition of ground-down BHs drawn from a distribution identical to the IMF.

The spins of the BHs also evolve in R1 over time. Fig. 5 shows the evolution of the spins of the initial population on prograde orbits after 1Myr. The black solid line in Fig. 5 shows the initial distribution of spins drawn from a uniform, flat spin distribution. The red solid line in Fig. 5 shows the final state of the spin distribution among this population after 1Myr. The retrograde orbiters are assumed to not accrete from the gas at any significant rate and their rate of interaction with migrating prograde orbiters is very small. Note that the final distribution only gives the magnitude of the spin, it does not take into account the fact that ϕ may have flipped sign (i.e. that the spin points in the opposite direction).

Fig. 6 shows the number of black hole mergers as a function of time in R1. As can be seen from Fig. 3 most mergers occur early on

(<0.1 Myr). We assumed an initial binary fraction of zero, but the random distribution of BHs in the disc corresponds to an effective initial binary fraction of $f_{\text{bin}} \sim 0.15$. These BHs merge quickly in our prescription but then must migrate within the disc to find more partners. Thus, we find a cascade of mergers (growing $\propto t^1$) in the first ~ 0.01 Myr and then growing more like $\propto t^{1/4}$ from 0.01 to 1Myr. This may suggest that AGN discs are most efficient at black hole mergers early on in their lifetimes. A somewhat counterintuitive point then emerges: if AGN discs are short-lived, they may end up increasing the rate of black hole mergers detectable with LIGO. This is because shorter lived AGN discs imply a large rate of AGN episodes per galactic nucleus (for a constant fraction of active nuclei per volume). The AGN ‘grinder’ then gets multiple opportunities to accelerate mergers of BHs within the overdense central parsec.

Fig. 7 shows the mass ratio for mergers in R1 as a function of time. In black are mergers away from the merger trap and in red are




Figure 3. Masses of BHs involved in mergers over time from R1. The black crosses are BHs involved in mergers not in the migration trap. The red filled circles are BHs involved in mergers in the migration trap. Of 252 mergers total in the run, 227 mergers occur away from the migration trap in the bulk of the disc and 25 mergers occur at the migration trap. More than half the mergers in this run (121) occur in the first ~ 0.1 Myr, with 99 in the first 0.05 Myr. Mergers at the migration trap only begin after 0.12 Myr.




Figure 4. Evolution of mass spectrum of BHs in the disc as a result of mergers and gas accretion for R1. The black solid line corresponds to input initial mass function (IMF). The red solid line corresponds to the mass spectrum of BHs in the disc after 0.95 Myr. Ground-down BHs added to the disc over time are drawn from a distribution identical to that for the IMF.

mergers at the migration trap. The mass ratio of mergers in the bulk population spans $q \sim [0.1, 1]$.

Fig. 8 shows the generations of BH involved in mergers in the bulk of R1. $1g+1g$ mergers dominate (78 per cent), with $1g+2g$ mergers making up most of the rest (17 per cent) of the mergers in R1.

Finally, Fig. 9 shows the distribution of χ_{eff} for all the mergers in R1. The χ_{eff} distribution given by the black solid curve corresponds to the distribution for mergers in the bulk of the disc. The red solid line corresponds to the distribution of χ_{eff} for mergers at the trap. Both merger distributions appear centred approximately around $\chi_{\text{eff}} \sim 0$, with the possibility of a bimodal distribution in χ_{eff} in the trap distribution. The red vertical dashed lines confine ≈ 90 per cent of the bulk distribution. Our initial results confirm our expectations from Section 2.1 that the χ_{eff} distribution should be biased towards




Figure 5. Evolution of spin spectrum of BHs in the disc as a result of mergers and gas accretion for R1. The black solid line corresponds to input initial spin distribution, drawn from a uniform, flat distribution. The red solid line corresponds to the spin spectrum of BHs in the disc after 0.95 Myr. Ground-down BHs added to the disc over time are drawn from a spin distribution identical to the initial spin distribution. The final distribution has not been corrected for spin flips due to retrograde binary orbital angular momentum. The angle ϕ contains that information in our simulations.

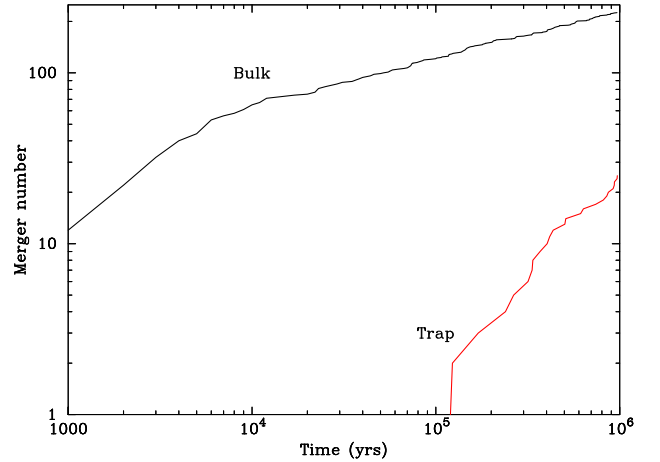


Figure 6. The integrated number of black hole mergers as a function of time in R1 for bulk (black) and trap (red). A majority of mergers (> 50 per cent) occur early on (≤ 0.1 Myr). The initial number of mergers grows as $\sim t^1$ at < 0.01 Myr and then grows as $\sim t^{1/4}$ at > 0.01 Myr. The population of ground-down orbits provides a support for the merger rate at later times.

$\chi_{\text{eff}} \sim 0$ at least in the bulk distribution, although we require a larger scale simulation (below) to confirm this statistically.

4.2 The effect of changing input parameters

From Table 2 we varied a number of our input parameters to illustrate the impact of changes in individual parameters. These results are merely to help guide the reader in what to expect from the full Monte Carlo runs. Here we note some of the gross features of the results, but refrain from discussion of the distributions until we have much larger samples.

In R2 we changed the ratio of gas hardening time-scales between retro- and pro-grade binaries from $t_-/t_+ = 5$ (Baruteau et al. 2011) in R1 to $t_-/t_+ = 1$. As a result, we remove a source of very rapid




Figure 7. The mass ratios of mergers as a function of time in R1. In black are the mass ratios of mergers in the bulk of the disc and in red are the mass ratios of mergers at the migration trap.




Figure 8. Generations of BHs involved in bulk mergers from R1. The black curve is the 1g+ng mergers, where $n \geq 1$. The red curve is 2g+mg mergers where $m \geq 2$. The blue curve is 3g+kg mergers where $k \geq 3$. For example, there are 38 1g+2g mergers in this instance of R1. The majority of mergers are by far 1g+1g.

mergers early on and so it is no surprise the number of mergers in the bulk of R2 is lower than in R1. In R3 we reduced the number N_{BH} from R1 and in Table 2 we can see a much lower number of mergers. The IMBH takes longer to start assembling at the migration trap in R3 compared to R1. However, in the end, R3 yields a similar mass IMBH to that in R1.

In R4 we changed the IMF index to $\gamma = 2$. The changed power-law normalization led to a slightly smaller $N_{\text{BH}} = 851$ compared to R1. After 1 Myr we find a similar number of mergers, but a smaller IMBH and higher \bar{q} than in R1 or R3. In R5, we extended the R4 run to 5 Myr, resulting in substantially more mergers in the bulk of the disc, but an order of magnitude increase in number of mergers at the trap. As a result, R5 yields the largest IMBH of all the runs. Note that the IMBH mass growth from $\sim 30 M_{\odot}$ to $\sim 2000 M_{\odot}$ in 5 Myr corresponds to >6 mass doublings in $\sim 1/8$ Eddington mass-doubling times. Thus, the IMBH grows at a remarkable $\times 40$ – 50 Eddington via collisions, as we expect (McKernan et al. 2012). Also in R5, the fraction of hierarchical mergers (ngmg/1g1g, where $n > 1$, $m \geq 1$) is higher in the bulk than in runs R1–R4.

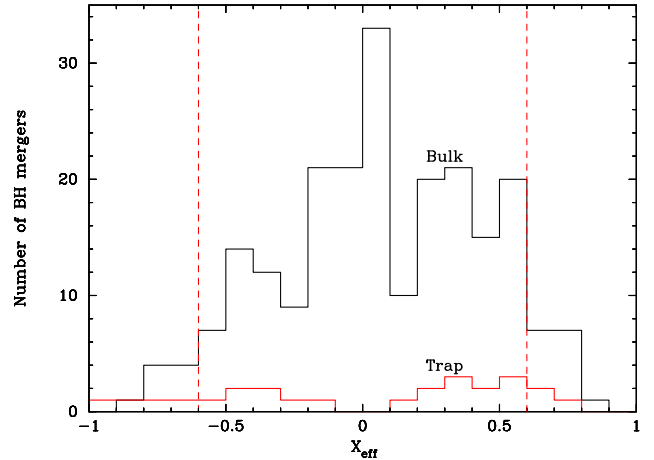


Figure 9. The distribution of χ_{eff} for the black hole mergers in R1. The black solid curve is the χ_{eff} distribution for mergers in the bulk of the disc. The red solid curve is the distribution for mergers at the trap, and hints at bi-modality. The vertical red dashed lines correspond to the bounds of 90 per cent of the mergers. The bulk χ_{eff} distribution appears to be approximately centred around $\chi_{\text{eff}} \sim 0$.

In R6, we changed the IMF to match the masses of BHs observed in our own Galaxy ($[5, 15] M_{\odot}$) which yields a low-mass IMBH after 1 Myr and higher \bar{q} . The low fraction of hierarchical mergers could correspond to the generally lower rates of migration. In R7 we changed the initial spin magnitude distribution to $a(1 - |a|)$, where a is drawn from $a = [-1, +1]$. This narrower spin distribution could physically correspond to a low natal spin population of BHs in the IMF. This results in a much narrower χ_{eff} distribution. In R8 and R9 we investigated the effect of grind-down of orbits into the disc. By removing grind-down, the black hole mass spectrum is no longer supported at the low-mass end with time. Also in R8 after 1 Myr, from Table 2, there is no IMBH. This suggests that orbital grind-down at small radii is a key driver of early IMBH formation at a migration trap. In R9 and R10, we extend the disc lifetime from R8 to 5 Myr, keeping no grind-down, and now we grow an IMBH at the trap. Thus, if τ_{AGN} is long enough, a large IMBH mass can build at a migration trap, independent of grind-down efficiency. The fraction of hierarchical mergers in the bulk increases substantially with no grind-down and longer disc lifetime.

In R11 we removed the migration trap. This prevents the build up of a single massive IMBH, with the largest post-merger BH only $\sim 139 M_{\odot}$.

Finally, in R12, we replaced the disc model of Sirko & Goodman (2003) with Thompson et al. (2005), which orbits a $10^9 M_{\odot}$ SMBH and is thinnest around $10^3 r_g$, with generally lower density, leading to a migration time longer on average than in Sirko & Goodman (2003). In this disc model, the migration trap lies at $500 r_g$, near the thinnest part of the disc where migration is fastest. As a result, the migration trap in a Thompson et al. (2005) disc encounters many BH. Thus, if AGN discs are generally more like the Thompson et al. (2005) model rather than the Sirko & Goodman (2003) model, contain a migration trap and live ≥ 1 Myr, we expect the formation of many massive IMBH in AGN discs.

One immediate conclusion from the final two columns in Table 2 is that if disc lifetimes are actually very short (~ 0.1 – 0.2 Myr), there is very little time to build much mass at migration traps. Thus, ongoing upper limits on BH mergers masses from LIGO are very important astrophysically, since this will put strong limits on disc

Table 3. Large-scale runs to evaluate parameter distributions for mergers in the bulk of the disc. Column 1 lists the run from Table 1, which is then run with different numerical seeds 10^2 times. Column 2 lists the median mass ratio (\tilde{q}) and standard deviation in the merger population. Column 3 lists the full range of q for each run. Column 4 lists the median χ_{eff} and standard deviation for the merger population. Column 5 lists the full range of χ_{eff} for the run. Column 6 lists the median mass and standard deviation of black hole masses involved in mergers. If the standard deviation is greater than the difference between the median mass and the minimum mass in the run, we list the latter instead as the lower limit. Column 7 lists the maximum merged mass and in brackets maximum merged mass when largest mass merger occurs within a Hill sphere. Column 8 lists the total number of mergers for the run. Column 9 lists the percentage of black holes with mass $> 50 M_{\odot}$ involved in all mergers, with in brackets the percentage when largest mass merger occurs within a Hill sphere.

Run	\tilde{q}	$[q_{\min}, q_{\max}]$	$\tilde{\chi}_{\text{eff}}$	$[\chi_{\min}, \chi_{\max}]$	$\tilde{M}_{\text{merger}}$ (M_{\odot})	M_{max} (M_{\odot})	N_{mergers}	$N_{>50 M_{\odot}}$ (per cent)
R1	0.48 ± 0.25	[0.02, 1.0]	-0.01 ± 0.33	[-0.98, 0.98]	$16.3^{+15.0}_{-11.3}$	258.6(270.5)	17 601	1.0(3.0)
R2	0.48 ± 0.25	[0.08, 1.0]	0.00 ± 0.33	[-0.98, 0.97]	$17.0^{+13.4}_{-12.0}$	93.5(236.3)	15 563	1.0(3.0)
R3	0.57 ± 0.26	[0.10, 1.0]	-0.04 ± 0.34	[-0.91, 0.84]	$14.0^{+12.8}_{-9.0}$	92.7(135.2)	761	0.1(3.0)
R4	0.56 ± 0.25	[0.10, 1.0]	-0.01 ± 0.32	[-0.95, 0.98]	$9.0^{+9.1}_{-4.0}$	87.1(164.9)	14 963	0.1(0.7)
R5	0.56 ± 0.24	[0.10, 1.0]	-0.02 ± 0.43	[-0.98, 0.98]	$9.0^{+22.8}_{-4.0}$	144.3(243.3)	29 405	0.2(1.0)
R6	0.71 ± 0.19	[0.33, 1.0]	-0.03 ± 0.31	[-0.98, 0.95]	$7.0^{+2.9}_{-2.0}$	33.3(90.7)	14 298	0.0(0.1)
R7	0.56 ± 0.25	[0.10, 1.0]	-0.01 ± 0.11	[-0.34, 0.64]	$9.0^{+9.2}_{-4.0}$	89.8(151.0)	15 004	0.1(0.8)
R8	0.56 ± 0.25	[0.10, 1.0]	-0.02 ± 0.33	[-0.97, 0.94]	$9.0^{+9.3}_{-4.0}$	91.8(219.4)	14 001	0.2(0.6)
R9	0.55 ± 0.24	[0.10, 1.0]	0.00 ± 0.41	[-0.98, 0.98]	$9.0^{+9.7}_{-4.0}$	83.8(183.0)	19 716	0.3(0.9)
R10	0.55 ± 0.24	[0.10, 1.0]	0.05 ± 0.42	[-0.98, 0.98]	$9.0^{+9.5}_{-4.0}$	92.4(412.0)	18 276	0.2(0.9)
R11	0.56 ± 0.25	[0.10, 1.0]	-0.01 ± 0.32	[-0.98, 0.97]	$9.0^{+9.2}_{-4.0}$	93.8(201.0)	15 311	0.2(0.7)
R12	0.64 ± 0.25	[0.05, 1.0]	-0.10 ± 0.32	[-0.98, 0.94]	$8.1^{+13.2}_{-3.1}$	97.8(120.8)	3992	0.2(0.5)

lifetime, disc structure (including the existence of migration traps), and orbital grind-down efficiency.

4.3 Monte Carlo simulations

We ran individual runs R1–R12 for 100 iterations, each with different initial random numerical seeds, to extract large distributions of BH masses and χ_{eff} . Separately, we also ran versions of runs R1–R12 that incorporated a modified merger condition: where more than two BH lay within a mutual Hill sphere, the two most massive BH were assumed to merge. Table 3 shows the results of Monte Carlo simulations of the various individual runs R1–R12 from Table 1 for mergers in the bulk of the disc. Table 4 shows the results of Monte Carlo simulations of the various individual runs for mergers at the migration trap.

From Table 3, in general, across a wide range of parametrization, the χ_{eff} distribution for mergers in the bulk of the disc is centred on $\tilde{\chi}_{\text{eff}} \approx 0$ with a standard deviation of $\sigma_{\chi} \sim 0.3$. The very small negative bias in runs R1, R3–R8, R10–R12 is due to enhanced production of negative, anti-aligned spin BH by our choice of $t_{-}/t_{+} = 5$. This is confirmed in runs R2 and R10, where $t_{-}/t_{+} = 1$, and $\tilde{\chi}_{\text{eff}} = 0.00, 0.05$ respectively. The χ_{eff} distribution is relatively broad, with $\sigma_{\chi} \sim 0.3$ generally. The only exception is in R7, where a narrow distribution of χ_{eff} with $\sigma_{\chi} \sim 0.1$ corresponds to a narrow, non-uniform, initial spin magnitude distribution. If the distribution of mergers follows a normal distribution, ~ 68 per cent of mergers have $|\chi_{\text{eff}}| < 0.3$ for a flat, uniform, distribution of initial spin magnitudes. As LIGO continues to operate, the width of the χ_{eff} distribution and a net negative bias in $\tilde{\chi}_{\text{eff}}$ may allow us to distinguish between a pure dynamics channel and the AGN disc channel.

Also from column 6 in Table 3, a flatter IMF ($M^{-\gamma}$) with $\gamma \sim 1$ in R1–R3 yields a higher median mass involved in mergers than in R4–R12. Because LIGO is much more sensitive to high-mass binaries, this steeper power-law index produces a detected distribution that is consistent with early LIGO observations; see Fig. 12. Both power-law distributions $\gamma = 1, 2$ are consistent with inferences about

the selection-corrected spectrum of BH masses in BH–BH binaries (Abbott et al. 2019). From column 7 in Table 3, the maximum mass produced in bulk mergers is $O(100\text{--}200 M_{\odot})$ generally. The higher maximum mass is generally associated with assuming the most massive BH within a mutual Hill sphere will merge. Changing the condition on mergers within the Hill sphere made no significant difference to the parameter distributions in R1–R12 otherwise. This is a reasonable mass upper limit to expect if migration traps do not generally occur in AGN discs. Finally, column 9 in Table 3 shows the percentage of black holes in all mergers in the bulk with mass $> 50 M_{\odot}$. If we assume the nearest BH within a mutual Hill sphere merge, then in no run does this percentage rise above 1 per cent. If we change the merger condition to require that the two most massive BH within a mutual Hill sphere to merge, then the percentage of mergers involving a merger in the upper mass gap rises by a factor ~ 3 for M^{-1} IMF and by a factor $\sim 3\text{--}8$ for M^{-2} IMF. Choice of IMF dominates the upper mass-gap percentage of mergers. Nevertheless, it seems most mergers are between the more numerous, lower mass BH in our distributions.

From Table 4, we can see that the median mass ratio for mergers is usually much smaller at the trap ($\tilde{q} \sim 0.1$) than in the bulk of the disc ($\tilde{q} \sim 0.5\text{--}0.7$). We can also see that the median χ_{eff} is modestly larger at the migration trap than in the bulk of the disc mergers. From Section 2.1 above, we expect that later generations of mergers should tend towards alignment and anti-alignment, and this seems consistent with the results for mergers at the migration trap. Changing the merger condition within the Hill sphere so that the two most massive BH merge, made no significant difference to results for the trap. Comparing Tables 3 and 4, we also find the mass of an incoming black hole arriving at the migration trap is generally larger than the median mass of black holes merging in the bulk of the disc. This is because more massive black holes migrate more rapidly than less massive BH (equation 3) so we expect that more massive BH should arrive at the migration trap first. The results from Table 4 seem inconsistent at the $\sim 2\sigma$ level with the results observed with LIGO so far. Future results from O3 may therefore

Table 4. Large-scale runs to test parameter distributions for mergers at the migration trap (where it exists). Column 1 lists the run from Table 1, which is then run 10^2 times. Column 2 lists the median mass ratio (\tilde{q}) from mergers and the associated standard deviation in the merger distribution. If the standard deviation is greater than the lower bound for the run, we use the lower bound for the run instead. Column 3 lists the median χ_{eff} and standard deviation for the mergers. Column 4 lists the full range of χ_{eff} for the run. Column 5 lists the median mass of black holes involved in mergers at the trap together with the standard deviation. Column 6 lists the maximum mass in mergers at the migration trap. Column 7 lists the median masses of the incoming BH merging with the mass at the trap, with the standard deviation of that distribution. If the standard deviation is greater than the difference between the median mass and the minimum mass in the run, we list the latter instead as the lower limit. Column 8 lists the total number of mergers at the migration trap for the run.

Run	\tilde{q}	$\tilde{\chi}_{\text{eff}}$	$[\chi_{\text{min}}, \chi_{\text{max}}]$	\tilde{M}_{trap} (M_{\odot})	M_{max} (M_{\odot})	\tilde{M}_{merge} (M_{\odot})	N_{mergers}
R1	$0.08^{+0.19}_{-0.07}$	0.01 ± 0.52	$[-0.96, 0.98]$	263.6 ± 157.7	1255.8	$19.2^{+16.0}_{-14.2}$	2333
R2	$0.09^{+0.18}_{-0.08}$	-0.06 ± 0.51	$[-0.96, 0.96]$	262.6 ± 141.5	677.0	$19.2^{+14.7}_{-14.2}$	2349
R3	$0.15^{+0.21}_{-0.14}$	-0.05 ± 0.48	$[-0.97, 0.96]$	164.5 ± 90.2	493.5	19.5 ± 13.9	1368
R4	$0.12^{+0.21}_{-0.11}$	0.06 ± 0.52	$[-0.96, 0.96]$	119.1 ± 65.9	330.0	$11.1^{+10.2}_{-6.1}$	1494
R5	$0.02^{+0.09}_{-0.01}$	0.00 ± 0.59	$[-0.98, 0.98]$	996.3 ± 625.7	2469.1	$12.1^{+11.8}_{-7.1}$	15 015
R6	$0.16^{+0.21}_{-0.15}$	0.04 ± 0.48	$[-0.95, 0.95]$	62.1 ± 33.8	168.6	$8.0^{+3.2}_{-3.0}$	1409
R7	$0.13^{+0.22}_{-0.12}$	0.01 ± 0.38	$[-0.76, 0.79]$	116.3 ± 71.8	418.3	$11.1^{+10.6}_{-6.1}$	1487
R8	0.29 ± 0.26	0.09 ± 0.52	$[-0.90, 0.97]$	58.1 ± 34.1	169.1	$10.1^{+9.7}_{-5.1}$	502
R9	$0.14^{+0.20}_{-0.13}$	0.13 ± 0.62	$[-0.95, 0.98]$	169.7 ± 138.0	920.6	$19.5^{+17.3}_{-14.5}$	1852
R10	$0.15^{+0.19}_{-0.14}$	-0.02 ± 0.62	$[-0.94, 0.98]$	154.5 ± 125.1	610.0	$18.7^{+15.4}_{-13.7}$	1819
R12	$0.01^{+0.11}_{-0.01}$	0.01 ± 0.54	$[-0.98, 0.98]$	782.3 ± 473.9	2194.7	$9.2^{+11.4}_{-4.2}$	14 810




Figure 10. Normalized χ_{eff} distribution for R1 carried out 10^2 times with different numerical seeds. The black solid curve shows the distribution of mergers in the bulk of the disc. The red solid curve shows the distribution of mergers at the trap. Approximately 90 per cent of mergers in the bulk occur between the vertical red dashed lines.

rule out the presence of migration traps in AGN discs. If this is the case then either AGN discs do not possess steep changes in aspect ratio or density in the inner disc, or AGN discs are generally short-lived ($\leq 10^5$ yr) unstable configurations, nested within a fuel reservoir or torus.

Fig. 10 shows the total normalized χ_{eff} distribution from the R1 Monte Carlo run. 90 per cent of the bulk mergers have $|\chi_{\text{eff}}| < 0.6$, but the trap distribution seems bimodal, driven by spin flipping of mergers at the trap due to random binary orbital angular momentum. By contrast, R7 is the one run that shows a consistently narrower χ_{eff} distribution and indeed in Fig. 11 we see a far narrower distribution of χ_{eff} among mergers in the bulk of the disc. The trap distribution remains bimodal, but the peaks are also narrower than in Fig. 10. In Fig. 11 ~ 95 per cent of the bulk mergers occur with $|\chi_{\text{eff}}| < 0.2$




Figure 11. As Fig. 10 except for R7. Approximately 95 per cent of mergers in the bulk of the AGN disc occur between the vertical red dashed lines.

(between the vertical red dashed lines). As we can see, a narrower spin input distribution (e.g. from a narrow natal spin distribution) clearly leads to a narrow χ_{eff} distribution in this channel.

5 IMPLICATIONS FOR O3 AND BEYOND

From Section 4.3 there are several constraints that we hope O3 will provide for this channel. First, a continuing build-up of the χ_{eff} distribution for binary black hole mergers will allow us to restrict the allowed black hole IMF ($M^{-\gamma}$, M_{min} , M_{max}) in galactic nuclei and the allowed initial spin magnitude distribution. Secondly, a build-up of the merger mass ratio (q) distribution and a merging mass upper bound allows us to constrain both the IMF and the possibility of migration traps in AGN discs. This latter constraint, if real, will tell us either: there are no steep gradients in density or aspect ratio in AGN discs, or that AGN discs are typically lower density than

in a Sirko & Goodman (2003) model, or that AGN discs are very short-lived ($\leq 10^5$ yr) instabilities in a nuclear accretion flow.

We expect that the merger rate from this channel should increase with redshift (z) out to $z \sim 2$ (McKernan et al. 2018), which is the presumed peak redshift for AGN activity, as well as being the peak redshift for (massive) star formation. This expectation is directly testable via LIGO observations. In tandem with LIGO observations, in the electromagnetic band, we can search for IMBH build up in AGN discs as well as EMRIs embedded in AGN discs by searching for Hill sphere shocks (McKernan et al. 2019), periodic wobbles, or short-lived ripples in the broad component of the FeK α line (McKernan et al. 2013, 2015) or in analogous oscillations in blurred soft X-ray features. Constraints on such effects will allow us to test the distribution of embedded masses in AGN discs as well as the gas torques that we expect to operate there.

6 GRAVITATIONAL WAVE PHENOMENOLOGY FOR THE AGN-DISC CHANNEL

As our understanding of this channel improves and as observations rapidly accumulate, we expect that direct comparisons between evolutionary models and observations will allow us to place sharp constraints on the underlying physics involved. At present, however, we adopt a phenomenological approach for the time-averaged mass and spin distribution of merging BHs formed in AGN discs.

Figs 12 and 13 motivate our phenomenology. These figures show the (cumulative) distribution of masses m_1 and m_2 due to mergers in the bulk and in the trap, averaged over time. All mass distributions are well-approximated by simple power laws. In the trap, the distribution of primary masses m_1 is dominated by IMBHs, which grow linearly with time in the trap up to a large (and physics-dependent) cutoff and which therefore have a uniform time-averaged mass distribution. The secondary mass distribution, by contrast, is a power law bounded by the same mass limits as the input BH population. The relatively stable upper mass limit is striking and important: despite some growth, and despite mechanisms that modify the power law to more strongly favour higher mass secondaries m_2 in the trap, the secondary masses in BBH trap mergers do not significantly extend to higher masses. In the bulk, the distribution of primary and secondary masses are dominated by the input population; for example, when the input mass distribution has $\gamma = 2$, the primary mass distribution also a power law $\gamma_{\text{eff}} \simeq 2$. Migration modifies the mass distribution for the secondary population, which has an effective power law $\gamma_{\text{eff}} \simeq 4$.

This phenomenology does not illustrate a notable signature of 1g or 2g components assembled in the bulk of the disc. While our previous calculations demonstrate such mergers occur, when viewed via GW alone averaged over time and without a priori knowledge of the IMF in the disc, we do not have a compelling *phenomenological* reason to include them as a separate population. Instead, for the immediate future and given the sparsity of the current observational sample, we will model the joint bulk and 1g/2g population into a single component, with a common power law with variable exponent and upper limit.

Motivated by our calculations, we adopt a simple two-component mass model, superposing phenomenological approximations to the distribution of ‘trap’ mergers and ‘bulk’ mergers. We treat these components as two completely independent power-law distributions, not imposing any consistency requirements on the relative rates or power laws aside from the requirement that they share




Figure 12. *Top panel* (trap) shows the CDF for m_1 in the trap. The dashed lines show an approximately uniform distribution $P(< m_1) \propto (m_1 - m_{\min})/(m_{\max} - m_{\min})$ for suitable m_{\max} , m_{\min} . *Bottom panel* (bulk) shows the CDF for m_1 in the bulk. The dashed lines show an approximation of a uniform distribution in $\log m_1$ (i.e. as if $\gamma = 1$).

a common maximum mass $m_{\text{bh, max}}$ for the progenitor population. Specifically, we use the following merger rate model:

$$\rho(m_1, m_2) = \frac{dN}{dV dt} = \rho_{\text{bulk}} + \rho_{\text{trap}} \quad (8)$$

$$\rho_{\text{bulk}} = \mathcal{R}_{\text{bulk}} \mathcal{P}(m_1 | \alpha_{\text{bulk}}, m_{\min}, m_{\max, \star}) \mathcal{P}(m_2 | \alpha_{\text{bulk}}, m_{\min}, m_{\max, \star}) \quad (9)$$

$$\rho_{\text{trap}} = \mathcal{R}_{\text{trap}} \mathcal{U}(m_1 | m_{\min}, m_{\max, \text{IMBH}}) \mathcal{P}(m_2 | \alpha_{\text{trap}}, m_{\min}, m_{\max, \star}), \quad (10)$$

where $\mathcal{U}(x | x_{\min}, x_{\max})$ denotes a uniform distribution in x , on the interval $[x_{\min}, x_{\max}]$, and $\mathcal{P}(x | \alpha, x_{\min}, x_{\max})$ denotes a power-law distribution proportional to $x^{-\alpha}$, on the interval $[x_{\min}, x_{\max}]$. In all runs, we set the minimum black hole mass to $m_{\min} = 5 M_{\odot}$ and allow the maximum IMBH mass $m_{\max, \text{IMBH}}$ to vary freely, with a uniform-in-log prior. We then have one run with the stellar-mass black hole mass fixed to $m_{\max, \star} = 50 M_{\odot}$, to match our simulations, as well as an additional run where we allow it to vary freely, to account for uncertainties in our simulations, with a uniform-in-log prior. Notably, the ‘bulk’ component of this model closely resembles power-law mass distributions previously used to interpret GW source populations; see e.g. Abbott et al. (2016b), Wysocki et al. (2019). For simplicity, because our BH spins distributions are broad

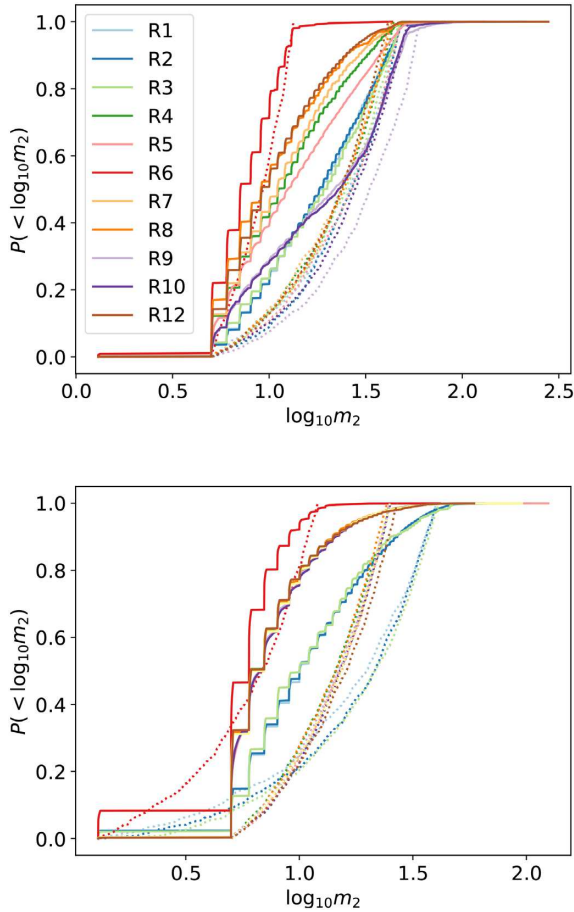


Figure 13. *Top panel* (trap) shows the CDF for m_2 in the trap. The dashed lines show an approximation with $\gamma_{\text{eff}} = 1$. *Bottom panel* (bulk) shows the CDF for m_2 in the bulk. The dashed lines show an approximation with $\gamma_{\text{eff}} = 4$. Note that in all cases, m_2 is almost always bounded by the maximum mass in the progenitor population – in other words, essentially no second-generation objects are secondary masses m_2 .

and do not show strong imprints of hierarchical growth, we adopt a fiducial spin distribution in our analysis and recovery, effectively ignoring the prospects for using spin to distinguish between different formation options.

We compare this phenomenological model to synthetic and real GW observations using the POPMODELS code from Wysocki et al. (2019). This code infers source populations, given observations, and a model for detector sensitivity. Unless otherwise noted, we will use a simplified model for detector sensitivity that accounts for BH mass and (non-precessing) spin, based on a single interferometer’s sensitivity to non-precessing black holes and assuming an analysis minimum frequency of $f_{\text{min}} = 10$ Hz.

A direct observational constraint on AGN disc BBH formation, Fig. 14 shows our inferences about phenomenological AGN disc population parameters given current reported GW observations to date. In the absence of reported IMBH observations which could distinctively identify a trap component, our conclusion is an upper limits on the AGN disc component of $161 \text{ Gpc}^{-3} \text{ yr}^{-1}$ from the trap (alone) and $143 \text{ Gpc}^{-3} \text{ yr}^{-1}$ from the bulk disc (alone); see Table 5.

To assess the prospects for constraining AGN disc formation in the near future, we generate several synthetic binary black hole populations drawn from our simulations and from our phenomenological models, and apply the same procedure. To better understand

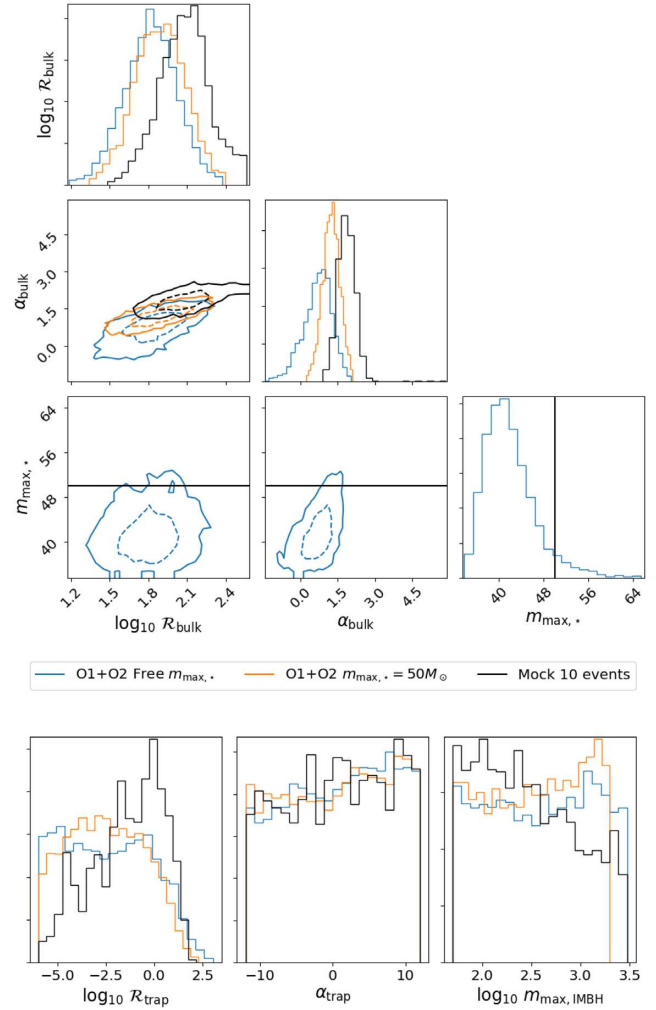


Figure 14. Inferences about a phenomenological AGN disc population, using GW observations available in Abbott et al. (2018), as well as mock data from a representative realization of this model. Each panel shows either 1D or 2D marginal distributions for our models’ phenomenological parameters: the merger rates R_{bulk} , R_{trap} ; the power-law index α in the bulk and trap; and the upper mass limits $m_{\text{max},*}$ for stellar- and intermediate-mass BHs. See Table 5 for credible intervals for these parameters.

our results, the top panel of Fig. 15 shows the interferometer network sensitive volume as a function of total mass for binaries with $m_2 = 50 M_{\odot}$ versus m_1 , for a few choices of BH spin. While network sensitivity increases rapidly at low mass, eventually as the primary black hole mass grows larger, the associated GW merger signal is produced at or below the minimum sensitive frequency viable for analysis at present. As shown in the bottom panel of Fig. 15, the strong mass dependence of network sensitivity influences the *detection-weighted* mass distribution (i.e. the mass distribution of observed BHs), so much so that the largest IMBH binaries produced in each population are not reliably observed. In other words, because of the low-frequency sensitivity of currently operating GW detector networks, the most massive IMBHs $m_1 > 300 M_{\odot}$ by this channel are hard to find when paired with stellar-mass companions as predicted here. For this reason, constraints on the maximum IMBH mass are expected to be weak, unless that mass limit falls below $300 M_{\odot}$.

Within this phenomenological approach, only by finding IMBHs can we identify the unique signature of the trap component and

Table 5. Inferences about our AGN disc phenomenological model, given reported observations in Abbott et al. (2019) and illustrated in Fig. 5. Limits are expressed as 90 per cent credible intervals on each population hyperparameter, except for $m_{\text{max, IMBH}}$ which is a 95 per cent lower limit. All rates are in $\text{Gpc}^{-3} \text{yr}^{-1}$, and all masses in M_{\odot} .

Model	$\mathcal{R}_{\text{trap}}$		α_{trap}		$\mathcal{R}_{\text{bulk}}$		α_{bulk}		$m_{\text{max, *}}$		$m_{\text{max, IMBH}}$
Fixed	0.00	4.95	−10.82	10.90	34.91	165.71	0.61	1.81	—	—	60.17
Free	0.00	14.02	−10.64	10.92	28.77	143.09	−0.40	1.56	36.07	51.39	60.01




Figure 15. *Top panel:* Plot of network sensitivity VT versus m_1 , assuming $m_2 = 50 M_{\odot}$ (blue) or $30 M_{\odot}$ (red), and a low-frequency analysis cutoff of $f_{\text{min}} = 10 \text{ Hz}$. The solid lines show zero BH spin; dotted lines assume the IMBH has spin $a_1 = 0.8$, and the merger occurs in a prograde orbit. Due to limited low-frequency sensitivity, GW detectors may have strongly suppressed sensitivity to the most high-mass and asymmetric sources produced in AGN migration traps, unless the IMBHs are spinning rapidly. *Bottom panel:* Detection-weighted cumulative distribution of m_1 for BBH mergers in the trap, accounting for BH spins, to contrast to Fig. 12. Due to limited sensitivity to $m_1 > 300 M_{\odot}$ indicated in the top panel, the IMBH-BH mergers observed in GW in the near future will not always probe the most massive IMBHs produced in AGN discs, without lower frequency analysis.

thus measure $\mathcal{R}_{\text{trap}}$. As illustrated by the detection volume and detection-weighted mass distributions in Fig. 15, however, ground-based sensitivity to very massive IMBHs in asymmetric and often high mass ratio binaries is relatively small. The ratio of detected binaries in the trap versus the bulk is proportional to their detection-volume-weighted merger rates, with a ratio $F \simeq V(m_2 \simeq 50, m_1) \mathcal{R}_{\text{trap}} / [V(m_1, m_2 \simeq 30 M_{\odot}) \mathcal{R}_{\text{bulk}}]$, where $m_1 \simeq 200 M_{\odot}$ is a typical detectable IMBH mass. The ratio will be further reduced

by approximately a factor $\mathcal{G} = 300 M_{\odot} / m_{\text{IMBH, max}}$, to account for the non-detectability of very massive high- q IMBH binaries; for our reference population, with $m_{\text{IMBH, max}} = 2500$, this factor is $\mathcal{G} \simeq 0.1$. If GW observatories continue to exclude IMBHs, then $\mathcal{R}_{\text{trap}} / \mathcal{R}_{\text{bulk}}$ is constrained by $F \lesssim 1 / N_{\text{obs}} \mathcal{G}$. To illustrate this plausible scenario, we generated a synthetic population with $\mathcal{R}_{\text{bulk}} = 150 \text{ Gpc}^{-3} \text{yr}^{-1}$ and $\mathcal{R}_{\text{trap}} = 30 \text{ Gpc}^{-3} \text{yr}^{-1}$, with the mass distribution of the bulk BBH population drawn from previously reported inferences about observational results. For simplicity and to be pessimistic, we performed our analysis assuming BH spins do not impact detector network sensitivity. In this synthetic scenario, no binaries with an IMBH are found within the first 95 observations, as they make up $\mathcal{O}(1 \text{ per cent})$ of detections under this model. Therefore our analyses of these events (or any smaller subset) only recover the properties of the injected bulk population, similar to Fig. 14. Because many of the synthetic trap IMBH binaries are very massive, our constraints on the overall merger rate are weak.

7 CONCLUSIONS

As LIGO identifies the distribution of BHs involved in mergers in our Universe, we gain increasingly strong astrophysical constraints on AGN discs. The merger of BHs in AGN discs directly probes physical parameters like gas density, gas distribution, and disc lifetime, which are hard to extract from electromagnetic observations. For example, LIGO is presently telling us that low ionization nuclear emission regions (LINERs) cannot be mostly optically thick advection-dominated accretion flows (McKernan et al. 2018; Ford & McKernan 2019). Future LIGO population statistics will constrain these astrophysical parameters directly and LIGO’s results over the next decade will have broad astrophysical implications for models of Λ CDM feedback, AGN formation, fuelling, and extinction.

In this work we provide the mass and χ_{eff} distributions under various assumptions about the physics involved in this compact binary formation channel. For the bulk of the disc, the mass ratio distribution among merging binaries is centred around $\tilde{q} \sim 0.5 - 0.7$. Even if we adopt a black hole IMF consistent with Galactic BH masses ($5 - 15 M_{\odot}$), we also find a modestly asymmetric population $\tilde{q} \sim 0.7$. The migration trap will steadily grow an IMBH, whose maximum mass reflects the disc lifetime and natal BH population. Mergers of BH binaries in that trap are much more asymmetric than mergers in the bulk of the disc, with typical mass ratios $\tilde{q} \sim 0.1$. Thus, ongoing observational constraints on the high-mass BH population from LIGO put limits on AGN disc lifetime and disc structure (in particular the existence of migration traps) as well as orbital grind-down efficiency. Conversely, ongoing LIGO non-detections of black holes $> 10^2 M_{\odot}$ puts strong limits on the presence of migration traps in AGN discs and the AGN disc lifetime.

The AGN disc channel can impart a distinctive spin distribution to the IMBH population. That said, for the majority of mergers, the spin distribution reflects assumptions about the natal spin distribution, convolved with the spin alignment physics discussed in the text. In our calculations, the χ_{eff} distribution from BH mergers in the bulk

of the disc is naturally centred around ($\tilde{\chi}_{\text{eff}} \approx 0$). The median of $\tilde{\chi}_{\text{eff}}$ is slightly negative if gas hardening of retrograde binaries is more efficient than for prograde binaries, while the width of the χ_{eff} distribution is significantly narrower for narrow spin magnitude distributions. Inverting from the observed population as LIGO continues to operate, an observed bias towards small positive or negative $\tilde{\chi}_{\text{eff}}$ may allow us to distinguish the efficiency of gas hardening of prograde or retrograde binaries.

The rate of black hole mergers is highest early on (~ 0.1 Myr) in the AGN disc lifetime. The highest merger rate occurs for this channel if AGN discs are relatively short-lived (~ 1 Myr) so multiple AGN episodes can happen per Galactic nucleus in a Hubble time. Finally, we wish to emphasize that LIGO observations going forward will allow us to dramatically constrain models of AGN discs in the next few years.

ACKNOWLEDGEMENTS

BM and KESF are supported by NSF grant 1831412. ROS and DW are supported by NSF PHY-1707965; DW also thanks the RIT COS and CCRG for support. We acknowledge very useful discussions with Bernard Kelly, Cole Miller, Yuri Levin, and Phil Armitage.

REFERENCES

- Aasi J. et al., 2015, *Class. Quantum Gravity*, 32, 074001
 Abbott B. P. et al., 2016a, *ApJ*, 833, L1
 Abbott B. P. et al., 2016b, *Phys. Rev. X*, 6, 1015
 Abbott B. P. et al., 2019, *Phys. Rev. X*, 9, 031040
 Abbott B. P. et al., 2019, *ApJ*, 882, L24
 Acernese F. et al., 2015, *Class. Quantum Gravity*, 32, 024001
 Antonini F., Rasio F., 2016, *ApJ*, 831, 187
 Bartos I. et al., 2017, *ApJ*, 835, 165
 Baruteau C., Cuadra J., Lin D. N. C., 2011, *ApJ*, 726, 28
 Belczynski K. et al., 2010, *ApJ*, 714, 1217
 Bellovary J. et al., 2016, *ApJ*, 819, L17
 Bogdanovic T., Reynolds C. S., Miller M. C., 2007, *ApJ*, 661, L147
 deMink S., Mandel I., 2016, *MNRAS*, 460, 3545
 Fishbach M., Holz D. E., Farr B., 2017, *ApJ*, 840, L24
 Ford K. E. S., McKernan B., 2019, *MNRAS*, 490, L42
 Generozov A., Stone N. C., Metzger B. D., Ostriker J. P., 2018, *MNRAS*, 478, 4030
 Gerosa D., Berti E., 2017, *Phys. Rev. D*, 95, 124046
 Hailey C. J. et al., 2018, *Nature*, 556, 70
 Leigh N. W. C. et al., 2018, *MNRAS*, 474, 5672
 LIGO, VIRGO Scientific Collaborations, 2019, *ApJ*, 882, L24
 McKernan B. et al., 2012, *MNRAS*, 425, 460
 McKernan B. et al., 2013, *MNRAS*, 432, 168
 McKernan B. et al., 2014, *MNRAS*, 441, 900
 McKernan B. et al., 2015, *MNRAS*, 452, L1
 McKernan B. et al., 2018, *ApJ*, 866, 66
 McKernan B. et al., 2019, *ApJ*, 884, L50
 Miralda-Escudé J., Gould A., 2000, *ApJ*, 545, 847
 Morris M., 1993, *ApJ*, 408, 496
 Rodriguez C. et al., 2016, *Phys. Rev. D*, 93, 084029
 Secunda A. et al., 2019, *ApJ*, 878, 85
 Sigurdsson S., Phinney E. S., 1993, *ApJ*, 415, 631
 Sirko E., Goodman J., 2003, *MNRAS*, 341, 501
 Stone N. C. et al., 2017, *ApJ*, 464, 946
 Subr L., Karas V., 2005, *A&A*, 433, 405
 Syer D., Clarke C., Rees M. J., 1991, *MNRAS*, 250, 505
 Thompson T. A., Quataert E., Murray N., 2005, *ApJ*, 630, 167
 Tichy W., Marronetti P., 2008, *Phys. Rev. D*, 78, 1501
 Varma V., Gerosa D., Stein L., Hebert F., Zhang H., 2019, *Phys. Rev. Lett.*, 122, 1101
 Wysocki D., Lange J., O’Shaughnessy R., 2019, *Phys. Rev. D*, 100, 043012
 Yang Y. et al., 2019, *Phys. Rev. Lett.*, 123, 1101

This paper has been typeset from a \LaTeX file prepared by the author.