



Testing the Binary Hypothesis: Pulsar Timing Constraints on Supermassive Black Hole Binary Candidates

Alberto Sesana¹, Zoltán Haiman^{2,3} , Bence Kocsis⁴, and Luke Zoltan Kelley⁵

¹ School of Physics and Astronomy and Institute of Gravitational Wave Astronomy, University of Birmingham, Edgbaston B15 2TT, UK

² Department of Astronomy, Columbia University, 550 West 120th Street, New York, NY 10027, USA

³ Department of Physics, New York University, New York, NY 10003, USA

⁴ Institute of Physics, Eötvös University, Pázmány P. s. 1/A, Budapest, 1117, Hungary

⁵ Harvard University, Center for Astrophysics, Cambridge, MA 02138, USA

Received 2017 April 5; revised 2018 February 1; accepted 2018 February 2; published 2018 March 22

Abstract

The advent of time domain astronomy is revolutionizing our understanding of the universe. Programs such as the Catalina Real-time Transient Survey (CRTS) or the Palomar Transient Factory (PTF) surveyed millions of objects for several years, allowing variability studies on large statistical samples. The inspection of ≈ 250 k quasars in CRTS resulted in a catalog of 111 potentially periodic sources, put forward as supermassive black hole binary (SMBHB) candidates. A similar investigation on PTF data yielded 33 candidates from a sample of ≈ 35 k quasars. Working under the SMBHB hypothesis, we compute the implied SMBHB merger rate and we use it to construct the expected gravitational wave background (GWB) at nano-Hz frequencies, probed by pulsar timing arrays (PTAs). After correcting for incompleteness and assuming virial mass estimates, we find that the GWB implied by the CRTS sample exceeds the current most stringent PTA upper limits by almost an order of magnitude. After further correcting for the implicit bias in virial mass measurements, the implied GWB drops significantly but is still in tension with the most stringent PTA upper limits. Similar results hold for the PTF sample. Bayesian model selection shows that the null hypothesis (whereby the candidates are false positives) is preferred over the binary hypothesis at about 2.3σ and 3.6σ for the CRTS and PTF samples respectively. Although not decisive, our analysis highlights the potential of PTAs as astrophysical probes of individual SMBHB candidates and indicates that the CRTS and PTF samples are likely contaminated by several false positives.

Key words: black hole physics – gravitational waves – quasars: general – surveys

1. Introduction

In the past two decades, it has been realized that supermassive black holes (SMBHs) are a fundamental ingredient of galaxy formation and evolution (see, e.g., Kauffmann & Haehnelt 2000; Croton et al. 2006), and it is now well established that possibly all massive galaxies host an SMBH at their center (see Kormendy & Ho 2013 and references therein). In the standard Λ CDM cosmology, structure formation proceeds in a hierarchical fashion, whereby galaxies frequently merge with each other, progressively growing their mass (White & Rees 1978). Following the merger of two galaxies, the SMBH hosted in their nuclei sink to the center of the merger remnant because of dynamical friction (DF), eventually forming an SMBH binary (SMBHB).

The evolution of SMBHBs was first sketched out by Begelman et al. (1980). After the initial DF phase, the two SMBHBs become bound at parsec scales forming a Keplerian system. At this point, DF ceases to be efficient and in the absence of any other physical mechanism at play, the binary would stall. Because the average massive galaxy suffers more than one major merger in its assembly, in this scenario virtually all galaxies would host parsec scale SMBHBs. Galactic nuclei, however, are densely populated with stars and also contain gas. It has been shown that both three-body ejection of ambient stars (Berczik et al. 2006; Khan et al. 2011; Preto et al. 2011;

Vasiliev et al. 2015; Sesana & Khan 2015), interaction with gaseous clumps (Goicovic et al. 2016) or with a massive circumbinary disk (Escala et al. 2005; MacFadyen & Milosavljević 2008; Cuadra et al. 2009; Hayasaki 2009; Nixon et al. 2011; Roedig et al. 2011, 2012; Shi et al. 2012; Farris et al. 2014; Miranda et al. 2017; Tang et al. 2017; see also the reviews by Dotti et al. 2012; Mayer 2013), or interactions between SMBH triplets (Bonetti et al. 2017; Ryu et al. 2018) may provide efficient ways to shrink SMBHBs down to centiparsec scales, where efficient gravitational wave (GW) emission takes over, leading to final coalescence. Still, binary hardening timescales can be as long as a gigayear (Sesana & Khan 2015; Vasiliev et al. 2015), implying a substantial population of subparsec SMBHB lurking in galactic nuclei (see also Kelley et al. 2017a).

Subparsec SMBHBs are extremely elusive objects (see Dotti et al. 2012; Komossa & Zensus 2016, for recent reviews). At extragalactic distances, their angular size is well below the resolution of any current instrument, making direct imaging impossible except possibly in radio VLBI observations (D’Orazio & Loeb 2017). Conversely, there is an increasing number of detections of SMBH pairs (i.e., SMBHBs that are still not gravitationally bound to each other) at hundred parsec to kiloparsec separations in merging galaxies, which are their natural progenitors (e.g., Hennawi et al. 2010; Comerford et al. 2013). With direct imaging impractical, other avenues to observe subparsec SMBHBs have been pursued, namely spectroscopic identification and time variability. Because of the high ($>1000 \text{ km s}^{-1}$) typical orbital velocities, SMBHBs have tentatively been identified with systems showing



Original content from this work may be used under the terms of the [Creative Commons Attribution 3.0 licence](https://creativecommons.org/licenses/by/3.0/). Any further distribution of this work must maintain attribution to the author(s) and the title of the work, journal citation and DOI.

significant offsets of the broad-line emission lines compared to the reference provided by narrow emission lines, and/or with frequency shifts of the broad lines over time (Tsalmantza et al. 2011; Eracleous et al. 2012; Decarli et al. 2013; Ju et al. 2013; Shen et al. 2013; Runnoe et al. 2017; Wang et al. 2017). The latter, in fact, are generated within the host galaxies at hundreds of parsecs from the central SMBHs, whereas the former are generated by gas bound to the SMBHs. If the SMBHs have a significant velocity compared to the galaxy rest frame, and the broad emission region is bound to the individual SMBHs, then broad lines will have an extra redshift/blueshift compared to their narrow counterparts. Note, however, that for very compact SMBHBs those broad lines might in fact be generated within the circumbinary disk (Lu et al. 2016), questioning this interpretation.

With the advent of time domain astronomy, the identification of SMBHBs via periodic variability has been put forward (Haiman et al. 2009). The rationale for selecting candidates in this way is that if gas is being accreted onto an SMBHB, the orbital period of the system could translate into periodic variability of the emitted luminosity. In fact, detailed hydrodynamical simulations show that SMBHBs embedded in circumbinary disks carve a cavity in the gas distribution and the gas streams from the cavity edge onto the binary in a periodic fashion (e.g., Artymowicz & Lubow 1996; Hayasaki et al. 2007; MacFadyen & Milosavljević 2008; Cuadra et al. 2009; Roedig et al. 2011; Noble et al. 2012; Shi et al. 2012; D’Orazio et al. 2013; Farris et al. 2014; Shi & Krolik 2015).

Whether such periodic streaming translates into periodicity in the emitted luminosity is less clear. Moreover, it has been pointed out that such periodicity would mostly impact the UV and X-ray emission from the binary, whereas the optical emission coming from the circumbinary disk might be relatively steady (Sesana et al. 2012; Farris et al. 2015), except for the most massive ($M \gtrsim 10^9 M_\odot$) SMBHBs for which the optical emission arises from gas bound to the individual black holes (D’Orazio et al. 2015b). Despite these uncertainties, active galactic nuclei (AGNs) showing periodic variability are viable candidates for hosting SMBHBs. Based on this hypothesis, Graham et al. (2015; hereafter G15) proposed 111 SMBHB candidates by inspecting the light curves of 243,486 quasars identified in the Catalina Real-time Transient Survey (CRTS, Drake et al. 2009). In a similar effort, Charisi et al. (2016; hereafter C16) identified 33 SMBHB candidates, among 35,383 spectroscopically confirmed quasars in the Palomar Transient Factory (PTF, Rau et al. 2009) survey, with somewhat fainter magnitudes and shorter periods than G15. Both groups presented a thorough analysis demonstrating a large excess of periodic sources at odds with the expectations of standard AGN-variability models.

The statistical significance of the detected periodicity depends strongly on the stochastic noise model for underlying quasar variability. For example, Vaughan et al. (2016) have recently shown that Gaussian red noise models can naturally lead to frequent false positives in periodicity searches, especially at inferred periods comparable to the length of the data stream. At the same time, observations are expected to yield binaries preferentially at lower frequencies where they spend the largest fraction of their lifetimes. On the other hand, pure red noise (i.e., a single f^{-2} power-law noise power spectrum) is a poor description of quasar variability in general.

A better fit to the power spectra observed for large quasar samples is the damped-random walk (DRW) noise model, in which the red noise power spectrum flattens to white noise ($\propto f^0$) at low frequencies (e.g., MacLeod et al. 2010). Adopting a DRW reduces the stochastic noise power at low frequencies, and therefore translates to a higher significance of an observed periodicity. As a result, the significance of the inferred periodicities depend strongly on the poorly constrained underlying noise model (see a further discussion of this point in, e.g., Kelly et al. 2014; Charisi et al. 2016; Vaughan et al. 2016; Kozłowski 2017).

In this paper, we test the SMBHB hypothesis on physical grounds. SMBHBs are powerful emitters of nHz GWs, which are currently probed by pulsar timing arrays (PTAs, Foster & Backer 1990). There are three “regional” PTAs currently in operation: the European Pulsar Timing Array (EPTA, Desvignes et al. 2016), the Australian Parkes Pulsar Timing Array (PPTA, Reardon et al. 2016) and the North American Nanohertz Observatory for Gravitational Waves (NANOGrav, Arzoumanian et al. 2015). These three collaborations share data under the aegis of the International Pulsar Timing Array (IPTA; Verbiest et al. 2016), with the goal of improving their combined sensitivity. The collective signal from a cosmic population of SMBHBs results in a stochastic GW background (GWB, Rajagopal & Romani 1995; Jaffe & Backer 2003; Wyithe & Loeb 2003; Sesana et al. 2008), but particularly massive/nearby systems might emit loud signals individually resolvable above the background level (Sesana et al. 2009).

Recent PTA efforts resulted in several upper limits both for a GWB (Lentati et al. 2015; Shannon et al. 2015; Arzoumanian et al. 2016) and for individual sources (Arzoumanian et al. 2014; Zhu et al. 2014; Babak et al. 2016). Both G15 and C16 showed that each SMBHB candidate identified in CRTS or in PTF is individually compatible with current PTA upper limits. This is partly because timing residuals from those individual SMBHBs are simply too small and partly because their GW emission lies at frequencies above $>10^{-8}$ Hz, which is above the PTA detection “sweet spot” at $\approx 10^{-8.3}$ Hz. The 10^{-8} Hz lower limit to the GW frequency, however, is only a selection effect due to the length of the CRTS and PTF data stream (9 years and 4 years respectively).

Here we show that, when properly converted into an SMBHB merger rate and extrapolated to lower frequencies, the CRTS and PTF samples are in tension with current PTA measurements. This is particularly true when virial mass estimates of the candidates are taken at face value; in this case, both the CRTS and the PTF sample are severely inconsistent with PTA upper limits. Virial mass estimates are, however, known to be biased high (Shen et al. 2008). Correcting for this bias alleviates the inconsistency of the samples with PTA data, but tension persists at the $>2\sigma$ level for both. This indicates that both the CRTS and PTF samples are contaminated by several false positives, whose light-curve variability therefore must have a different physical origin. We show that our conclusions are not severely affected by physical processes potentially capable of suppressing the low-frequency GW signal, such as significant eccentricities, or strong environmental coupling.

This paper is organized as follows. We concentrate our investigation on the G15 SMBHB candidate sample, which is presented in Section 2 and used in Section 3 to reconstruct the merger rate of SMBHBs throughout cosmic history. The

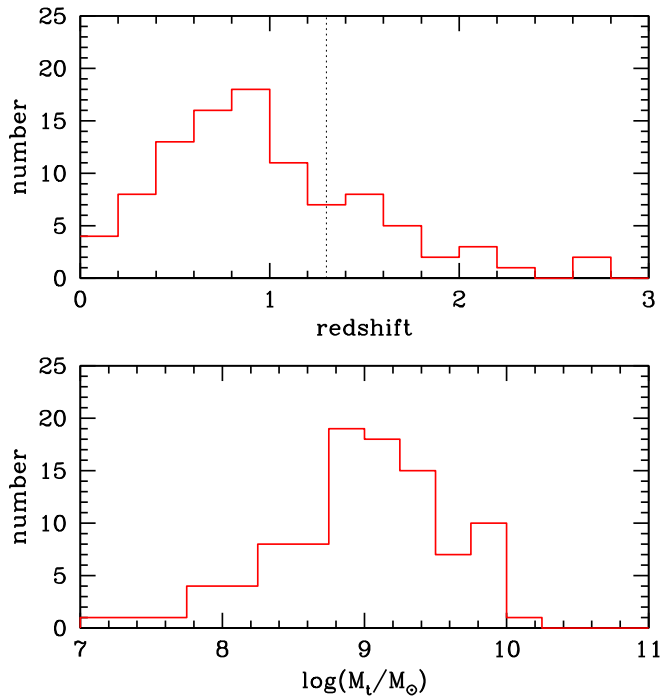


Figure 1. Redshift (upper panel) and total mass (lower panel) distribution of the SMBHB candidates identified by Graham et al. (2015) in the CRTS. Out of the 111 candidates, only the 98 with reported black hole mass estimates are shown. The dotted vertical line in the upper panel at $z = 1.3$ marks the higher redshift considered in SMBHB population models by Sesana (2013b).

implied GWB is computed in Section 4 and compared to current PTA upper limits. In Section 5, we apply the same formalism to the C16 sample. The main results are summarized in Section 6. Throughout this paper, we adopt a concordance Λ CDM cosmology with $h = 0.679$, $\Omega_M = 0.306$, and $\Omega_\Lambda = 0.694$ (Planck Collaboration XIII 2016).

2. The Catalina Survey SMBHB Candidate Sample

CRTS is a time-domain survey periodically scanning 33,000 deg^2 (80% of the whole sky). The data release used by G15 contains the light curves of millions of objects monitored over nine years. Objects have been cross-checked with the 1M quasar catalog⁶ to identify more than 300 k spectroscopically confirmed quasars, of which 243,486 have sufficient light-curve coverage for a periodicity search. Among these sources, 111 have been flagged as periodically varying and have been proposed as potential SMBHB candidates. For most of these systems, the values of the estimated SMBHB mass, redshift, and orbital period is provided by G15 in tabular form.

2.1. Observational Properties and Intrinsic Mass Estimates

Of the 111 candidates presented by G15, we consider only the 98 with a reported measurement of the mass. In the following, we will conservatively assume that this is the total SMBHB mass ($M_t = M_1 + M_2$). The mass and redshift distributions of these 98 systems, as reported in G15, are plotted in Figure 1. Interestingly, the number of sources peaks at $z \approx 1$, rapidly declining to zero at $z > 2$. In the binary hypothesis, this can be explained by a selection effect. G15’s search is sensitive to observed periods between 20 and 300 weeks

and has a fixed magnitude limit. As a result, at higher redshifts, they could only find binaries with shorter rest-frame orbital periods (< 2 years at $z > 2$), and with increasingly large masses. These massive, compact binaries have decoupled from their circumbinary disks, and are likely well inside the GW-driven inspiral regime (Haiman et al. 2009), with very short ($\lesssim 10^4$ years) inspiral times. As a result, they will be exceedingly rare. Additionally, binaries can take several gigayears to overcome the “final parsec problem,” so that the distribution of SMBHBs might peak at lower redshift with respect to the quasar luminosity function.

In performing our analysis, a correct estimate of the mass of the systems is of paramount importance. The values of M_t reported in G15 and plotted in Figure 1 are taken from the catalog compiled by Shen et al. (2008), which provides “single epoch mass measurements” using virial BH mass estimators based on the luminosities and the $H\beta$, MgII , and CIV emission lines. We will refer to those mass estimates as virial masses. Let us consider an object with true mass M , with a virial mass estimate M_e and let us define $m = \log M$, and $m_e = \log M_e$. Let us further make the long-standing simplifying assumption (e.g., Eddington 1913) that the virial mass estimator, at fixed true mass, follows the log-normal distribution,

$$p(m_e|m) = \frac{1}{\sqrt{2\pi\sigma^2}} e^{-\frac{(m-m_e)^2}{2\sigma^2}}, \quad (1)$$

where $\sigma = 0.3$ dex is the measured intrinsic scatter. Equation (1) describes the probability of measuring a mass M_e given a true mass M . We are, however, interested in an estimate of the true mass M , which can be obtained using Bayes theorem

$$p(m|m_e) \propto p(m_e|m)p(m), \quad (2)$$

where $p(m)$ is the prior distribution of the true masses and we omitted a constant normalization factor. In our case, $p(m)$ represents the prior knowledge of the SMBH mass function $\phi(m)$, leading to the product

$$p(m|m_e) \propto \phi(m) e^{-\frac{(m-m_e)^2}{2\sigma^2}}. \quad (3)$$

For a bottom-heavy mass function, this leads to a Malmquist-type bias in the estimate of the true mass M , i.e., $\langle m \rangle < \langle m_e \rangle$ (e.g., Lynden-Bell et al. 1988). In our calculation, we consider the SMBH mass function obtained by Hopkins et al. (2007). We note that the normalization of the mass function is irrelevant; our results depend only on the shape of the mass function. The key factor is that $\phi(m)$ is generally described as a broken power law, with a steep slope $\alpha \approx -2.5$ at the high mass end. In that mass range, when combined with $\phi(m)$, a measured m_e implies an underlying true mass m that is approximately 0.5 dex smaller (see Shen et al. 2008). This is of capital importance, because it means that at the high mass end, the M_t values reported in G15 are biased high by a factor of ≈ 3 . This is illustrated in Figure 2, where the mass function at different redshifts is plotted along with the median bias $\Delta \log(M)$ computed by taking the difference between the virial masses reported in G15 and the median of the true mass distribution obtained via Equation (3). We checked that Equation (3) implies a bias of 0.55 dex when combined with a single power law with $\alpha \approx -2.6$, consistent with the calculation in Shen et al. (2008).

⁶ <http://quasars.org/milliquas.htm>

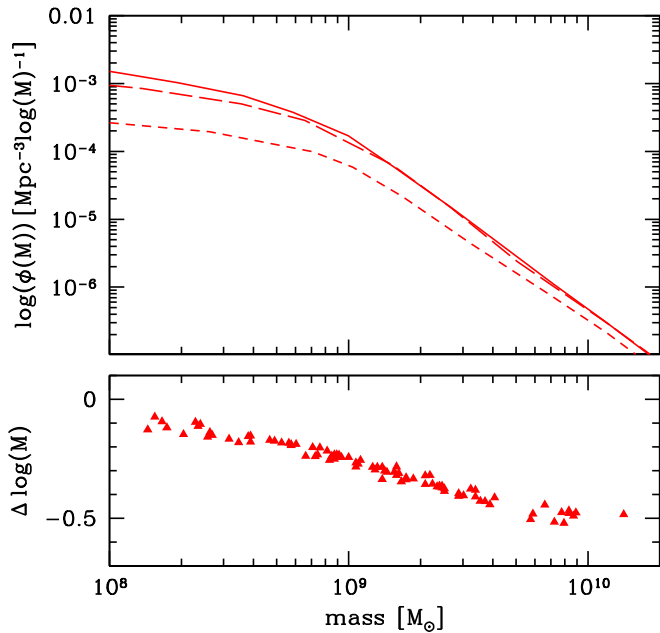


Figure 2. Top panel: SMBH mass function derived by Hopkins et al. (2007) at $z = 0$ (solid), $z = 1$ (long-dashed), and $z = 2$ (short-dashed). Bottom panel: median bias of the virial masses of all the candidates reported in G15 (see the main text for a full description).

We also checked that our results are robust against the specific $\phi(m)$ choice: using mass functions derived by Shen & Kelly (2012) and Kelly et al. (2009) yields quantitatively similar results. This is not surprising, since all mass functions share a consistent steep decay at the high mass end. In the following, we therefore consider two models:

1. **Model True:** “true” masses are generated by drawing randomly from the probability distribution in Equation (3), where $\phi(m)$ is taken from Hopkins et al. (2007). This is our fiducial model.
2. **Model Vir:** we assume that virial masses given by Shen et al. (2008) and reported in G15 are unbiased and we apply a further uncertainty of $\sigma = 0.3$ dex to the measured value.

The second scenario is included for two reasons. First, comparing the “True” and the “Vir” models allows us to quantify the importance of correcting for a systematic bias in the virial mass estimates. This correction depends on the assumed shape of the intrinsic mass function and is often neglected in the literature. Second, the approach in the “True” models further assumes that the scatter observed in the virial mass estimates, measured in practice at fixed luminosity (L_λ) and line width (FWHM of broad lines), represents the scatter at a fixed true mass. This is the most common interpretation, which leads to Equation (3) (see, e.g., Shen et al. 2008). However, we note that virial mass estimates M_e are calibrated against masses M_r determined from reverberation mapping for a subset of quasars (e.g., Peterson 2014). If one makes the extreme assumption that the reverberation masses are indeed the true masses, with no errors, then the scatter observed in the 2D plane of (M_e, M_r) , measured with respect to the line $M_e = M_r$, can be interpreted as the scatter in the true mass M_r at

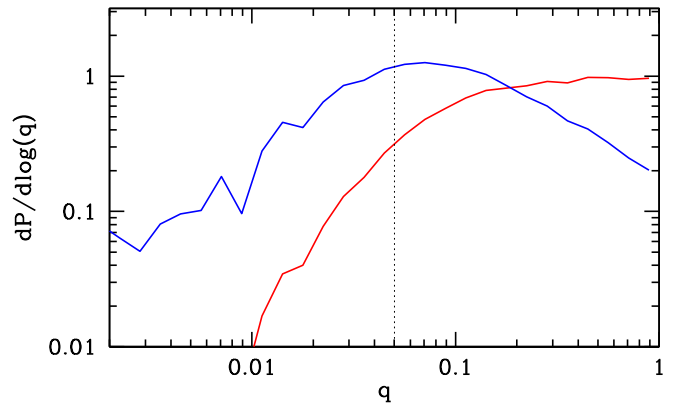


Figure 3. Probability distribution of the mass ratio of merging SMBHBs in the Millennium Run considering only systems with $M_r > 3 \times 10^8 M_\odot$ merging at $z < 2$ (red line) compared to the q probability distribution of an SMBHB observed in a selected frequency range (blue line). The vertical dotted line marks $q = 0.05$, above which accretion-induced periodicity is seen in numerical simulations of SMBHBs in circumbinary disks (see the text for details).

fixed M_e . In this extreme limiting case, the “Vir” models would give the correct probability distribution of the true masses.

2.2. Assigning Individual SMBH Masses

The procedure detailed above provides only an estimate of the true total mass of each individual SMBHB candidate. To compute the associated GW signal, the mass ratio $q = M_2/M_1$ (where by definition $M_1 > M_2$) of the sources is also needed. Inferring the q distribution of observable SMBHBs is far from being a trivial task, and we will see that it is an important factor in assessing PTA constraints. We describe in the following some of the subtleties that come into play, which led us to consider three different scenarios.

In comparing their sample with PTA upper limits, G15 assumed a typical $q = 1/2$. Although this might appear to be a reasonable choice, galaxy formation models predict indeed a larger range of q , as shown in Figure 3. Here we consider all the N_m SMBHBs with $M_r > 3 \times 10^8 M_\odot$ merging at $z < 2$ in a modified version of the galaxy formation models of Guo et al. (2013) implemented in the Millennium Simulation (Springel et al. 2005). Modifications included populating merging galaxies with SMBHBs according to a specific scaling relation (in this case, the M – σ relation from Gültekin et al. 2009), and adding appropriate delays between galaxy mergers and SMBHB mergers to account for DF and SMBHB hardening. Plotted is the probability distribution $\mathcal{P}(\log q)$ (normalized so that $\int \mathcal{P}(\log q) d \log q = 1$) of those merging systems. The distribution is essentially flat down to $q = 0.1$, and drops at lower mass ratios. We see, however, that mergers down to $q \approx 0.01$ are still possible. Similar distributions, albeit sometimes flatter and extending to even lower q , have been found in semi-analytic models (e.g., Barausse 2012) and are produced by full cosmological hydrodynamical simulations like Illustris (Kelley et al. 2017a).

If we take a snapshot of the sky, the q distribution of SMBHB observed at a given frequency is not the same as that of the merging systems. In fact, if we assume for simplicity that all binaries are circular and purely GW driven, the residence time at a given frequency is proportional to (see,

e.g., Sesana et al. 2008),

$$\frac{dt}{df} = \frac{5}{96\pi^{8/3}} \left(\frac{GM}{c^3} \right)^{-5/3} f^{-11/3}, \quad (4)$$

where $\mathcal{M} = M_1^{3/5} M_2^{3/5} / M_t^{1/5} = M_1 q^{3/5} / (1+q)^{1/5}$ is the binary chirp mass. Therefore, for a given M_1 (and ignoring the $1+q \approx 1$ factor) $dt/df \propto q^{-1}$. Since the number of observable binaries in a given frequency window is $dN/df \propto N_m \times dt/df$, their actual q distribution is proportional to $\mathcal{P}(\log q)/q$. This is shown by the blue line in Figure 3: the q distribution is now skewed toward lower values, peaking at $q \lesssim 0.1$. Longer-lived, low- q binaries should therefore be quite common. Whether they are observable as periodic quasars, however, is less clear. In the standard picture, periodicity in the light curve is associated with variability in the accretion onto an SMBHB as $L(t) \propto \dot{M}(t)$. The typical variability of 0.2–0.5 magnitudes observed in CRTS thus corresponds to accretion fluctuations (and relative luminosity fluxes) of about 20%–60%, necessitating a relatively high q . For example, Farris et al. (2014) and D’Orazio et al. (2016) find that $q > 0.05$ is needed to get a distinctive variability pattern in the accretion rate (which is marked in Figure 3). On the other hand, D’Orazio et al. (2015b) proposed that the sinusoidal behavior of the light curve is given by relativistic Doppler boosting. This explanation is preferred at low-mass ratios ($q \lesssim 0.1$), because low q (i) increases the secondary’s orbital velocity and the amplitude of the Doppler modulation, while (ii) reducing the hydrodynamical variability. Thus, the Doppler variability is complementary to the accretion-induced variability. Besides the technical details of the source of variability, the large accretion rates implied by the CRTS sample of bright quasars are generally found to be triggered by major mergers (e.g., Croton et al. 2006). Therefore, the q distribution of SMBHBs hosted in bright quasars might be biased high.

In light of these uncertainties, we construct three mass ratio models:

1. Model hiq: q drawn from a log-normal distribution with $\sigma_{\log q} = 0.5$ dex, peaked at $\log q = 0$ (and we consider only $\log q < 0$). This model is representative of a biased high q distribution, and is similar to the $q = 1/2$ study case considered in G15.
2. Model fid: q is drawn from the distribution plotted in Figure 3 (blue curve), but with a minimum cutoff placed at $q = 0.05$. This model is inspired by accretion-induced periodicity favoring higher q .
3. Model loq: q is drawn from the distribution plotted in Figure 3 (blue curve), without any cutoff. This model preserves binaries with $q \lesssim 0.1$ and is inspired by the Doppler boosting scenario, favoring lower q .

Coupled with the two mass-models described in the previous section, this gives us a total of six models that we label Model_True_hiq, Model_True_fid, Model_True_loq, Model_Vir_hiq, Model_Vir_fid, Model_Vir_loq. “True” models, that correct for the intrinsic bias in virial mass measurement, are our default models and the following discussion will concentrate on them. Among those, we pick Model_True_fid as our fiducial model, since it combines corrected mass estimates with a mass ratio distribution derived from cosmological simulations coupled with a minimum cut that is motivated by systematic

hydrodynamical simulations. The main limitation of this model is that q is evaluated by considering the SMBH masses before the galaxy merger. It has been shown that in the merger process, this mass ratio is expected to change, with the smaller black hole likely to accrete more (Farris et al. 2014; Capelo et al. 2015). Although this tends to favor larger values of q than those assumed in our “fid” model, quantifying the overall impact on the q distribution is beyond the scope of this study. We therefore stress that the mass ratio distributions listed above, although encompassing a wide range of physically motivated possibilities, are illustrative.

Note that Model_True_loq and Model_Vir_loq imply a moderate inclination with respect to the observer line of sight (otherwise the boosting would not be efficient enough, see details in D’Orazio et al. 2015b). This implies that the CRTS observed sources would be an incomplete sample (including only systems with favorable inclination) of the underlying SMBHB population. We do not attempt to address this incompleteness in the following analysis, which makes our estimates for those two models conservative, and accounts for the possibility of accretion-induced variability at low q (Farris et al. 2014).

3. Building Mock SMBHB Populations

For each of the six models enumerated in the previous section, we take the 98 CRTS candidates with estimated total mass and we assign them M_t and q by drawing from the respective distributions. We repeat the procedure 1000 times, to get a statistical ensemble of SMBHB samples under each of the model prescriptions.

Each individual realization of the 98 CRTS SMBHB candidates can be used to infer an SMBHB merger rate as follows. Neglecting, for the moment, completeness issues, the output of the CRTS can be treated as an all sky population of binaries with the given masses and redshifts, emitting in some frequency range. We note that, because of the limited timespan of the data, CRTS is sensitive only to SMBHBs with observed orbital periods shorter than the threshold value $\tilde{P}_{\text{orb}} = 6$ years. This corresponds to a rest-frame GW frequency $\tilde{f}_r = 2 \times \tilde{P}_{\text{orb}}^{-1} (1+z)$. The rest-frame coalescence timescale of a (circular, GW driven) SMBHB emitting at \tilde{f}_r is (Peters 1964)

$$T_c(\tilde{f}_r) = \frac{5}{256} \left(\frac{GM}{c^3} \right)^{-5/3} (\pi \tilde{f}_r)^{-8/3}. \quad (5)$$

Suppose now we have N identical (in mass and redshift) binaries emitting at frequencies $f_r > \tilde{f}_r$ identified in the CRTS sample. By virtue of the continuity equation for SMBHB evolution, we can convert this number, into a merger rate, using

$$\dot{N} \equiv \frac{dN}{dt_r} = \frac{N(f_r > \tilde{f}_r)}{T_c(\tilde{f}_r)}. \quad (6)$$

We stress once again that $T_c(\tilde{f}_r)$ is the coalescence time at the longest orbital period probed by CRTS (\tilde{P}_{orb}), which is the timescale defining the sample population, and not at the period of each specific binary candidate. We can then generalize the argument to a distribution of SMBHBs with different masses and redshifts and numerically construct a binned distribution $\Delta^2 \dot{N} / (\Delta M \Delta z)$ by summing up the sources in the CRTS

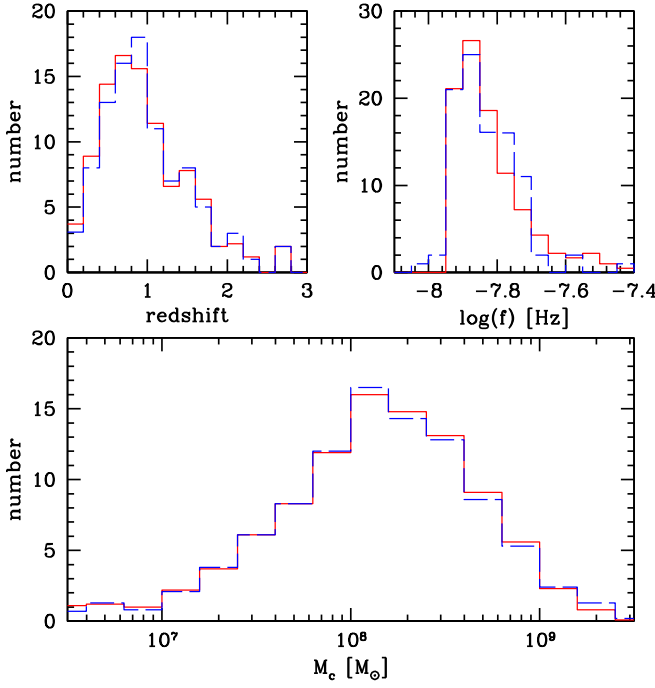


Figure 4. Comparison between the 98 SMBHB candidates in the CRTS, and a mock population of SMBHBs constructed from the merger rate implied by the CRTS systems. Upper left: redshift distribution. Upper right: rest-frame frequency distribution. Lower panel: chirp mass distribution. The chirp mass distribution is constructed assuming Model_True_fid (see the text for details). In all panels, solid red histograms are the CRTS sample, dashed blue histograms are an average over 10 Monte Carlo realizations of Model_True_fid.

sample as follows:

$$\frac{\Delta^2 \dot{N}}{\Delta \mathcal{M} \Delta z} = \sum_{i \in \Delta \mathcal{M}} \sum_{i \in \Delta z} \frac{1}{\Delta \mathcal{M} \Delta z T_{c,i}(\tilde{f}_r)}. \quad (7)$$

Here $i = 1, \dots, 98$ is an index identifying the SMBHB candidates, and $T_{c,i}(\tilde{f}_r)$ is the coalescence timescale of the i -th binary according to Equation (5). Note that although Equation (7) depends on the choice of the bins $\Delta \mathcal{M}$ and Δz , when computing the total merger rate via Equation (9) or the GW signal via Equation (10) below, we integrate over $d\mathcal{M}$ and dz , and the final results are bin-independent. To ease the notation, we switch to the continuum and use the differential form $\Delta^2 \dot{N}/(\Delta \mathcal{M} \Delta z) \rightarrow d^3 \dot{N}/(d\mathcal{M} dz dt_r)$. We stress, however, that all computations have been performed numerically on binned distributions and the robustness of the results have been checked against the choice of bin sizes.

To check whether our merger rate is consistent with the observed CRTS SMBHB population, we can now construct the expected $d^3 \dot{N}/(d\mathcal{M} dz df_r)$ as

$$\frac{d^3 \dot{N}}{d\mathcal{M} dz df_r} = \frac{d^3 \dot{N}}{d\mathcal{M} dz dt_r} \frac{dt_r}{df_r}, \quad (8)$$

where dt_r/df_r is the same as Equation (4) but evaluated in the source rest-frame. We can then perform a Monte Carlo sampling of the distribution given by Equation (8) and compare it to the original CRTS SMBHB candidate ensemble. The comparison is given in Figure 4, where we show the chirp mass, rest-frame frequency, and redshift distribution averaged over 10 realizations of Model_True_fid. The inferred merger

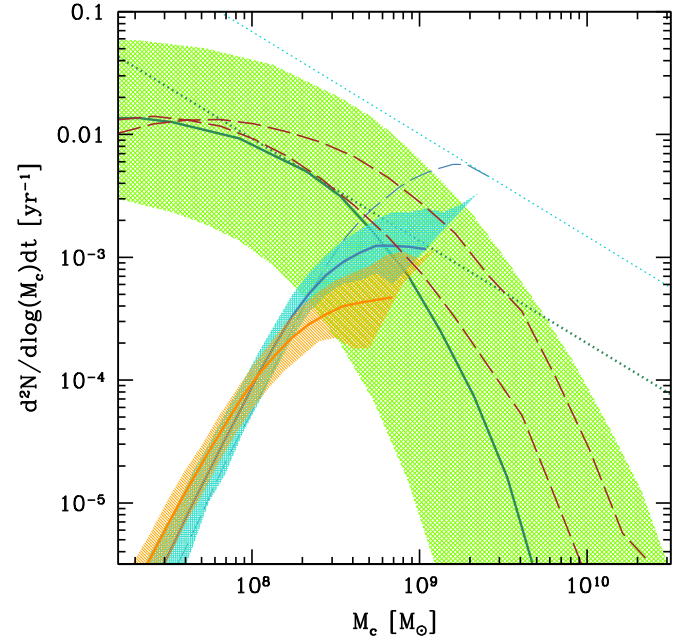


Figure 5. Merger rate mass function comparison. The green shaded area is the 95% interval produced by theoretical models presented in Sesana (2013b) and the solid dark green line is the median value. The blue and orange shaded areas are the 95% intervals based on the observed samples of periodic AGNs produced by Model_True_hiq and Model_True_loq described in the text, and the thick blue and orange lines are the respective median values. Model_True_fid lies in between the two and it is omitted to preserve figure readability. We also show the median merger rate obtained in Model_Vir_hiq (blue long-dashed line). The two dashed brown lines are two models selected from Sesana et al. (2009). The dotted curves with a $-5/6$ slope are tangent to the median values predicted by selected models and show which chirp mass contributes the most to the GW signal (see the text for details).

rate numerically constructed as explained above is perfectly consistent with the observed CRTS SMBHB population. Equation (7), therefore, provides a reliable estimate of the SMBHB merger rate implied by the observed candidates. We caution that this might differ from the intrinsic SMBHB merger rate as, in practice, completeness may depend on M , q , f_r , which distorts the distribution as discussed below.

3.1. Coalescence Rates

We can now compute the observed differential merger rate per unit chirp mass, by integrating Equation (7) in a given redshift range,

$$\frac{d^2 \dot{N}}{d\mathcal{M} dt} = \int_{z_{\min}}^{z_{\max}} \frac{d^3 \dot{N}}{d\mathcal{M} dz dt_r} \frac{dt_r}{dt} dz. \quad (9)$$

Results are shown in Figure 5, where the merger rates of Model_True_hiq and Model_True_loq are integrated in the redshift range $0 < z < 1.3$ and are compared to theoretical estimates from the literature (integrated in the same redshift range). In particular, we consider an updated version of the observation-based models of Sesana (2013b), and two models extracted from the Millennium simulation (Springel et al. 2005) and described in Sesana et al. (2009). The sharp decline at $\mathcal{M} < 3 \times 10^8 M_\odot$ of Model_True_hiq and Model_True_loq is likely due to CRTS incompleteness; lighter SMBHBs are intrinsically fainter and might be missing in the relatively shallow Catalina survey. Both models are consistent with

theoretical estimates, even though incompleteness of the CRTS sample can significantly increase the underlying merger rate, which might create tension at high masses, especially for Model_True_hiq. We will discuss the effect of incompleteness on our results in Section 4.3.

Figure 5 also shows the median merger rate obtained when Model_Vir_hiq is assumed (blue long-dashed line). In this case, the CRTS population is inconsistent at least at a 95% level at $\mathcal{M} > 2 \times 10^9 M_\odot$, with respect to expectations from hierarchical clustering. Note that incompleteness of the CRTS sample can only increase the merger rate, which makes this discrepancy worse at high masses; in fact, although undetected low-mass systems might close the gap at $\mathcal{M} < 10^8 M_\odot$, any correction for incompleteness can only make the mismatch at the high mass end more severe. In particular, we will see in Equation (10) below that for a given merger rate, the GW strain is $h_c \propto \mathcal{M}^{5/6}$. This means that the largest contribution to the GW signal comes from the value of \mathcal{M} where the merger rate distribution is tangent to a line with $-5/6$ slope in the log-log plane, as depicted in Figure 5. We see that the line tangent to the median of Model_Vir_hiq gives a systematically higher rate by an order of magnitude than the merger rate estimates from theoretical models, which results in a large overprediction of the GW signal (as is later shown in Figure 7). This highlights that a proper estimate of the true masses of the CRTS candidates is of paramount importance.

4. PTA Implications of the Catalina Sample

With a reliable estimate of the SMBHB merger rate in hand, we can now proceed to the computation of the expected GW signal. We start by considering circular, GW-driven binaries. Following Phinney (2001) and Sesana et al. (2008), we can write the overall GW signal as

$$h_c^2(f) = \frac{4}{\pi f^2} \int_0^\infty dz \int_0^\infty d\mathcal{M} \frac{d^2n}{dzd\mathcal{M}} \frac{1}{1+z} \frac{dE_{\text{gw}}(\mathcal{M})}{d \ln f_r}. \quad (10)$$

The energy radiated per logarithmic frequency interval is

$$\frac{dE_{\text{gw}}}{d \ln f_r} = \frac{\pi^{2/3}}{3} \mathcal{M}^{5/3} f_r^{2/3}, \quad (11)$$

and $d^2n/dz d\mathcal{M}$ is the comoving number density of mergers per unit redshift and chirp mass and is related to the overall cosmic merger rate of Equation (7) via

$$\frac{d^2n}{dz d\mathcal{M}} = \frac{d^3N}{d\mathcal{M} dz dt_r} \frac{dt_r}{dz} \frac{dz}{dV_c}. \quad (12)$$

The functions dt_r/dz and dz/dV_c are given by standard cosmology. We can therefore construct the expected GW signal by using Equation (10) for all our models. For each model, we compute h_c from the 1000 realizations of the CRTS ensemble, to get a measurement of the uncertainty in the predicted signal amplitude.

Note that in the circular GW-driven scenario, the GW signal computed through Equation (10) is a simple power law with a spectral slope $-2/3$ (Phinney 2001). The GW amplitude is thus

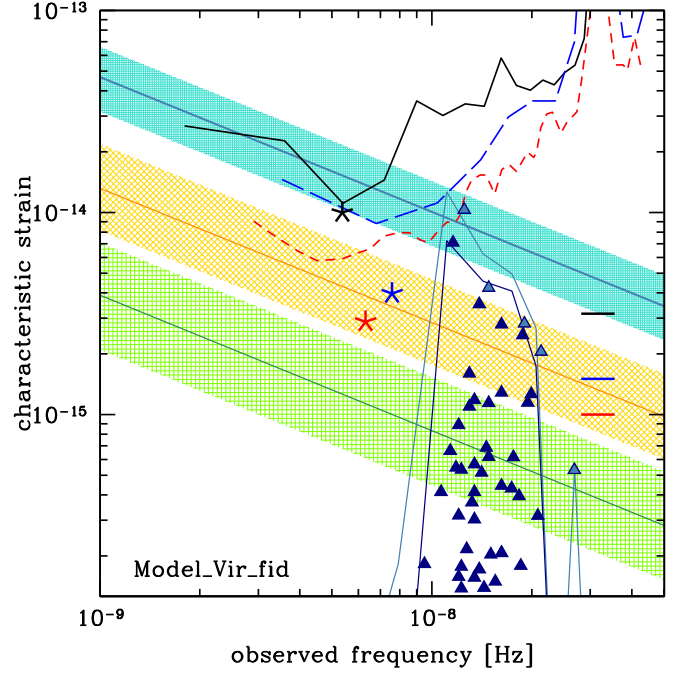


Figure 6. Example of the characteristic amplitude of the GW signal for circular, GW-driven binaries. The orange and green shaded areas are 68% intervals produced by selected models presented in Sesana (2013b) and (Sesana et al. 2016; see the text for details); the solid orange and green lines are their respective median values. The blue shaded area is the 68% interval produced by Model_Vir_fid and the solid blue line is the median value. The dark blue triangles illustrate one Monte Carlo realization of individual GW signals from the G15 sample. The light blue triangles are sources that are individually louder than all other binaries in the same frequency bin (and therefore, potentially resolvable individually by a putative PTA with enough sensitivity), taken here to be $\Delta f = 3.17$ nHz (i.e., 10 yr^{-1}). The corresponding light blue line is the resulting overall signal, and the dark blue line is the level of the GWB after these putative resolvable sources are subtracted. The solid black, long-dashed blue and short-dashed red lines represent the sensitivities to a stochastic GW signal as a function of frequency for the EPTA, NANOGrav, and PPTA respectively. The corresponding stars represent the 95% upper limit to the integrated GWB from each PTA, placed at the frequency at which it is most sensitive. The horizontal ticks are the extrapolation of those limits to a frequency of 1 yr^{-1} .

customarily written as

$$h_c = A \left(\frac{f}{1 \text{ yr}^{-1}} \right)^{-2/3}, \quad (13)$$

and PTA results are quoted as limits on the amplitude normalization A at a nominal frequency of 1 yr^{-1} .

4.1. Model_Vir: Severe Inconsistency with PTA Measurements

Figure 6 shows the outcome of the calculation just described for Model_Vir_fid. h_c given by Equation (10) is compared with the expected signal from observation-based models featuring different SMBH-galaxy scaling relations. In particular, we pick from Sesana (2013b) a fairly “optimistic” model in which the SMBH mass correlates with the host bulge mass following the relation proposed by Kormendy & Ho (2013). Note that the GW amplitude predicted by this model, $-15.1 < \log A < -14.7$ at 68% confidence, is already in tension with the current best 95% GWB upper limit at $\log A < -15$ (Shannon et al. 2015). We also consider an alternative model featuring a much more conservative SMBHB-host relation proposed by Shankar et al. (2016). This

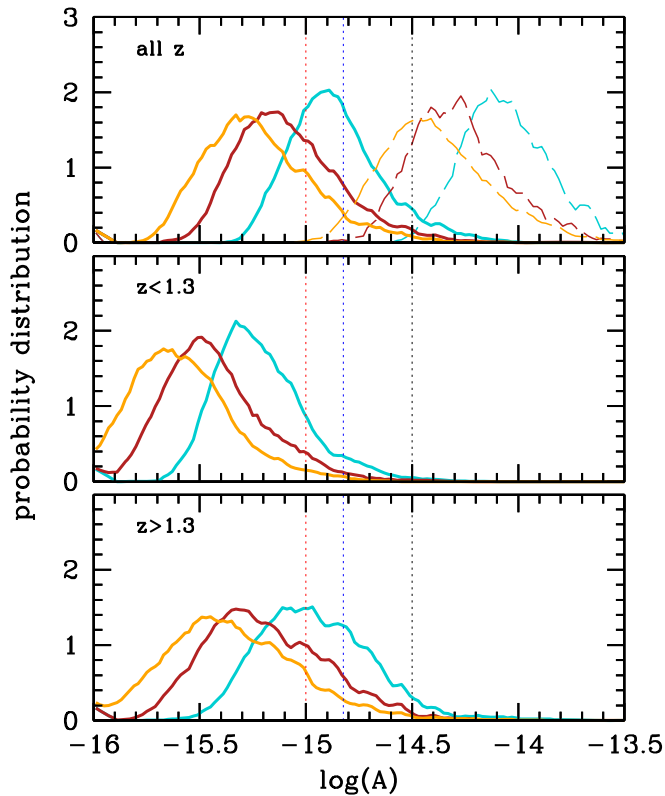


Figure 7. Distribution of the stochastic GWB characteristic amplitudes A at the fiducial frequency of 1 yr^{-1} inferred from the CRTS sample under different assumptions about the masses and mass ratios of the SMBHBs. The blue curves are distributions for Model_Vir_hiq (dashed) and Model_True_hiq (solid), brown curves are for Model_Vir_fid (dashed) and Model_True_fid (solid), and orange curves are for Model_Vir_loq (dashed) and Model_True_loq (solid). The vertical dotted lines are 95% upper limits given by EPTA (black), NANOGrav (blue), and PPTA (red). The three panels show the overall amplitude distribution considering all CRTS candidates (top) and the contribution of candidates at $z < 1.3$ (center) and $z > 1.3$ (bottom). In the latter two panels, only the “True” models are shown.

model predicts a GW amplitude of $-15.8 < \log A < -15.2$ at 68% confidence (Sesana et al. 2016).

The figure highlights the key idea behind our calculation. Although signals from individual candidates are perfectly consistent with current PTA limits, even when virial masses are taken at face value, the implied stochastic GWB extrapolated at lower frequencies is strongly inconsistent with current IPTA upper limits (Lentati et al. 2015; Shannon et al. 2015; Arzoumanian et al. 2016). This is true for all “Vir” models. In the top panel of Figure 7, we show the GWB amplitude A at the fiducial frequency of 1 yr^{-1} , obtained from the 1000 Monte Carlo realizations of all the “Vir” models (dashed lines on the right). We obtain $\log(A_{50\%}) = -14.11, -14.32, -14.45$ for Model_Vir_hiq, Model_Vir_fid, and Model_Vir_loq, respectively, which are severely inconsistent with the PTA limits.

4.2. Model_True: Lower Signal Normalization

Taking virial mass measurements at face value, the CRTS sample is inconsistent with PTA limits, implying that for the majority of the candidates, inferred variability cannot be linked to the presence of an SMBHB. Those mass estimates are, however, known to be biased high, as described in Section 2.1. We therefore consider the “True” models, where this bias has been corrected via Equation (3), to achieve more robust

conclusions. The GWB amplitude A from 1000 Monte Carlo realizations in each model is shown as solid lines in the top panel of Figure 7, and highlights the critical importance of taking the bias into account. The distributions are now all at least marginally consistent with the most stringent current PTA upper limit (Shannon et al. 2015) and we obtain $\log(A_{50\%}) = -14.88, -15.13, -15.25$ for Model_True_hiq, Model_True_fid and Model_True_loq respectively. The ≈ 0.8 dex difference with respect to the “Vir” models is easy to explain. From Equation (10), h_c is proportional to the square root of the merger rate times the energy spectrum, both of which are proportional to $\mathcal{M}^{5/3}$. Therefore, a reduction of about 0.5 dex in the mass estimate (see Figure 2) results in a corresponding reduction of ≈ 0.8 dex in the GWB.

Even though the bulk of the GWB is theoretically expected to come from SMBHBs at $z < 1.5$ (Sesana et al. 2008; Ravi et al. 2015; Simon & Burke-Spolaor 2016), the G15 sample features several fairly massive candidates at higher redshifts. Note that, according to Equation (10), the GW contribution depends on the differential density of mergers per unit redshift and chirp mass, $d^2n/(dzd\mathcal{M})$. Each individual candidate in the CRTS sample contributes to this quantity via Equations (5) and (12). The latter includes the dz/dV_c factor accounting for the comoving volume shell accessible at a given redshift: a single binary observed at higher redshift implies a lower merger rate density because of the larger comoving volume available. One might therefore think that higher z systems do not contribute significantly to the GW signal. Since, for a fixed observed frequency, binaries at higher redshift are emitting at higher rest-frame frequencies; however, their intrinsic coalescence time is shorter, and the inferred merger rate larger according to Equation (5). It turns out that this latter fact compensates for averaging over a larger volume shell, and the GW signal is dominated by SMBHBs at $z > 1.3$. This is shown in the center and bottom panels of Figure 7, where we break down the contribution to the GWB from sources at $z < 1.3$ and $z > 1.3$. The latter contribute, on average, about two-thirds of the total GWB, in striking contrast with theoretical expectations.

To single out the individual candidates contributing the most to the GWB, we performed the following experiment. Under the assumption of the fiducial Model_True_fid, we generated 1000 realizations of M_1 and M_2 of each of the 98 CRTS binaries. We then used Equations (5) and (12) to compute their contribution to the SMBHB merger density, and folded the result into Equation (10) to compute the associated GW signal. We stress again that although the merger rate obtained via Equation (7) depends on the chosen bin size, this is compensated by the integral over the size of the bin in Equation (10). We then rank the sources by individual contribution to the GWB for each realization of the signal, progressively removing sources one by one from the loudest to the quietest. The result is shown in Figure 8 for the progressive removal of the loudest 50 systems. The overall signal drops steeply as the first 10 sources are removed. Their median contribution to the GWB for this specific model is listed in the last column of Table 1. Note that their combined contribution alone accounts for about 70% of the overall GWB and note also that the five loudest candidates are at $z > 1.5$, where we do not expect a significant contribution to the GWB from theoretical models. This analysis puts the CRTS candidate distribution strongly at odds with theoretical expectations, even while the distribution is marginally consistent with PTA observations alone.

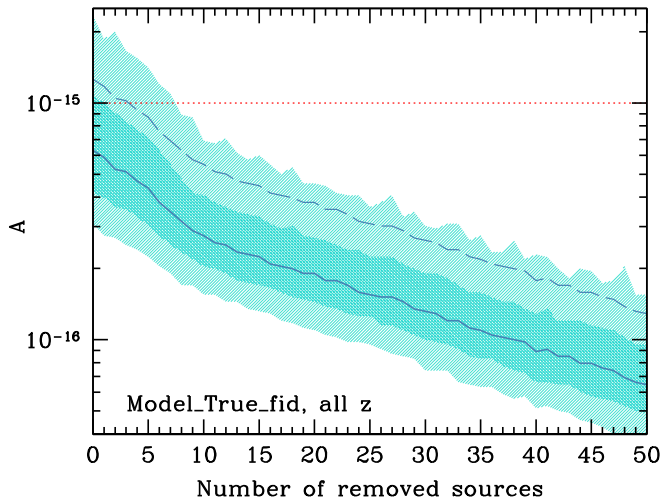


Figure 8. Build-up of the GWB signal. Plotted is the characteristic strain of the GWB at $f = 1 \text{ yr}^{-1}$, A , vs. the number of removed CRTS candidates. Candidates are ranked and removed, starting with the loudest, from left to right. Shaded areas mark the 68% and 95% confidence region of the amplitude computed over 1000 realizations of the signal, and the solid blue line is the median value. The dashed blue line is the median shifted upward by a factor of two to show a conservative estimate of incompleteness (the real impact of incompleteness is likely larger; see the discussion in Section 4.3). The horizontal dotted red line is the PPTA upper limit. Model_True_fid is assumed.

4.3. Correcting for Incompleteness: Tension between Model_True and PTA Upper Limits

Although the CRTS sample seems consistent with current PTA upper limits, we did not yet consider the fact that the SMBHB sample is necessarily incomplete in several ways. First, CRTS covers only 80% of the sky (about $33,000 \text{ deg}^2$). Although this is a significant fraction of the whole sky, most of the selected quasars are identified in the SDSS survey (78 out of 111), which covers only about one-quarter of the sky. Assuming a complete identification in the SDSS field of view, we would therefore expect $78 \times 4 = 312$ objects in the whole sky. We conclude that the G15 catalog is at best only 35.5% complete based on sky coverage arguments only. Moreover, of the 334 k identified quasars, 83 k (25%) were rejected because of sparse sampling. Finally, we dropped 13 candidates with no mass estimate (12%) from the sample. Adding everything up, the sample we use is at best only 24% complete. Since the signal is proportional to the square root of the coalescence rate, all the signal estimates shown so far must be multiplied by at least a factor of two.

The effect of incompleteness is shown in Figure 9. The GWB predicted from the G15 sample, corrected for incompleteness, starts now to be in tension with current PTA measurements, in particular, with the PPTA limit. In fact, we find $\log(A_{50\%}) = -14.60, -14.82, -14.97$ for Model_True_hiq, Model_True_fid, and Model_True_loq, respectively, all of which are above the 95% upper limit at $\log(A) = -15$, published by PPTA. When virial mass measurements are assumed (“Vir” models), the predicted amplitude is severely inconsistent with all PTA upper limits.

Note that, besides these simple “counting arguments,” which can be easily corrected for, there are potentially many more sources of incompleteness in the G15 SMBHB sample. First, the redshift distribution of the CRTS candidate quasars shows a prominent peak around $z \approx 0.8$, hinting to incompleteness at

higher redshifts. Second, the candidates are all spectroscopically confirmed type 1 quasars; numbers should therefore be corrected for the fraction of obscured type 2 quasars, which is poorly known but can easily double the sample (Lusso et al. 2013). Third, the search is constructed to look for sinusoidal periodicities and would be much less sensitive to eccentric binaries. Finally, not all SMBHBs have to be active in the first place, especially in gas-poor mergers that are rather frequent at low redshift. In any case, the factor of two corrected signal shown in Figure 9 is necessarily a lower limit to the intrinsic GWB implied by the CRTS SMBHB candidate population.

4.4. Bayesian Model Selection: Preference for the Null Hypothesis

We now wish to properly quantify the concept of “tension” between the CRTS sample and the PTA measurements. To do so, we employ the concept of Bayesian model selection considering two competitive hypotheses: the null hypothesis N that the CRTS candidates are not binaries, versus the hypothesis B that the candidates are indeed binaries. When considering nonparametric models, the odds ratio of model N over model B is simply given by

$$\Lambda_{NB} = \frac{p(D|N) p(N)}{p(D|B) p(B)}, \quad (14)$$

where the model likelihood $p(D|X)$ is the likelihood of the observed data D under the X hypothesis, and $p(X)$ is the prior probability assigned to model X . If we make the agnostic assumption that models N and B are a priori equally probable, then the odds ratio reduces to computing the likelihood ratio of model N over B . In a pairwise comparison, it is then possible to associate a probability $p_N = p(D|N)/(p(D|N) + p(D|B))$ to the null hypothesis, and a probability $p_B = 1 - p_N$ to the binary hypothesis.

Given the data, the model likelihoods can be evaluated as the integral of the amplitude distribution predicted by a specific model times the posterior amplitude distribution derived by the data. The likelihood of model X is thus

$$p(D|X) = \int \mathcal{F}(A) p_X(A) dA. \quad (15)$$

We consider here the PPTA data used in Shannon et al. (2015). The posterior amplitude distribution is fairly well described by a Fermi function of the form

$$\mathcal{F}(A) = \frac{C_1}{e^{(A-C_2)/C_3} + 1}, \quad (16)$$

with $C_1 = 1.63$, $C_2 = 1.2 \times 10^{-16}$, and $C_3 = 2.6 \times 10^{-16}$ (Middleton et al. 2017), and is shown in Figure 9. We need now to specify the amplitude probability of the competing models. In the null hypothesis, no binaries are present in the data, the amplitude distribution is therefore a delta function centered at $A = 0$, $p(A) = \delta_0$. We consider pairwise comparisons with each of the binary models proposed in this paper, so model B will be Model_True_hiq, Model_True_fid, Model_True_loq, Model_Vir_hiq, Model_Vir_fid, and Model_Vir_loq, and the associated amplitudes are those computed from the 1000 realizations of the GWB, corrected for incompleteness in each case, as shown in Figure 9.

Table 1
Top 10 **G15** Candidates Providing the Largest Contribution to the Expected GWB

Name	z	$P[\text{day}]$	$\log\left(\frac{M_t}{M_\odot}\right)$	$A_s[10^{-16}]$
SDSS J164452.71+430752.2	1.715	2000	10.15 (9.66)	1.644
HS 0926+3608	2.150	1562	9.95 (9.47)	0.718
SDSS J140704.43+273556.6	2.222	1562	9.94 (9.45)	0.698
SDSS J092911.35+203708.5	1.845	1785	9.92 (9.44)	0.698
SDSS J131706.19+271416.7	2.672	1666	9.92 (9.45)	0.679
HS 1630+2355	0.821	2040	9.86 (9.34)	0.608
SDSS J114857.33+160023.1	1.224	1851	9.90 (9.38)	0.544
SDSS J134855.27-032141.4	2.099	1428	9.89 (9.41)	0.515
SDSS J160730.33+144904.3	1.800	1724	9.82 (9.37)	0.515
SDSS J081617.73+293639.6	0.768	1162	9.77 (9.28)	0.474

Note. Columns give the object identification name, redshift, observed orbital period in days, total mass as given in **G15** (median mass estimate when bias is included is given in parenthesis), and individual contribution to the GWB $A_s(f=1\text{yr}^{-1})$ in Model_True_fid.

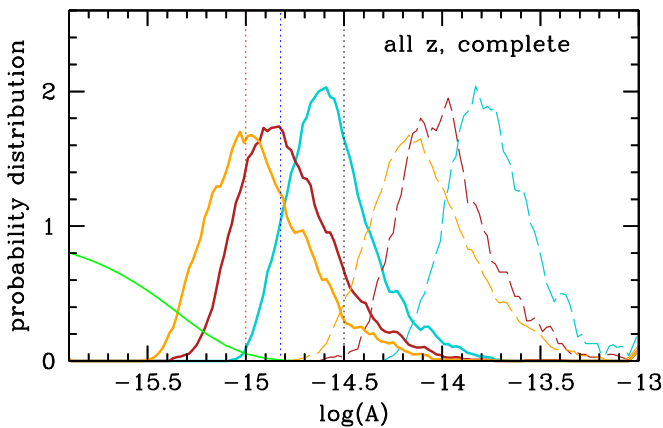


Figure 9. Same as the top panel of Figure 7, but correcting for completeness of the CRTS sample. Line style as in Figure 7, the extra green line represents the posterior sample distribution of the PPTA amplitude measurement from which the 95% upper limit of $A = 10^{-15}$ is derived.

Note that $\mathcal{F}(0) = 1$ so that the likelihood of the null hypothesis is trivially $p(D|N) = 1$, whereas the likelihood of the B hypothesis has to be computed via numerical integration of Equation (15). Results are shown in Table 2. It is clear that the “Vir” models are inconsistent with PTA limits and are in fact strongly disfavored compared to the null hypothesis. The situation is not so conclusive when the fiducial “True” models are considered. $\log\Lambda_{NB}$ are now in the range 1.1–2.5, indicating substantial preference for the null hypothesis (Kass & Raftery 1995). Depending on the mass ratio distribution, the null hypothesis is preferred at the 99.7%, 97.1%, and 92.7% confidence level. Translated in the familiar “ σ -level” jargon, the null hypothesis is preferred over our fiducial model (Model_True_fid) at $\approx 2.3\sigma$.

4.5. Possible Impact of SMBHB Coupling with Stars and Gas

In the previous discussion, we considered circular GW-driven binaries. However, it has been shown that both coupling with the environment and large eccentricities can, in principle, suppress the GW signal at nHz frequencies (see, e.g., Enoki & Nagashima 2007; Kocsis & Sesana 2011; Sesana 2013a; Ravi et al. 2014).

Since the **G15** SMBHB candidates are quasars, it is important to explore the case of strong coupling with a gaseous environment. Relevant models were constructed by

Table 2
Model Selection Results: Pair Comparisons between the Null Hypothesis and All the Investigated Models Based on the CRTS Sample

Model Pair	Odds Ratio and Probabilities		
	$\log\Lambda_{NB}$	p_N	p_B
N/True_hiq	2.61	0.9976	0.0024
N/True_fid	1.52	0.9709	0.0291
N/True_loq	1.10	0.9267	0.0733
N/Vir_hiq	11.1	>0.9999	< 10^{-11}
N/Vir_fid	6.67	>0.9999	< 10^{-6}
N/Vir_loq	5.22	>0.9999	< 10^{-5}

Note. For each two-model comparison, we report the log of the likelihood ratio Λ_{NB} , and the probability of the null hypothesis (p_N) and the binary hypothesis (p_B).

Kocsis & Sesana (2011). They studied the SMBHB-disk coupling assuming different thin disk models (Shakura & Sunyaev 1973) described in Haiman et al. (2009). In particular, they considered steady-state solutions by Syer & Clarke (1995), assuming both α and β disks, and the self-consistent nonsteady solution of Ivanov et al. (1999). They found that only α disks with a large viscosity parameter ($\alpha_{SS} = 0.3$, in that work) and large accretion rate ($\dot{M}/\dot{M}_{\text{Edd}} = 0.3$, in that work) can significantly suppress the signal at the frequencies where current PTAs are sensitive. We therefore consider this model here, even though it might be unlikely to occur in nature because the secular thermal-stability of α -disks is uncertain (see, e.g., Hirose et al. 2009; Jiang et al. 2013). Moreover, we include eccentricity, according to the findings of Roedig et al. (2011), where binaries have $e \approx 0.6$ as long as they are coupled with the disk, and they progressively circularize due to GW emission after decoupling. Results are shown in the left panel of Figure 10 for Model_True_fid. Even considering such an extreme coupling, the implied GW signal is reduced only by a factor of ≈ 1.5 at 6 nHz (and ≈ 7 at 1 nHz) and is still generally in tension with PTA upper limits. We stress that any other model (i.e., alpha disks with a smaller α -viscosity, any β disk, or any nonsteady self-consistent disk) would result in a negligible suppression of the signal in the PTA frequency range.

Coupling with stars can also accelerate SMBHB evolution and promote eccentricity growth (see, e.g., Quinlan 1996; Sesana 2010), both of which suppress the low-frequency GW

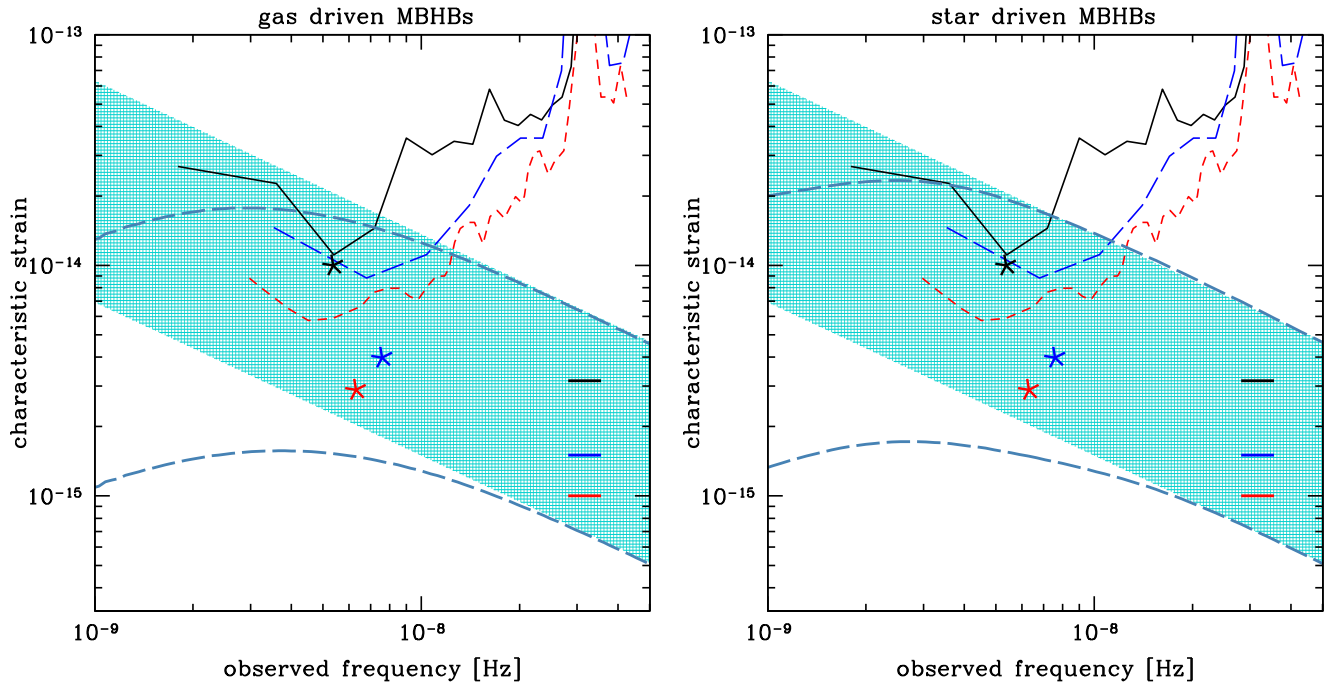


Figure 10. Effect of gas (left) and stellar (right) dynamics on the GW signal implied by the CRTS SMBHB sample (see the text for details of the models of SMBHB-environment couplings). In each panel, the blue shaded area is the 95% confidence intervals produced by the fiducial Model_True_fid. The area was computed, including G15 SMBHBs at all redshifts and correcting for incompleteness of the sample (see the text for full details). The long-dashed blue lines represent the 95% predicted range for h_c under the effect of environmental coupling. The representation of the PTA sensitivity curves and upper limits is the same as that in Figure 6.

signal. Sesana (2010) modeled the stellar environment as a broken power law, with a central cusp of density $\rho \propto r^{-\gamma}$ joining an external isothermal sphere at the SMBHB influence radius. In those models, the SMBHB loss cone is assumed to be always full, which implies that the binary evolves at a pace that is dictated by the stellar density ρ_0 at the influence radius. Since the isothermal sphere is quite compact, especially compared to observed inner density profiles of massive ellipticals, these models can overestimate ρ_0 by more than an order of magnitude, and are therefore aggressive in terms of the efficiency of the SMBHB-stellar coupling. We consider here the model with $\rho \propto r^{-1}$ and $e_0 = 0.6$ (e_0 is the eccentricity at binary formation). Results are shown in the right panel of Figure 10 again for Model_True_fid. Also, in this case, the suppression of the GWB is mild, and the result is still in tension with PTA upper limits.

In principle, a larger suppression is possible if SMBHBs are more eccentric. However, as an example, the Bayesian analysis performed by D’Orazio et al. (2015b) on one of the G15 binary candidates, the $z = 0.3$ quasar PG 1302-102, showed that its eccentricity is at most ≈ 0.2 and it is consistent with zero. All SMBHBs in CRTS have been identified by their sinusoidal behavior, which implies small eccentricities. The models considered in this section are consistent with SMBHBs having $e < 0.25$ at $f > 10^{-8}$ Hz, i.e., in the frequency range probed by CRTS. For example, a stellar driven model with $e_0 = 0.8$ would result in a population of fairly eccentric binaries at $f > 10^{-8}$ Hz (e.g., Kelley et al. 2017b), inconsistent with the sinusoidal light curves observed in the CRTS sample.

We caution that our computation underestimates the signal in the case of extreme environmental couplings. In fact, the normalization of the signal is set by the merger rate that is computed according to Equation (7) assuming circular GW-driven binaries. If SMBHBs are still coupled with the

environment, then the coalescence time $T_c(f)$ would be shorter, resulting in a higher merger rate and, in turn, in a higher signal normalization.

5. Application to the PTF Sample

We next apply the same methodology described in Section 3 to the PTF sample identified by C16. By studying a sample of 35,383 spectroscopically confirmed quasars from PTF, they construct a sample of 33 binary candidates with periods $\lesssim 500$ days. Masses reported in C16 are either from Shen et al. (2008) or estimated by the authors using a similar technique based on the luminosity alone (in those cases when Shen et al. 2008 measurements were unavailable; note that the luminosity-based estimates have a slightly larger uncertainty). These are therefore again virial masses, and a systematic bias, as in the CRTS sample, should be present. We therefore consider also in this case the same “Vir” and “True” models, constructed as described in Section 2.1, assuming $\tilde{P}_{\text{orb}} = 500$ days, consistent with the shorter timespan of the PTF data set.

We also note that the sample is affected by severe incompleteness. Of the 33 candidates, 28 fall in the SDSS footprint. Considering that the PTF and SDSS footprints overlap over 9700 deg^2 , one would expect $28 \times (41253/9700) = 119$ systems for a uniform sky coverage. Moreover, of the 278,740 spectroscopically confirmed quasars in PTF, only 35,383 passed the selection criteria in terms of observational cadence. Considered together, these two facts imply that the C16 sample is only $\lesssim 3\%$ complete. This implies an upward correction factor of $\approx \sqrt{40} \approx 5.8$ when computing the associated GWB.

We find that, when corrected for incompleteness, the predicted GWB amplitudes in the “Vir” models are inconsistent with PTA upper limits by more than an order of magnitude. “True” models are also inconsistent with the data, as shown in

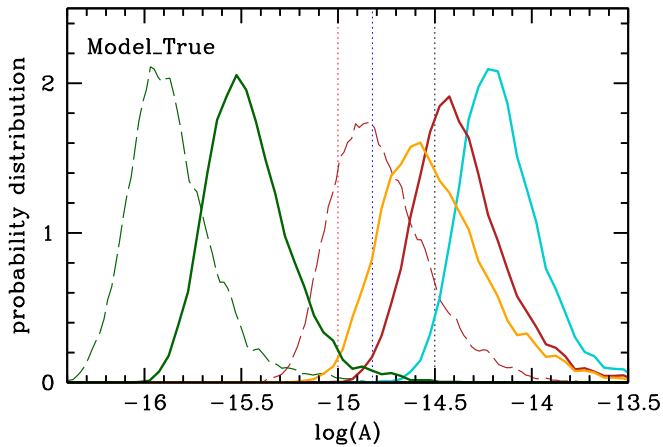


Figure 11. Distribution of the stochastic GWB characteristic amplitudes inferred from the PTF sample. Shown are the default Model_True_hiq (light blue), Model_True_fid (brown), and Model_True_loq (orange) alongside an extreme model in which all candidates have $q = 0.01$ (thick solid green). Thin dashed lines show for comparison the GWB produced by the CRTS sample under Model_True_fid (brown) and assuming $q = 0.01$ (green). The vertical dotted lines are 95% upper limits given by EPTA (black), NANOGrav (blue), and PPTA (red).

Figure 11. The average GWB amplitude is estimated to be $\log(A_{50\%}) = -14.20, -14.42, -14.57$ for Model_True_hiq, Model_True_fid, and Model_True_loq respectively. If we apply the analysis described in Section 4.4, those models are inconsistent with the PPTA measurement at a level of 4.5σ , 3.6σ , and 3.0σ as reported in Table 3.

When analyzing this sample, Charisi and coauthors find that the observed period distribution best matches the expected SMBHB coalescence time distribution if their mass ratio is $q \approx 0.01$. We therefore constructed an alternative model where we assumed “True” mass estimates but an extreme mass ratio $q = 0.01$ for all candidates. The expected GWB amplitude (after correcting for incompleteness) is also displayed in Figure 11 in green. Under this ansatz, the PTF sample would be consistent with current PTA data. Note that the same would be true for the CRTS sample (also shown in the figure).

A few things should be noted though. A mass ratio distribution of $q \approx 0.01$ might be expected if the periodicity is due to Doppler boosting, as also noted by C16. If this is the case, however, there is an observational bias in favor of SMBHBs with small inclination angles (i.e., nearly edge on). A proper estimate of the GWB within this interpretation should take this bias into account. Moreover, we did not consider any correction for missing SMBHBs with $q > 0.01$; if Doppler boosting manifests itself only at these small q , then there must be a large population of undetected SMBHBs with larger q contributing to the GWB, likely boosting it to unacceptable levels. Alternatively, one has to envisage a model whereby only $q \approx 0.01$ binaries form, which would be unexpected in the context of hierarchical build-up of SMBHBs.

Summarizing:

1. the PTF sample is 3% complete;
2. after correcting for incompleteness, all Model_Vir are more than an order of magnitude inconsistent with the PTA upper limit;
3. the fiducial model Model_Vir_fid is inconsistent with the PPTA upper limit at 3.6σ ;
4. if we assume $q \approx 0.01$ as in C16, then the PTF sample is consistent with PTA data. However, the GWB computed

Table 3
Same as Table 2 for the PTF Sample

Model Pair	Odds Ratio and Probabilities		
	$\log \Lambda_{NB}$	p_N	p_B
N/True_hiq	5.55	>0.9999	$<10^{-5}$
N/True_fid	3.60	0.9997	0.0003
N/True_loq	2.57	0.9973	0.0027

Note. Only “true” models are displayed.

in this way is necessarily a lower limit to the total GWB because of (i) inclination bias and (ii) higher q binaries being missed in the sample.

6. Conclusions

We have shown that the list of SMBHB candidates identified by Graham et al. (2015) in the CRTS and by Charisi et al. (2016) in the PTF are both in tension with current pulsar timing array limits on a stochastic GW background at nHz frequencies. The bulk of our analysis focused on the CRTS sample. Although none of the candidates, taken individually, is inconsistent with PTA measurements, they can be collectively used to construct a cosmic SMBHB merger rate and the associated GWB at nHz frequencies. The GWB computation implies the knowledge of the chirp masses of the candidates; we therefore need to assign a total mass and a mass ratio to each object. We considered two mass models that we labeled Model_Vir and Model_True. In the former, virial masses are directly taken from Shen et al. (2008); whereas in the latter, prior knowledge of the SMBH mass function is used to infer the “true” mass of each candidate from its virial mass, thus correcting the intrinsic bias in individual virial mass estimates. For each of the mass models, we explored three different possible mass ratio (q) distributions: one preferentially high (hiq), another low (loq), and a fiducial intermediate model (fid). This gives a total of six models: Model_True_hiq, Model_True_fid, Model_True_loq, Model_Vir_hiq, Model_Vir_fid, and Model_Vir_loq.

Our main findings can be summarized as follows:

1. The GWB calculation critically depends on the mass estimates of the SMBHB candidates. Each Model_Vir results in a GWB that is ≈ 0.8 dex larger than the corresponding Model_True.
2. The mass ratio distribution has a milder impact, affecting the GWB at a factor of ≈ 2 level. This is because the GW signal depends on the chirp mass, $\mathcal{M} = Mq^{3/5}$, that has a milder dependence on q .
3. All Model_Vir are severely inconsistent with all PTA measurements. This is, however, understandable since virial mass estimates are known to be biased high.
4. When corrected for sample incompleteness, all Model_True show significant tension with the PPTA upper limit (the most stringent to date), at a level that ranges between 1.8σ and 3σ depending on the mass ratio distribution. Our fiducial Model_True_fid is inconsistent with the PPTA measurement at 2.3σ .
5. About two-thirds of the GWB is generated by sources at $z > 1.3$ in stark contrast with theoretical expectations; this is likely due to the presence of selection effects in the CRTS sample.

6. On the other hand, about 70% of the signal is contributed by a set of 10, mostly high-redshift SMBHB candidates (see Figure 8 and Table 1). Misidentification of a few particularly loud systems as binaries can therefore have a significant impact on the expected GWB.
7. Environmental coupling can decrease the signal by at most a factor of 1.5 at 6 nHz (i.e., where current PTAs are most sensitive), only partly alleviating the tension with PTA measurements. We stress that this estimate is optimistic, because it does not take into account that coupling with the environment, which would shorten the coalescence timescale of the SMBHB candidates, resulting in a higher merger rate and, in turn, in a higher GWB normalization (which we ignored in Section 4.5).

We stress that we implicitly assumed that all SMBHBs with orbital periods of a few years are periodic quasars. If only a fraction of those SMBHBs exhibit AGN activity then the predicted stochastic GW background would severely violate the current PTA upper limits.

These results show that the SMBHB hypothesis for the full CRTS sample, as presented in G15, is in tension with PTA measurements. To get some insight on how this candidate sample can be reconciled with current PTA upper limits, we turn to Equation (10). This equation shows that h_c is proportional to the square root of the SMBHB merger rate, which is in turn proportional to the number of SMBHB candidates as per Equation (6). An average factor of ≈ 2 suppression (required to make our fiducial Model_True_fid fully consistent with PTA data) is therefore achievable if at least 75% of the systems are not SMBHBs. Note, however, that the contribution to the GWB varies significantly across candidates (see Figure 8 and Table 1). If the five loudest systems are excluded, the GWB implied by the rest of the sample would be consistent with current PTA upper limits even after incompleteness correction (see Figure 8).

If we want to leave the number of SMBHB candidates untouched, the other parameter to look into is the inferred mass of each system. Equation (10) implies in this case that $h_c \propto \mathcal{M}^{5/3}$. A factor $\mathcal{M}^{5/6}$ comes from the square root of $dE/d\ln f$ and another $\mathcal{M}^{5/6}$ contribution comes from Equations (5) and (6). If the candidate masses are lower, the coalescence timescale is longer and the implied merger rate smaller. So, in principle, a mild reduction of the true mass estimate by a mere factor of 1.5 would suffice to reconcile the CRTS sample with PTA data. Alternatively, the mass ratio distribution of all these systems can be severely biased toward $q \ll 1$. However, a scenario in which all SMBHBs have $q \ll 1$ is difficult to accommodate in current galaxy formation models. Moreover, it is not clear whether low mass ratio perturbers would result in a luminosity modulation as large as 50%, as observed in the CRTS sample (D’Orazio et al. 2016).

Alternatively, it has been proposed that the optical variability of SMBHB may be related to the appearance of an $m = 1$ mode at the inner edge of the circumbinary disk (D’Orazio et al. 2015a; Miranda et al. 2017). In this interpretation, however, the binary orbital frequency would be ≈ 5 times higher than the observed variability in the CRTS sample, implying typical coalescence rates a factor of $5^{8/3} \approx 70$ larger and an associated GWB louder by almost an order of magnitude. It is clear that this interpretation of the CRTS candidates is not viable. Conversely, C16 mentioned the possibility that the periodicity might be related to higher harmonics of the orbital periods.

This would of course imply much longer periods and coalescence times, making PTA GWB upper limits less constraining.

A similar analysis applied to the PTF sample identified by C16, yields comparable results. When the sample is corrected for incompleteness and virial mass bias, the resulting GWB is inconsistent with PTA upper limits at 4.5σ , 3.6σ , and 3.0σ , depending on the postulated q distribution. Conversely, the alternative model proposed by C16, whereby all candidates have $q \approx 0.01$, is consistent with current PTA constraints. This is justified under the assumption that periodicity is due to Doppler boosting, which introduces a bias toward such small q . We note, however, that in this interpretation the GWB has to be corrected for the missing fraction of SMBHBs with larger mass ratios. Alternatively, one has to put forward a model whereby only binaries with $q \approx 0.01$ form, which would be unexpected in the current framework of SMBH assembly.

In summary, we conclude that both the CRTS and the PTF candidate samples are in moderate tension with the current PTA measurements. Possibilities to alleviate this tension preserving the binary hypothesis include (i) even after being corrected for the known virial mass estimate bias, typical true SMBH masses have been overestimated by a factor of ($\gtrsim 1.5$), (ii) the typical mass ratio is lower than theoretically expected ($q \lesssim 0.1$), or else (iii) that the loudest GW sources are preferentially false positives.

While these results question the viability of SMBHB identification via periodicity studies alone, they demonstrate the status of PTAs as important astrophysical probes. Even without a direct GWB detection, PTA upper limits can put stringent constraints on interesting candidate objects. Further studies of these systems are required to identify the true origin of their periodic variability.

We thank Scott Tremaine and Maria Charisi for useful discussions. A.S. is supported by a University Research Fellowship of the Royal Society. Financial support was provided by NASA through ATP grant NNX15AB19G and ADAP grant NNX17AL82G and by NSF grant 1715661 (Z.H.). Z.H. also gratefully acknowledges support from a Simons Fellowship in Theoretical Physics. This project has received funding from the European Research Council (ERC) under the European Union’s Horizon 2020 research and innovation programme (grant agreement No. 638435; GalNUC) and from the Hungarian National Research, Development, and Innovation Office grant NKFIH KH-125675 (to B.K.).

ORCID iDs

Zoltán Haiman  <https://orcid.org/0000-0003-3633-5403>

References

- Artymowicz, P., & Lubow, S. H. 1996, *ApJ*, **467**, L77
- Arzoumanian, Z., Brazier, A., Burke-Spolaor, S., et al. 2014, *ApJ*, **794**, 141
- Arzoumanian, Z., Brazier, A., Burke-Spolaor, S., et al. 2015, *ApJ*, **813**, 65
- Arzoumanian, Z., Brazier, A., Burke-Spolaor, S., et al. 2016, *ApJ*, **821**, 13
- Babak, S., Petiteau, A., Sesana, A., et al. 2016, *MNRAS*, **455**, 1665
- Barausse, E. 2012, *MNRAS*, **423**, 2533
- Begelman, M. C., Blandford, R. D., & Rees, M. J. 1980, *Natur*, **287**, 307
- Berczik, P., Merritt, D., Spurzem, R., & Bischof, H.-P. 2006, *ApJ*, **642**, L21
- Bonetti, M., Sesana, A., Barausse, E., & Haardt, F. 2017, *MNRAS*, submitted (arXiv:1709.06095)
- Capelo, P. R., Volonteri, M., Dotti, M., et al. 2015, *MNRAS*, **447**, 2123
- Charisi, M., Bartos, I., Haiman, Z., et al. 2016, *MNRAS*, **463**, 2145

- Comerford, J. M., Schluns, K., Greene, J. E., & Cool, R. J. 2013, *ApJ*, **777**, 64
- Croton, D. J., Springel, V., White, S. D. M., et al. 2006, *MNRAS*, **365**, 11
- Cuadra, J., Armitage, P. J., Alexander, R. D., & Begelman, M. C. 2009, *MNRAS*, **393**, 1423
- D’Orazio, D. J., Haiman, Z., Duffell, P., Farris, B. D., & MacFadyen, A. I. 2015a, *MNRAS*, **452**, 2540
- D’Orazio, D. J., Haiman, Z., Duffell, P., MacFadyen, A., & Farris, B. 2016, *MNRAS*, **459**, 2379
- D’Orazio, D. J., Haiman, Z., & MacFadyen, A. 2013, *MNRAS*, **436**, 2997
- D’Orazio, D. J., Haiman, Z., & Schiminovich, D. 2015b, *Natur*, **525**, 351
- D’Orazio, D. J., & Loeb, A. 2017, *ApJ*, submitted (arXiv:1712.02362)
- Decarli, R., Dotti, M., Fumagalli, M., et al. 2013, *MNRAS*, **433**, 1492
- Desvignes, G., Caballero, R. N., Lentati, L., et al. 2016, *MNRAS*, **458**, 3341
- Dotti, M., Sesana, A., & Decarli, R. 2012, *AdAst*, 2012, 940568
- Drake, A. J., Djorgovski, S. G., Mahabal, A., et al. 2009, *ApJ*, **696**, 870
- Eddington, A. S. 1913, *MNRAS*, **73**, 359
- Enoki, M., & Nagashima, M. 2007, *PThPh*, **117**, 241
- Eracleous, M., Boroson, T. A., Halpern, J. P., & Liu, J. 2012, *ApJS*, **201**, 23
- Escala, A., Larson, R. B., Coppi, P. S., & Mardones, D. 2005, *ApJ*, **630**, 152
- Farris, B. D., Duffell, P., MacFadyen, A. I., & Haiman, Z. 2014, *ApJ*, **783**, 134
- Farris, B. D., Duffell, P., MacFadyen, A. I., & Haiman, Z. 2015, *MNRAS*, **446**, L36
- Foster, R. S., & Backer, D. C. 1990, *ApJ*, **361**, 300
- Goicovic, F. G., Sesana, A., Cuadra, J., & Stasyszyn, F. 2016, arXiv:1602.01966
- Graham, M. J., Djorgovski, S. G., Stern, D., et al. 2015, *MNRAS*, **453**, 1562
- Gültekin, K., Richstone, D. O., Gebhardt, K., et al. 2009, *ApJ*, **698**, 198
- Guo, Q., White, S., Angulo, R. E., et al. 2013, *MNRAS*, **428**, 1351
- Haiman, Z., Kocsis, B., & Menou, K. 2009, *ApJ*, **700**, 1952
- Hayasaki, K. 2009, *PASJ*, **61**, 65
- Hayasaki, K., Mineshige, S., & Sudou, H. 2007, *PASJ*, **59**, 427
- Hennawi, J. F., Myers, A. D., Shen, Y., et al. 2010, *ApJ*, **719**, 1672
- Hirose, S., Blaes, O., & Krolik, J. H. 2009, *ApJ*, **704**, 781
- Hopkins, P. F., Richards, G. T., & Hernquist, L. 2007, *ApJ*, **654**, 731
- Ivanov, P. B., Papaloizou, J. C. B., & Polnarev, A. G. 1999, *MNRAS*, **307**, 79
- Jaffe, A. H., & Backer, D. C. 2003, *ApJ*, **583**, 616
- Jiang, Y.-F., Stone, J. M., & Davis, S. W. 2013, *ApJ*, **778**, 65
- Ju, W., Greene, J. E., Rafikov, R. R., Bickerton, S. J., & Badenes, C. 2013, *ApJ*, **777**, 44
- Kass, R. E., & Raftery, A. E. 1995, *J. Am. Stat. Assoc.*, **90**, 773
- Kauffmann, G., & Haehnelt, M. 2000, *MNRAS*, **311**, 576
- Kelley, L. Z., Blecha, L., & Hernquist, L. 2017a, *MNRAS*, **464**, 3131
- Kelley, L. Z., Blecha, L., Hernquist, L., Sesana, A., & Taylor, S. R. 2017b, *MNRAS*, **471**, 4508
- Kelly, B. C., Becker, A. C., Sobolewska, M., Siemiginowska, A., & Uttley, P. 2014, *ApJ*, **788**, 33
- Kelly, B. C., Vestergaard, M., & Fan, X. 2009, *ApJ*, **692**, 1388
- Khan, F. M., Just, A., & Merritt, D. 2011, *ApJ*, **732**, 89
- Kocsis, B., & Sesana, A. 2011, *MNRAS*, **411**, 1467
- Komossa, S., & Zensus, J. A. 2016, in IAU Symp. 312, *Star Clusters and Black Holes in Galaxies across Cosmic Time*, ed. Y. Meiron et al. (Cambridge: Cambridge Univ. Press), 13
- Kormendy, J., & Ho, L. C. 2013, *ARA&A*, **51**, 511
- Kozłowski, S. 2017, *A&A*, **597**, A128
- Lentati, L., Taylor, S. R., Mingarelli, C. M. F., et al. 2015, *MNRAS*, **453**, 2576
- Lu, K.-X., Li, Y.-R., Bi, S.-L., & Wang, J.-M. 2016, *MNRAS*, **459**, L124
- Lusso, E., Hennawi, J. F., Comastri, A., et al. 2013, *ApJ*, **777**, 86
- Lynden-Bell, D., Faber, S. M., Burstein, D., et al. 1988, *ApJ*, **326**, 19
- MacFadyen, A. I., & Milosavljević, M. 2008, *ApJ*, **672**, 83
- MacLeod, C. L., Ivezić, Ž., Kochanek, C. S., et al. 2010, *ApJ*, **721**, 1014
- Mayer, L. 2013, *CQGra*, **30**, 244008
- Middleton, H., Chen, S., Del Pozzo, W., Sesana, A., & Vecchio, A. 2017, arXiv:1707.00623
- Miranda, R., Muñoz, D. J., & Lai, D. 2017, *MNRAS*, **466**, 1170
- Nixon, C. J., Cossins, P. J., King, A. R., & Pringle, J. E. 2011, *MNRAS*, **412**, 1591
- Noble, S. C., Mundim, B. C., Nakano, H., et al. 2012, *ApJ*, **755**, 51
- Peters, P. C. 1964, *PhRv*, **136**, 1224
- Peterson, B. M. 2014, *SSRv*, **183**, 253
- Phinney, E. S. 2001, arXiv:astro-ph/0108028
- Planck Collaboration XIII 2016, *A&A*, **594**, A13
- Preto, M., Berentzen, I., Berczik, P., & Spurzem, R. 2011, *ApJL*, **732**, L26
- Quinlan, G. D. 1996, *NewA*, **1**, 35
- Rajagopal, M., & Romani, R. W. 1995, *ApJ*, **446**, 543
- Rau, A., Kulkarni, S. R., Law, N. M., et al. 2009, *PASP*, **121**, 1334
- Ravi, V., Wyithe, J. S. B., Shannon, R. M., & Hobbs, G. 2015, *MNRAS*, **447**, 2772
- Ravi, V., Wyithe, J. S. B., Shannon, R. M., Hobbs, G., & Manchester, R. N. 2014, *MNRAS*, **442**, 56
- Reardon, D. J., Hobbs, G., Coles, W., et al. 2016, *MNRAS*, **455**, 1751
- Roedig, C., Dotti, M., Sesana, A., Cuadra, J., & Colpi, M. 2011, *MNRAS*, **415**, 3033
- Roedig, C., Sesana, A., Dotti, M., et al. 2012, *A&A*, **545**, A127
- Runnoe, J. C., Eracleous, M., Pennell, A., et al. 2017, *MNRAS*, **468**, 1683
- Ryu, T., Perna, R., Haiman, Z., Ostriker, J. P., & Stone, N. C. 2018, *MNRAS*, **473**, 3410
- Sesana, A. 2010, *ApJ*, **719**, 851
- Sesana, A. 2013a, *CQGra*, **30**, 224014
- Sesana, A. 2013b, *MNRAS*, **433**, L1
- Sesana, A., & Khan, F. M. 2015, *MNRAS*, **454**, L66
- Sesana, A., Roedig, C., Reynolds, M. T., & Dotti, M. 2012, *MNRAS*, **420**, 860
- Sesana, A., Shankar, F., Bernardi, M., & Sheth, R. K. 2016, *MNRAS*, **463**, L6
- Sesana, A., Vecchio, A., & Colacino, C. N. 2008, *MNRAS*, **390**, 192
- Sesana, A., Vecchio, A., & Volonteri, M. 2009, *MNRAS*, **394**, 2255
- Shakura, N. I., & Sunyaev, R. A. 1973, *A&A*, **24**, 337
- Shankar, F., Bernardi, M., Sheth, R. K., et al. 2016, *MNRAS*, **460**, 3119
- Shannon, R. M., Ravi, V., Lentati, L. T., et al. 2015, *Sci*, **349**, 1522
- Shen, Y., Greene, J. E., Strauss, M. A., Richards, G. T., & Schneider, D. P. 2008, *ApJ*, **680**, 169
- Shen, Y., & Kelly, B. C. 2012, *ApJ*, **746**, 169
- Shen, Y., Liu, X., Loeb, A., & Tremaine, S. 2013, *ApJ*, **775**, 49
- Shi, J.-M., & Krolik, J. H. 2015, *ApJ*, **807**, 131
- Shi, J.-M., Krolik, J. H., Lubow, S. H., & Hawley, J. F. 2012, *ApJ*, **749**, 118
- Simon, J., & Burke-Spolaor, S. 2016, *ApJ*, **826**, 11
- Springel, V., White, S. D. M., Jenkins, A., et al. 2005, *Natur*, **435**, 629
- Syer, D., & Clarke, C. J. 1995, *MNRAS*, **277**, 758
- Tang, Y., MacFadyen, A. I., & Haiman, Z. 2017, *MNRAS*, **469**, 4258
- Tsalmantza, P., Decarli, R., Dotti, M., & Hogg, D. W. 2011, *ApJ*, **738**, 20
- Vasiliev, E., Antonini, F., & Merritt, D. 2015, *ApJ*, **810**, 49
- Vaughan, S., Uttley, P., Markowitz, A. G., et al. 2016, *MNRAS*, **461**, 3145
- Verbiest, J. P. W., Lentati, L., Hobbs, B., et al. 2016, *MNRAS*, **458**, 1267
- Wang, L., Greene, J. E., Ju, W., et al. 2017, *ApJ*, **834**, 129
- White, S. D. M., & Rees, M. J. 1978, *MNRAS*, **183**, 341
- Wyithe, J. S. B., & Loeb, A. 2003, *ApJ*, **590**, 691
- Zhu, X.-J., Hobbs, G., Wen, L., et al. 2014, *MNRAS*, **444**, 3709