

Accurate models for recoil velocity distribution in black hole mergers with comparable to extreme mass-ratios and their astrophysical implications

Tousif Islam^{1,*} and Digvijay Wadekar^{2,3}

¹*Kavli Institute for Theoretical Physics, University of California Santa Barbara, Kohn Hall, Lagoon Rd, Santa Barbara, CA 93106*

²*Department of Physics and Astronomy, Johns Hopkins University,
3400 N. Charles Street, Baltimore, Maryland, 21218, USA*

³*Weinberg Institute, University of Texas at Austin, Austin, TX 78712, USA*

(Dated: November 17, 2025)

Modeling the remnant recoil velocity (kick) distribution from binary black hole mergers is crucial for understanding hierarchical mergers in active galactic nuclei or globular clusters. Existing analytic models often show large discrepancies with numerical relativity (NR) data, while data-driven models are limited to mass ratios of $q \leq 8$ (aligned spins) and $q \leq 4$ (precessing spins) and break down when extrapolated outside their training ranges. Using ~ 5000 of NR simulations from the SXS and RIT catalogs up to $q = 128$ and ~ 100 black hole perturbation theory simulations up to $q = 200$, we present two classes of models: (i) `gwModel_kick_q200` (`gwModel_kick_q200_GPR`), an analytic (Gaussian process regression) model for aligned-spin binaries. (ii) `gwModel_kick_prec_flow`, a normalizing-flow model for kick distribution from precessing binaries with isotropic spins. Our approach combines analytic insights from post-Newtonian theory with data-driven techniques to ensure correct limiting behavior and high accuracy across parameter space. Both `gwModel_kick_q200` and `gwModel_kick_prec_flow` are valid from comparable to extreme mass ratios. Extensive validation shows all three models outperform existing ones within their respective domains. Finally, using both back-of-the-envelope estimates and 1404 detailed star cluster simulations incorporating our kick models, we find that the black hole retention probability in low mass globular clusters can vary noticeably when the `gwModel_kick_prec_flow` model is employed. The models are publicly available through the `gwModels`^a package.

I. INTRODUCTION

One of the most intriguing features of binary black hole (BBH) mergers is that the remnant black hole (BH) receives a recoil velocity—commonly referred to as a ‘kick’. Over the past two decades, numerous numerical and analytic studies have been conducted to investigate these kick velocities using several complementary frameworks, including numerical relativity (NR) [1–25], black hole perturbation theory (BHPT) [26–31], and post-Newtonian (PN) [15, 32–38] approximations. Results from NR simulations have demonstrated that the kick velocity can reach values as high as ~ 5000 km/s for certain spin orientations [39–46]. Recently, for more finely tuned spin configurations, kick velocities as large as $\sim 30,000$ km/s have been reported [47].

These recoil kicks have significant astrophysical implications. When two BHs merge, the magnitude of the kick determines whether the remnant remains to its host environment, such as a globular cluster, nuclear star cluster, or active galactic nucleus, or escapes it. If the kick velocity exceeds the escape velocity of the host, the remnant is ejected and can no longer participate in subsequent mergers within that environment [48, 49]. Conversely, if the kick is smaller, the remnant BH remains within its host and may undergo additional mergers with other BHs. Such successive or ‘hierarchical’ mergers provide a natural pathway for the formation of increasingly massive BHs [9, 48–53]. BHs formed through hierarchical mergers are expected to exhibit distinct mass and spin distributions compared to those originating from the direct collapse

of massive stars [52, 54, 55]. Interestingly, several BH candidates detected through gravitational-wave (GW) observations show properties consistent with this hierarchical formation channel [56]. Furthermore, for a few GW events, the remnant kick has already been inferred to be as large as ~ 1500 km/s.

Understanding the population of BHs formed through hierarchical mergers is also crucial for modeling the dynamical evolution of astrophysical systems such as globular clusters and active galactic nuclei (AGN) [57–60]. Such an understanding relies heavily on the availability of accurate models for the kick velocities produced in BBH mergers. Significant progress has been made in modeling kick velocities over the past decade [3, 12, 13, 19, 20, 24, 27, 28, 38, 61–66]. Nevertheless, a substantial modeling gap remains, particularly in accurately predicting kick velocities across the full parameter space of mass ratios and spin configurations.

Currently, the state-of-the-art models for kick velocities include four main frameworks. `NRSur7dq4Remnant` [63] and `NRSur3dq8Remnant` [64] are widely used NR surrogate models for precessing-spin and aligned-spin binaries, respectively. They are trained on NR simulations from the SXS catalog [67, 68] up to mass ratios of $q = 4$ and $q = 8$ respectively, although they are often used beyond their nominal training ranges, up to $q \sim 6$ and $q \sim 10$ respectively. We define the mass ratio $q := m_1/m_2$ where m_1 (m_2) is the mass of the larger (smaller) BH. On the other hand, `BHPTNRSurRemnant` [28] is a remnant-kick model based on BHPT simulations, trained over mass ratios in the range $3 \leq q \leq 1000$ and calibrated against NR data, but applicable only to non-spinning binaries. All three of these models are based on Gaussian process regression (GPR) and are therefore susceptible to extrapolation errors outside their training domains. In contrast, an analytic model, commonly referred to as the `HLZ` model [12, 13, 19, 20, 45], was

* tousifislam@ucsb.edu

^a <https://github.com/tousifislam/gwModels>




Figure 1. We show the kick-velocity predictions from the analytic `gwModel_kick_q200` model (blue solid line) along with the corresponding 95% confidence interval (blue shaded region) in the nonspinning limit, as a function of mass ratio from $q = 1$ to $q = 256$. For comparison, we include a combination of NR data from the SXS (circles) and RIT (squares) catalogs, as well as BHPT results (triangles). The red, orange, purple, and green lines correspond to predictions from the `NRSur7dq4Remnant`, `NRSur3dq8Remnant`, `BHPTNRSurRemnant`, and `HLZ` models, respectively, all shown in their nonspinning limits. The red and orange dashed lines indicate extrapolations of the `NRSur7dq4Remnant` and `NRSur3dq8Remnant` models beyond their region of validity. The inset shows the residuals of the `gwModel_kick_q200` predictions. More details are in Section II C.

developed using NR data from the RIT catalog [69, 70], combined with insights from PN theory and BHPT limits. While this model is, in principle, valid across a wide range of mass ratios and spin configurations, it tends to yield larger errors when compared to NR results in the comparable-mass regime.

In this paper, we present new models for the kick velocity that bridge the existing gap in the parameter-space coverage and improves model accuracy. We combine approximately 5000 NR simulations spanning mass ratios $q \in [1, 128]$ from both the SXS [67, 68, 71] and RIT catalogs [69, 70], together with about 100 BHPT simulations covering $q \in [40, 200]$. Our BHPT simulations are obtained using the time-domain Teukolsky solver developed in Refs. [27, 72–77]. These simulations are based on linear perturbation theory, where the smaller BH undergoes an adiabatic inspiral driven by the loss of energy and angular momentum through gravitational radiation. Recent studies have demonstrated that such simulations can reproduce NR results with high accuracy in the comparable-mass regime after minimal calibration and are expected to remain reliable in the intermediate-mass-ratio range ($q \geq 15$) [78–84].

By fusing analytic insights with data-driven techniques such as GPR and normalizing flows, we construct unified kick-velocity models that are valid from the comparable-mass to the extreme-mass-ratio regime. Specifically, we introduce two models for aligned-spin binaries: an analytic model, `gwModel_kick_q200`, and a data-driven model, `gwModel_kick_q200_GPR`. Both models are trained on a hybrid dataset combining NR and BHPT results across the parameter space $q \in [1, 200]$ and spin magnitudes $\chi_{1,2} \in [0, 1]$. Furthermore, we present a normalizing-flow-based probabilistic model for precessing-spin binaries, `gwModel_kick_prec_flow`, trained over the range $q \in [1, 100]$ and $\chi_{1,2} \in [0, 1]$. Details of the aligned-spin mod-

els are provided in Section II, while the precessing-spin model is described in Section III. We find that our models exhibit excellent extrapolation behavior and can be used even in the extreme mass ratio limit. We discuss the implications of these new models for cluster dynamics and hierarchical BH mergers in Section IV. Finally, we outline our future directions in Section V. All models presented in this work are publicly available through the `gwModels` package.

II. ALIGNED-SPIN MODELS

First, we construct two kick models for aligned-spin binaries that are valid up to a mass ratio of $q = 200$ for spin magnitudes $\chi_{1,2} \in [0, 1]$. One of the models is data-driven, while the other is analytic. We refer to the analytic model as `gwModel_kick_q200`, and to the data-driven model as `gwModel_kick_q200_GPR`, as it is trained using GPR.

Our training dataset consists of NR simulations from the SXS and RIT catalogs, together with BHPT waveforms. Specifically, we use 494 aligned-spin NR simulations from the SXS catalog with mass ratios up to $q = 20$, including 49 nonspinning cases. We also include 9 aligned-spin RIT simulations at $q = 16$ and $q = 32$, as well as 16 nonspinning RIT simulations spanning mass ratios from $q = 1.08$ to $q = 128$. In addition, we incorporate BHPT data at $q = 200$ for spin values $|\chi_1| = [0.8, 0.4, 0.0, -0.4, -0.8]$. In the nonspinning limit, multiple NR simulations are often available at the same mass ratio; in those cases, we take the average kick velocity and compute the standard deviation to estimate an uncertainty band. For the NR data, we use the kick velocities provided directly in the corresponding catalogs, where they are computed from either horizon quantities or waveform fluxes. For the BHPT

data, we compute the kick velocities using the `gw_remnant`¹ package [28] that relies on waveform fluxes.

To develop our models, we adopt the fitting ansatz proposed in Ref. [24], which includes additional terms compared to the formulations in Refs. [19, 20]. We first introduce a set of auxiliary variables. We define the symmetric mass ratio as $\eta = q/(1+q)^2$ and the mass-difference parameter as $\delta m = (q-1)/(q+1)$. Furthermore, we introduce two PN-inspired auxiliary quantities, the mass-weighted effective spin parameters:

$$\tilde{S}_{\parallel} = \frac{\chi_{1z} + q^2 \chi_{2z}}{(1+q)^2}, \quad \Delta_{\parallel} = \frac{\chi_{1z} - q \chi_{2z}}{1+q}. \quad (1)$$

Additionally, we define the effective spin parameters as

$$\chi_{\text{eff}} = \frac{q \chi_{1z} + \chi_{2z}}{1+q}, \quad \hat{\chi} = \frac{\chi_{\text{eff}} - 38 \eta (\chi_{1z} + \chi_{2z})/113}{1 - 76 \eta/113}, \quad (2)$$

and the anti-symmetric spin as

$$\chi_a = \frac{1}{2}(\chi_{1z} - \chi_{2z}). \quad (\text{S6})$$

A. GPR model for aligned-spin kicks

For the GPR model, we use the `scikit-learn` implementation with an optimized kernel of the form

$$k(\mathbf{x}, \mathbf{x}') = 0.939^2 k_{\text{RBF}}(\mathbf{x}, \mathbf{x}') + k_{\text{white}}(\mathbf{x}, \mathbf{x}'), \quad (3)$$

where the squared exponential (RBF) component is defined as

$$k_{\text{RBF}}(\mathbf{x}, \mathbf{x}') = \exp \left[-\frac{1}{2} \sum_{i=1}^3 \frac{(x_i - x'_i)^2}{\ell_i^2} \right], \quad (4)$$

with $\ell = [0.619, 2.57, 0.602]$. Here, ℓ denotes the characteristic length scales of the input parameters $\mathbf{x} = \{\log_2(q), \hat{\chi}, \chi_a\}$, which control how rapidly the correlation decays along each dimension. The amplitude term 0.939^2 sets the overall variance of the process. The white-noise component is given by

$$k_{\text{white}}(\mathbf{x}, \mathbf{x}') = 0.00187 \delta_{\mathbf{x}\mathbf{x}'}, \quad (5)$$

where $\delta_{\mathbf{x}\mathbf{x}'}$ is the Kronecker delta. The white-noise term accounts for uncorrelated statistical scatter or numerical noise in the training data. In this formulation, $\mathbf{x}$ and $\mathbf{x}'$ represent two points in the input feature space, and the kernel $k(\mathbf{x}, \mathbf{x}')$ quantifies their similarity, determining how strongly the corresponding outputs are correlated under the GPR model.

B. Analytic model for aligned-spin kicks

GPR models typically perform poorly when extrapolated beyond the range of the training data. To ensure reliable extrapolation and impose known physical constraints, such as




Figure 2. (Upper panel) We show the histogram of errors in kick-velocity predictions from `gwModel_kick_q200` (blue) and `gwModel_kick_q200.GPR` (black), computed with respect to NR and BHPT results for aligned-spin binaries across five different models in their aligned-spin limit. The NR error (gray) is estimated by comparing simulations at two different numerical resolutions when available. We find that `gwModel_kick_q200` performs as well as, or in some cases better than, the NR surrogate models `NRSur7dq4Remnant` (orange) and `NRSur3dq8Remnant` (red) within their respective training domains, and shows significantly improved agreement with the data compared to the `HLZ` model (green). The errors from our model are comparable to the intrinsic NR uncertainties. (Lower panel) We show the kick-velocity distribution as a function of the mass ratio q and the effective inspiral spin χ_{eff} . Dashed lines indicate regions where additional data are available from SXS NR, RIT NR, and BHPT simulations. More details are in Section II C.

the requirement of zero recoil in the equal-mass limit, we place greater emphasis on the analytic model. Following Refs. [12, 13, 19, 20, 45], we decompose the total kick velocity in aligned-spin binaries into two contributions:

$$V_{\text{kick,AS}}(q, \chi_{1z}, \chi_{2z}) = \sqrt{V_m^2 + V_{\perp}^2 + 2 V_m V_{\perp} \cos \xi}, \quad (6)$$

where V_m denotes the contribution due to mass asymmetry, $V_{\perp}$ is the in-plane recoil arising from spins aligned with the orbital angular momentum, and ξ is the angle between the unequal-mass and spin contributions to the kick in the orbital plane. The mass-asymmetry contribution is modeled as (following

¹ https://github.com/tousifislam/gw_remnant

Refs. [12, 13, 19, 20, 33])

$$V_m = A \eta^2 \delta m (1 + B \eta + C \eta^2), \quad (7)$$

with $A = (1.3753 \pm 0.0245) \times 10^4$ km/s, $B = -2.636 \pm 0.168$, and $C = 5.433 \pm 0.531$. The spin-induced contribution to the kick, arising from spins aligned with the orbital angular momentum axis, is modeled using the following polynomial expansion adopted from Ref. [24]:

$$\begin{aligned} V_{\perp} = & H \eta^2 (\Delta_{\parallel} + H_{2a} \tilde{S}_{\parallel} \delta m + H_{2b} \Delta_{\parallel} \tilde{S}_{\parallel} + H_{3a} \Delta_{\parallel}^2 \delta m \\ & + H_{3b} \tilde{S}_{\parallel}^2 \delta m + H_{3c} \Delta_{\parallel} \tilde{S}_{\parallel}^2 + H_{3d} \Delta_{\parallel}^3 + H_{3e} \Delta_{\parallel} \delta m^2 \\ & + H_{4a} \tilde{S}_{\parallel} \Delta_{\parallel}^2 \delta m + H_{4b} \tilde{S}_{\parallel}^3 \delta m + H_{4c} \tilde{S}_{\parallel} \delta m^3 \\ & + H_{4d} \Delta_{\parallel} \tilde{S}_{\parallel} \delta m^2 + H_{4e} \Delta_{\parallel} \tilde{S}_{\parallel}^3 + H_{4f} \tilde{S}_{\parallel} \Delta_{\parallel}^3), \end{aligned} \quad (8)$$

The best-fit coefficients for the in-plane recoil, obtained using the `scipy.curve_fit` [85] routine, are $H = (7.402 \pm 0.028) \times 10^3$ km s⁻¹, $H_{2a} = 5.891 \pm 0.042$, $H_{2b} = -0.737 \pm 0.019$, $H_{3a} = -0.608 \pm 0.040$, $H_{3b} = -1.407 \pm 0.098$, $H_{3c} = -1.462 \pm 0.078$, $H_{3d} = (-1.67 \pm 0.66) \times 10^{-2}$, $H_{3e} = 6.768 \pm 0.080$, $H_{4a} = -0.834 \pm 0.106$, $H_{4b} = -2.768 \pm 0.148$, $H_{4c} = 3.655 \pm 0.099$, $H_{4d} = -2.099 \pm 0.145$, $H_{4e} = 1.145 \pm 0.176$, and $H_{4f} = 0.283 \pm 0.054$. The angle ξ is given by:

$$\xi = a + b \tilde{S}_{\parallel} + c \delta m \Delta_{\parallel}, \quad (9)$$

with $a = 146.30^\circ \pm 0.37^\circ$, $b = 111.92^\circ \pm 1.80^\circ$, and $c = 136.36^\circ \pm 2.94^\circ$. Note that these coefficients are obtained from a 5-fold cross-validation procedure and are significantly different from the ones used in Ref. [24]. Furthermore, in Ref. [24], the coefficients entering Eq. (8) were fitted independently, while those appearing in Eq. (7) were adopted from Refs. [12, 13, 19, 20, 33]. In contrast, in this work, we fit all coefficients simultaneously.

C. Model accuracy for aligned-spin kicks

In Figure 1, we show the kick-velocity predictions in the non-spinning limit from the `gwModel_kick_q200` model, along with the corresponding 95% confidence interval, and compare them against a set of NR and BHPT data as a function of mass ratio ranging from $q = 1$ to $q = 256$. The 95% confidence limits are derived from the covariance matrix of the fitted parameters, thereby quantifying the uncertainty in the analytical fit. For comparison, we include predictions from four state-of-the-art models: `NRSur7dq4Remnant`, `NRSur3dq8Remnant`, `BHPTNRSurRemnant`, and `HLZ`. We find that both `gwModel_kick_q200` and `HLZ` reproduce the NR and BHPT data accurately across the entire mass-ratio range. The `NRSur7dq4Remnant` model agrees well with the data up to $q = 4$ (within its training domain); although it is sometimes used up to $q = 6$, Fig. 1 shows that its predictions begin to deviate beyond $q = 4$ and become unreliable for $q > 6$, indicating poor extrapolation capability. Interestingly, `NRSur3dq8Remnant` (an aligned-spin model) fails to reproduce the NR results between $q = 1.8$ and $q = 4$, missing the




Figure 3. We show the predictions from `gwModel_kick_q200` (solid lines) and `gwModel_kick_q200_GPR` (dashed lines) for five different aligned-spin configurations as a function of mass ratio, ranging from $q = 1$ to $q = 256$. We find that the analytic model `gwModel_kick_q200` exhibits smooth and physically consistent behavior across the full mass-ratio range, with the kick velocity decreasing monotonically with increasing q . In contrast, the data-driven `gwModel_kick_q200_GPR` model produces several unphysical features, even though it achieves a slightly lower validation error. Similar unphysical trends are also observed in one-dimensional parameter-space slices for other GPR-based models, such as `NRSur7dq4Remnant` and `NRSur3dq8Remnant`. An example of such behavior for `NRSur3dq8Remnant` is shown in Fig. 1. More details are in Section II C.

well-known peak in the nonspinning kick. Furthermore, its performance degrades significantly for $q > 8$ (the upper limit of its training range). Although the model is often used up to $q = 10$, Fig. 1 suggests that its predictions are inaccurate in the range $8 \lesssim q \lesssim 10$. Finally, `BHPTNRSurRemnant` fails to match the data for $q \lesssim 25$. Overall, in the nonspinning limit, `gwModel_kick_q200` provides the most accurate predictions among all models considered. The associated validation residuals, shown in the inset, remain mostly below ~ 25 km/s.

Next, in Figure 2 (upper panel), we show the model prediction errors with respect to NR and BHPT data. For `gwModel_kick_q200`, `gwModel_kick_q200_GPR`, and `HLZ`, we compute the errors against all available data up to $q = 200$ (shown in the lower panel of Fig. 2). For `NRSur7dq4Remnant` and `NRSur3dq8Remnant`, we restrict the comparison to data points within their respective training domains, i.e., $q \leq 4$ and $q \leq 8$. We make this choice because, as shown in Fig. 1, the models tend to be inaccurate in the mass-ratio ranges where they are commonly extrapolated, i.e. $4 \leq q \leq 6$ and $8 \leq q \leq 10$. For reference, we also show the numerical errors in NR data, obtained by comparing kick velocities from simulations performed at different numerical resolutions whenever available. We find that both `gwModel_kick_q200` and `gwModel_kick_q200_GPR` outperform all other models, with `gwModel_kick_q200_GPR` showing slightly better over-

all accuracy. Furthermore, the errors in these two models are comparable to the intrinsic NR uncertainties. For `gwModel_kick_q200` and `gwModel_kick_q200_GPR`, the predictions shown correspond to results from a 5-fold cross-validation procedure. The coefficient of determination (R^2) values for different models are as follows: 0.9958 for NR resolution errors, 0.9919 for `gwModel_kick_q200`, 0.9987 for `gwModel_kick_q200_GPR`, 0.9730 for `NRSur3dq8Remnant`, 0.9948 for `NRSur7dq4Remnant`, and 0.5553 for `HLZ`.

Even though the `gwModel_kick_q200_GPR` model performs slightly better than the analytic model `gwModel_kick_q200` in terms of the validation coefficient of determination, **we choose the analytic model as our default**. This choice is motivated by the fact that in the three-dimensional parameter space, several regions are sparsely populated with data. Although these regions formally lie within the training range, all models effectively extrapolate there. However, GPR-based models (such as `gwModel_kick_q200_GPR`, `NRSur7dq4Remnant`, and `NRSur3dq8Remnant`) can often yield unreliable or nonphysical predictions in these data-sparse regions.

We illustrate this issue in Figure 3, where we show predictions from `gwModel_kick_q200` and `gwModel_kick_q200_GPR` for five different aligned-spin configurations, some of which are not present in the training or validation sets, across mass ratios ranging from $q = 1$ to $q = 256$. We find that predictions from the analytic model `gwModel_kick_q200` vary smoothly and decrease with increasing mass ratio, as physically expected. In contrast, the `gwModel_kick_q200_GPR` model occasionally exhibits unphysical oscillations at large mass ratios and irregular behavior even in the comparable-mass regime. Similar artifacts are observed in the `NRSur7dq4Remnant` and `NRSur3dq8Remnant` models when inspecting one-dimensional slices of their predictions (e.g. in Fig. 1). These results provide a strong motivation to prefer analytic models whenever feasible.

III. PRECESSING-SPIN MODEL

Next, we focus on precessing-spin binaries. Modeling kicks in such systems is considerably more challenging than in aligned-spin cases, primarily because the parameter space becomes seven-dimensional owing to the four additional in-plane spin orientation angles ($\theta_1, \theta_2, \phi_1$ and ϕ_2). However, for many astrophysical applications—such as population synthesis, galaxy evolution, and cosmological studies—a single point estimate of the kick velocity for a specific precessing configuration is not typically required. Instead, these studies often rely on spin-angle-averaged recoil distributions. Motivated by this, we develop a model that fits the distribution of kick velocities resulting from varying spin orientations, given the mass ratio and spin magnitudes of the two black holes. To achieve this, we employ a probabilistic approach based on normalizing flows.

Our training dataset consists of a total of 2,866 SXS NR simulations covering mass ratios $q \in [1, 15]$, 1,881 RIT NR simulations over the same range, and 400 BHPT simulations




Figure 4. We show the kick-velocity distribution from precessing-spin SXS (blue stars) and RIT (orange circles) NR simulations up to $q = 15$ for different spin magnitudes and orientation configurations. Our training dataset also includes BHPT simulations spanning mass ratios from $q = 40$ to $q = 100$ (not shown). For reference, we plot the empirical upper envelope of the kick velocities as a dashed gray line, given by $V_{\text{upper}} = 7.6 \times 10^4 \eta^2$. The inset shows the distribution of these data points as a function of the spin magnitudes $|\chi_1|$ and $|\chi_2|$. More details are in Section III B.

spanning $q \in [40, 100]$. These datasets collectively sample a wide variety of spin configurations characterized by the spin orientation angles ($\theta_1, \theta_2, \phi_1, \phi_2$) and spin magnitudes ($|\chi_1|, |\chi_2|$). Interestingly, most of the NR data from the SXS and RIT catalogs are concentrated at relatively high spin magnitudes, $\chi_{1,2} \gtrsim 0.7$. By examining all available data, we first constrain an empirical upper bound on the kick velocities (for various spin magnitudes and orientations) as a function of the mass ratio. We express this relation as

$$V_{\text{upper}} = V_{\text{max}} \eta^2, \quad (10)$$

where $V_{\text{max}} = 7.6 \times 10^4$ km/s.

A. Normalizing flow model for precessing-spin kicks

To model the distribution of kick velocities for precessing-spin binaries, we employ a probabilistic normalizing-flow architecture implemented using the `nflows` package [86]. The goal of this model is to learn the conditional probability density

$$p(v_{\text{kick}} | \mathbf{c}), \quad (11)$$

where v_{kick} denotes the kick velocity (the target variable), and $\mathbf{c}$ represents the contextual parameters $\mathbf{c} = \{\log_2(q), |\chi_1|, |\chi_2|\}$ encoding the binary's mass ratio and individual spin magnitudes. For any specified $(q, |\chi_1|, |\chi_2|)$, the model predicts the full probability distribution of kick velocities marginalized over isotropically-oriented spin-angle configurations. This conditional setup enables the model to efficiently capture the




Figure 5. We show the training ($p_{\text{loss}}^{\text{train}}$; crimson solid line) and validation ($p_{\text{loss}}^{\text{val}}$; green dashed line) losses for the precessing-spin normalizing-flow model `gwModel_kick_prec_flow` as a function of training steps. Additionally, we plot the Wasserstein distance (W_p ; black solid line) between the precessing-spin training data and the corresponding `gwModel_kick_prec_flow` outputs at various training stages. The inset shows the distribution of these data points as a histogram (gray), along with the corresponding distribution predicted by the trained normalizing-flow model `gwModel_kick_prec_flow` (blue). More details are in Section III B.

statistical effects of spin precession on the kick distribution while remaining computationally tractable.

Before training, all quantities are standardized to zero mean and unit variance. To avoid discontinuities near the equal-mass boundary ($q = 1$), we duplicate the training samples with transformations ($\log_2 q \rightarrow -\log_2 q$, $|\chi_1| \leftrightarrow |\chi_2|$). This symmetrization ensures that the normalizing flow can learn the distribution smoothly across the parameter space. The resulting dataset is randomly shuffled and split into training (75%) and validation (25%) subsets.

The flow is defined as a composition of $n_{\text{layers}} = 2$ autoregressive transformations, each consisting of a Masked Affine Autoregressive Transform (MAF) preceded by a `ReversePermutation` layer to enhance expressivity:

$$f = f_2 \circ f_1 = (\mathcal{P}_2 \circ \mathcal{T}_2) \circ (\mathcal{P}_1 \circ \mathcal{T}_1), \quad (12)$$

where $\mathcal{P}_i$ denotes a permutation of the input features and $\mathcal{T}_i$ represents the autoregressive affine transformation. The overall flow thus defines an invertible mapping between a base latent variable $\mathbf{z}$ and the data variable v_{kick} :

$$v_{\text{kick}} = f(\mathbf{z}; \mathbf{c}), \quad \mathbf{z} \sim \mathcal{N}(0, I). \quad (13)$$

The conditional likelihood is then given by

$$\log p(v_{\text{kick}}|\mathbf{c}) = \log p_{\mathbf{Z}}(f^{-1}(v_{\text{kick}}; \mathbf{c})) + \log \left| \det \frac{\partial f^{-1}}{\partial v_{\text{kick}}} \right|. \quad (14)$$

Each autoregressive block employs a fully connected multilayer perceptron (MLP) with $d_{\text{hidden}} = 8$ neurons per layer




Figure 6. We show the distribution of kick velocities predicted by three precessing-spin models—`gwModel_kick_prec_flow` (blue), `NRSur7dq4Remnant` (red), and `HLZ` (green)—within their respective domains of validity, for four different mass ratios: $q = [1, 5, 10, 15]$ (with the $q = 1$ results shown in the inset), assuming spin magnitudes of $|\chi_1|, |\chi_2| = [0.6, 0.9]$. More details are in Section III B.

and Gaussian Error Linear Unit (GELU) activation functions. The context features $\mathbf{c}$ are embedded through linear conditioning networks with dimensionality $d_{\text{context}} = 3$. The base distribution is a standard normal, $p_{\mathbf{Z}} = \mathcal{N}(0, I)$.

The model is trained by minimizing the negative log-likelihood over the training set:

$$\mathcal{L} = -\mathbb{E}_{(v_{\text{kick}}, \mathbf{c})} [\log p(v_{\text{kick}}|\mathbf{c})]. \quad (15)$$

We use the Adam optimizer with a learning rate of 10^{-3} and a weight decay of 10^{-4} for regularization. Training is performed for 20,000 iterations with a batch size of 128. During training, we periodically evaluate the validation loss and compute the Wasserstein distance between the generated and target kick distributions to monitor convergence and sample quality. The best model checkpoint is selected based on the minimum validation loss. Note that we keep the flow model intentionally simple, employing only two layers. This design choice helps prevent overfitting and allows the model to learn the underlying physical manifold of the kick-velocity distribution more effectively.

The final model, denoted as `gwModel_kick_prec_flow`, captures the conditional distribution of kick velocities across a wide range of mass ratios ($q \leq 200$) and spin magnitudes. As this model is entirely data-driven, we also provide a simple single-precession model in Appendix 1 for intuition.




Figure 7. We show the Jensen–Shannon divergence (JSD; color-coded) between the kick-velocity distributions obtained for 1000 randomly chosen combinations of spin magnitudes $|\chi_1|$ and $|\chi_2|$. Spin orientation angles are drawn from an isotropic distribution. The left and center panels show the results for mass ratios q within the range $q \in [1, 4]$, where all three models are valid: `gwModel_kick_prec_flow` (referred to as `gwModel`), `NRSur7dq4Remnant` (referred to as `NRSur`), and `HLZ`. The symbol size scales with the mass ratio: smaller circles correspond to smaller q , while larger circles correspond to larger q . The predictions of `gwModel` and `NRSur` agree more closely compared to those of `HLZ`. More details are in Section III B.




Figure 8. Same as Fig. 7 but for more asymmetric mass ratios: $q \in [4, 100]$. We do not compare with `NRSur7dq4Remnant` as this range is beyond its domain of validity. Spin orientation angles are again drawn from an isotropic distribution. The symbol size scales with the mass ratio: smaller circles correspond to smaller q , while larger circles correspond to larger q . More details are in Section III B.

B. Model accuracy for precessing-spin kicks

In Figure 4, we show the kick-velocity distribution from precessing-spin SXS and RIT NR simulations up to $q = 15$, covering a wide range of spin magnitudes and orientation configurations. We also plot the empirical upper envelope of the kick velocities, given by Eq. (10).

In Figure 5, we present the training and validation losses for `gwModel_kick_prec_flow` as a function of training iterations. Both losses decrease steadily and reach a value of approximately 0.5% after about 5000 training steps, indicating

that the model has converged efficiently. The close agreement between the training and validation losses demonstrates that the model does not suffer from overfitting. In addition, we compute the Wasserstein distance between the precessing-spin training data and the corresponding `gwModel_kick_prec_flow` samples at different stages of training. This distance stabilizes after roughly 1000 iterations at a value of ~ 0.1 , further confirming that the trained model accurately reproduces the underlying data distribution. Furthermore, we display the distribution of the kick velocities as a histogram (in the inset), along with the corresponding distribution predicted by the trained normalizing-flow model `gwModel_kick_prec_flow`. We find that the predictions from `gwModel_kick_prec_flow` closely reproduce the training data across the entire parameter space.

Next, in Figure 6, we show the distribution of kick velocities predicted by three precessing-spin models—`gwModel_kick_prec_flow`, `NRSur7dq4Remnant`, and `HLZ`—within their respective domains of validity, for four different mass ratios: $q = [1, 5, 10, 15]$. We fix the spin magnitudes to $|\chi_1|, |\chi_2| = [0.6, 0.9]$ and draw the spin orientation angles from an isotropic distribution. Because `gwModel_kick_prec_flow` was trained using a different spin-orientation prior, we reweight its predictions to match an isotropic spin-orientation distribution. For $q = 1$, all three models produce nearly identical kick-velocity distributions. At $q = 5$, `gwModel_kick_prec_flow` and `NRSur7dq4Remnant` yield consistent results, whereas `HLZ` predicts a noticeably broader distribution, corresponding to larger recoil velocities. For higher mass ratios ($q = 10$ and $q = 15$), `NRSur7dq4Remnant` fails to provide meaningful predictions as these values lie beyond its training range, while the `HLZ` model continues to overestimate the width of the distribution compared to `gwModel_kick_prec_flow`. We also observe that for lower spin magnitudes, the discrepancies between model predictions become more pronounced. This behavior likely arises because the available NR training data in this regime are sparse, causing the models to extrapolate beyond

well-sampled regions of the parameter space.

We explore this comparison further in Figure 7 (first two panels). We sample 1000 different mass ratios within the comparable-mass regime, specifically for $q \in [1, 4]$, and spin magnitudes $|\chi_1|$ and $|\chi_2|$ uniformly distributed between 0 and 1. For each configuration, we compute the kick-velocity distributions obtained from isotropically distributed spin-orientation angles using three precessing-spin models: `gwModel_kick_prec_flow`, `NRSur7dq4Remnant`, and `HLZ`. To quantify the similarity between the kick-velocity distributions predicted by these models, we compute the Jensen–Shannon divergence (JSD). The JSD between two probability distributions p and q is defined as

$$\text{JSD}(p||q) = \frac{1}{2}\text{KL}(p||m) + \frac{1}{2}\text{KL}(q||m), \quad (16)$$

where $m = \frac{1}{2}(p + q)$ is the pointwise mean of p and q and KL denotes the Kullback–Leibler divergence. The JSD is symmetric and bounded as $0 \leq \text{JS}(P||Q) \leq \log 2 \approx 0.693$. Typically, $\text{JS} \lesssim 0.1$ indicates that the two distributions are statistically similar, while values $\text{JS} \gtrsim 0.1$ signify significant discrepancies between them.

We find that, across most of the parameter space, the kick-velocity distributions predicted by `gwModel_kick_prec_flow` and `NRSur7dq4Remnant` agree well, with $\text{JSD} \leq 0.1$ for the majority of sampled points. However, for smaller spin magnitudes, the distributions from these two models begin to deviate. This behavior arises partly because most of the available NR training data correspond to high spin magnitudes, making the models effectively extrapolate in the low-spin regime. As shown earlier, data-driven models (particularly those based on GPR) tend to perform poorly in data-sparse regions. When comparing `gwModel_kick_prec_flow` with `HLZ`, we find that the JSD values are significantly larger than in the previous case, indicating that the kick-velocity distributions predicted by these two models differ substantially across most of the parameter space. This comparison further suggests that `gwModel_kick_prec_flow` performs accurately in the comparable-mass regime and agrees closely with `NRSur7dq4Remnant` within this range. The advantage of `gwModel_kick_prec_flow` is that it extends to extreme mass ratios which `NRSur7dq4Remnant` does not. We also compare the distributions obtained from the `NRSur7dq4Remnant` and `HLZ` models and find that they differ in the regime where one of the spin magnitudes is small.

We repeat the same experiment for large mass ratios ($q \in [4, 100]$), computing the Jensen–Shannon divergence between the kick-velocity distributions predicted by `gwModel_kick_prec_flow` and `HLZ` (Figure 8). The distributions are obtained assuming isotropic spin orientations and spin magnitudes $|\chi_1|$ and $|\chi_2|$ uniformly sampled within $[0, 1]$. We find that, across most of the parameter space, the kick-velocity distributions from these two models differ substantially, yielding JSD values significantly greater than 0.1.

Table I. Median evaluation cost and standard deviation (in seconds) for computing kick-velocity distributions using a single CPU for 50 randomly sampled mass ratios within $q \in [1, 4]$ and spin magnitudes $|\chi_{1,2}| \in [-1, 1]$, with (2500) spin-orientation angles drawn from an isotropic distribution for three different models.

Model	Median [s]	Std. Dev. [s]
<code>gwModel_kick_prec_flow</code>	0.0488	0.0016
<code>HLZ</code>	0.00024	0.00001
<code>NRSur7dq4Remnant</code>	2.7281	0.1574

C. Computational speedup

We now evaluate the computational efficiency of our precessing-spin model `gwModel_kick_prec_flow` and compare it with two existing models: `NRSur7dq4Remnant` and `HLZ`. While `gwModel_kick_prec_flow` is a normalizing-flow-based model, `NRSur7dq4Remnant` relies on GPR, and `HLZ` is fully analytic.

On a single CPU, we randomly sample 50 mass ratios within $q \in [1, 4]$ and spin magnitudes $|\chi_{1,2}| \in [-1, 1]$. For each configuration, we compute the kick-velocity distribution resulting from 2,500 isotropically distributed spin orientations. Table I shows the distribution of evaluation times for all three models. On average, `gwModel_kick_prec_flow` requires ~ 0.05 seconds per evaluation, whereas `HLZ` takes only ~ 0.0025 seconds, and `NRSur7dq4Remnant` takes ~ 3 seconds. Thus, `gwModel_kick_prec_flow` (`HLZ`) is approximately 60 (1200) times faster than `NRSur7dq4Remnant`.

We further find that increasing the number of spin-angle samples causes the evaluation cost of `NRSur7dq4Remnant` to grow linearly, as each configuration is computed sequentially. In contrast, the evaluation times for `gwModel_kick_prec_flow` and `HLZ` remain nearly constant regardless of sample size. Note also that the evaluation cost of GPR model scales as $O(N_{\text{train}})$, where N_{train} is the number of training data samples, while the evaluation cost of `gwModel_kick_prec_flow` and `HLZ` is independent of N_{train} .

Our analytic aligned-spin model `gwModel_kick_q200` achieves comparable efficiency to `HLZ`, while the GPR-based `gwModel_kick_q200_GPR` incurs similar computational costs to `NRSur7dq4Remnant`. For both computational and extrapolation-related reasons, we do not recommend using `gwModel_kick_q200_GPR` for astrophysical applications.

D. Extrapolation capabilities

One of the key challenges with data-driven models is their tendency to perform poorly outside their training domain (often exhibiting divergent behavior) or in regions of parameter space with sparse data coverage. This issue becomes more pronounced in high-dimensional spaces such as for precessing binaries where the system is typically described by seven intrinsic parameters. Although our normalizing-flow model marginalizes over the spin-orientation angles, the underlying data still occupy this seven-dimensional space. In




Figure 9. Unlike the GPR based models, we find that the `gwModel_kick_prec_flow` model shows reasonable extrapolation outside its training parameter region ($q \leq 100$). Its kick velocity distribution (due to different spin orientation angles) predictions vary smoothly across different mass ratios in the extreme-mass-ratio regime, extending well beyond its training range. For this test, the spin magnitudes are fixed at $|\chi_1| = 0.8$ and $|\chi_2| = 0.9$. More details are in Section III D.

previous sections, we demonstrated that GPR-based models such as `gwModel_kick_q200`, `NRSur7dq4Remnant`, and `NRSur3dq8Remnant` often fail when extrapolated or evaluated in data-scarce regions of the parameter space. As shown in Figure 1 and Figure 3, these GPR models tend to deviate rapidly from the data beyond their training range, producing either unphysical blow-ups or non-smooth behavior.

To assess whether similar issues arise for our normalizing-flow model `gwModel_kick_prec_flow`, we perform extensive tests across a wide range of parameter-space points, including extreme extrapolations up to mass ratios as large as $q = 10^4$, even though the model was trained only up to $q = 100$ (Figure 9). We find that our model's predictions remain smooth and non-divergent, decreasing monotonically with increasing mass ratio as expected. One could further improve the flow model results to mimic the asymptotic η^2 scaling expected from analytic studies, however the current extrapolation kick values are sufficient for most astrophysical simulations as they are lower than the typical escape velocities in astrophysical systems. We emphasize that, to our knowledge, this is the first data-driven model in GW astrophysics that exhibits such stable and physically sensible extrapolation behavior far beyond its training domain.

IV. ASTROPHYSICAL IMPLICATIONS

We now explore one of the astrophysical implications of our models in the context of hierarchical mergers in dense astrophysical environments such as globular clusters, nuclear




Figure 10. Retention probability of first-generation binary black hole (BBH) mergers as a function of the host environment's escape velocity. The upper panel shows results for precessing BBHs with kick velocities computed using the `gwModel_kick_prec_flow` (blue), `NRSur7dq4Remnant` (red), and `HLZ` (green) models; the lower panel shows aligned-spin BBHs with kick velocities computed using the `gwModel_kick_q200` (blue), `NRSur3dq8Remnant` (orange), and `HLZ` (green) models. Typical escape-velocity ranges for globular clusters, nuclear star clusters, elliptical galaxies, and active galactic nuclei are indicated for reference. All results are obtained using the `gwGenealogy` package. More details are in Section IV A.

star clusters, and elliptical galaxies. We perform both back-of-the-envelope calculations and detailed semi-analytic N -body cluster simulations.

A. Back-of-the-envelope calculations for clusters

We first focus on the use of the precessing-spin kick model in the context of binaries in environments such as globular clusters, nuclear star clusters, and elliptical galaxies where binaries are mostly expected to be precessing. For this, we use




Figure 11. (*Upper panel*): Distribution of BH retention probabilities obtained from detailed stellar evolution and BBH assembly simulations for 180 different cluster models with varying initial number of stars, half-mass radius, metallicity, and natal spin of $1g$ BHs. We compare results using two kick-velocity models: `gwModel_kick_prec_flow` (blue) and `HLZ` (green). (*Lower panel*): Distributions of the differences in overall BH retention probability as a function of the initial cluster mass M_{cl} , computed using the same two kick-velocity models. Each marker is color-coded by the cluster half-mass radius r_h , and the symbol size scales with the natal black hole spin s_{BH} —smaller circles correspond to lower spins, while larger circles correspond to higher spins. Results are obtained using the `rapster` package. More details are in Section IV B.

`gwGenealogy` package². We generate an ensemble of 5000 BHs whose component masses $m_{1,2}$ follow a power-law distribution, $p(m_{1,2}) \propto m^{-\beta_m}$ (with $\beta_m = -2$) [87], within the range $[5, 50] M_\odot$ —representing a fiducial first-generation black hole population formed through stellar collapse. We assign spin magnitudes such that 50% of the BHs follow a Beta distribution, $p(|\chi_{1,2}|) \propto x^{\alpha-1}(1-x)^{\beta-1}$ (with $\alpha = 1.4$ and $\beta = 3.6$) [88], while the remaining 50% have spins drawn uniformly from $\mathcal{U}(0, 1)$. The spin orientation angles are chosen isotropically: the polar angles $\theta_{1,2}$ are sampled uniformly in $\cos(\theta_{1,2})$, and

the azimuthal angles $\phi_{1,2}$ are drawn uniformly from $[0, 2\pi)$. We then pair and merge these BHs randomly without imposing any ordering constraints. For each merger, we compute the recoil (kick) velocity using three precessing-spin models—`gwModel_kick_prec_flow`, `NRSur7dq4Remnant`, and `HLZ`. Although some mergers have mass ratios exceeding $q = 6$, we include them in the `NRSur7dq4Remnant` sample for completeness. Once the kick velocities are computed, we estimate the remnant retention probability by assuming that remnants with kick velocities smaller than the host environment’s escape velocity are retained.

In Figure 10, we show the retention probabilities as a function of the escape velocities. For reference, we also indicate the typical escape-velocity ranges of various astrophysical host environments, including globular clusters (up to $\sim 60 \text{ km s}^{-1}$) [89], nuclear star clusters (up to $\sim 600 \text{ km s}^{-1}$) [89], and elliptical galaxies (ranging from ~ 400 to $\sim 2400 \text{ km s}^{-1}$ [65]). We also mark the approximate escape velocity of the Milky Way ($\sim 500 \text{ km s}^{-1}$ [90]). We find that the retention probability of merger remnants increases when using the `gwModel_kick_prec_flow` model, while the `HLZ` model yields the lowest retention probabilities. This trend is consistent with our findings in Fig. 6 and Section III, where `HLZ` consistently predicted larger kick velocities. Similar results have also been reported in Ref. [49]. Overall, we find that using `gwModel_kick_prec_flow` can increase the remnant retention probability by approximately 5–10% for many binaries, which could significantly influence hierarchical merger rates in low-mass globular clusters.

Next, we focus on aligned-spin BBHs in AGNs. Several formation scenarios suggest that, due to interactions with the surrounding gas disk, the component spins in AGNs are more likely to align with the orbital angular momentum axis, leading predominantly to aligned-spin BBHs [91]. We adopt the same mass distribution as in the previous experiment. The spin components along the z -axis are now sampled uniformly between -1 and 1 . We then generate an ensemble of 5000 BBHs and merge them randomly. For each merger, we compute the remnant kick velocity using three aligned-spin models: `gwModel_kick_q200`, `NRSur3dq8Remnant`, and `HLZ`. For reference, we also indicate the typical escape velocities of various host environments, including globular clusters, nuclear star clusters, and particularly active galactic nuclei (between 200 km/s to 400 km/s [92]). We find that the `HLZ` model consistently predicts larger kicks, resulting in lower remnant retention probabilities. When comparing `gwModel_kick_q200` with `NRSur3dq8Remnant`, we find very similar retention probabilities, with the latter yielding slightly higher values. However, as demonstrated in Section II, `gwModel_kick_q200` provides more accurate predictions than `NRSur3dq8Remnant` in the aligned-spin regime.

While these back-of-the-envelope calculations provide some insight, they neglect many astrophysical effects that are present in star clusters and AGNs. It is therefore important to perform detailed cluster-evolution modeling that incorporates these effects. We do so in the next section.

² <https://github.com/tousifislam/gwGenealogy>

B. Semi-analytic cluster simulations

Next, we use the rapid cluster evolution code `rapster`³ [93] to simulate the evolution of massive star clusters and the assembly of BBHs within them, using more realistic prescriptions than those employed in Section IV A. Unlike detailed N -body simulation codes, which are computationally expensive, `rapster` adopts efficient semi-analytic prescriptions calibrated against full N -body results, allowing rapid yet physically motivated modeling of cluster evolution.

We simulate a total of 1404 clusters, varying the half-mass radius $r_h = [0.1, 0.2, 0.5, 0.7, 1, 2, 5, 7, 10]$ pc, metallicity $Z = [0.0002, 0.001, 0.02]$, and natal spin of first-generation (1g) BHs $s = [0.0, 0.2, 0.4, 0.6, 0.8]$. Furthermore, we choose 13 different values for the initial number of stars, N_{star} , within the range $[5, 100] \times 10^4$. This corresponds to an initial cluster mass M_{cl} in the range $\sim 3 \times 10^4 M_\odot$ to $\sim 6 \times 10^6 M_\odot$. The redshift of cluster formation is set to $z = 3$. The initial stellar population follows a zero-age main sequence (ZAMS) mass distribution $M_{\text{ZAMS}} \in [0.08, 150] M_\odot$, and the initial binary fraction is set to 0.1. Furthermore, for the 1g BHs, `rapster` assigns the spin orientation angles isotropically. By default, `rapster` computes remnant kick velocities using the `HLZ` prescription. We have now modified the code to incorporate our `gwModel_kick_prec_flow` model for kick-velocity calculations. In `rapster`, the escape velocity of each cluster is determined self-consistently using the virial theorem.

Out of these 1404 clusters, only 628 produced BBH mergers. In Figure 11 (upper panel), we show the distribution of BH retention probabilities across these 628 simulated cluster models when using two different kick-velocity prescriptions. We find that for many clusters, the retention probability can vary significantly depending on the chosen model. The median BH retention probabilities are 54.5% for `HLZ` and 52.7% for `gwModel_kick_prec_flow`, indicating only a modest overall difference. However, a larger fraction of clusters exhibit retention probabilities exceeding 80% when `gwModel_kick_prec_flow` is used.

Overall, we find that the retention probabilities can differ by as much as 70% across clusters (Figure 11, lower panel). To investigate whether any specific cluster property drives this difference, we plot the variation in retention probabilities as a function of the initial cluster mass and half-mass radius, where the size of each marker (Figure 11, lower panel) indicates the natal black hole spin. We find that in low-mass clusters, fewer black holes are produced, and for these systems, differences between the `HLZ` and `gwModel_kick_prec_flow` kick-velocity distributions almost always affect the retention rates. In contrast, for high-mass clusters ($M_{\text{cl}} \geq 10^5 M_\odot$), the differences vary considerably depending on other cluster parameters.

V. CONCLUDING REMARKS

In this paper, we presented two kick-velocity models for aligned-spin binaries: `gwModel_kick_q200` and `gwModel_kick_q200_GPR`. The former is an analytic model, while the latter is based on GPR. Both models are trained on a combination of NR and BHPT data up to a mass ratio of $q = 200$. We find that both models reproduce the training data accurately, as verified through 5-fold cross-validation, and outperform existing state-of-the-art NR surrogate models `NRSur7dq4Remnant` and `NRSur3dq8Remnant` and analytic `HLZ` model in the aligned-spin limit. Furthermore, our model errors are comparable to the intrinsic NR uncertainties up to $q = 128$. Finally, we demonstrate that GPR-based models—including `gwModel_kick_q200_GPR`, `NRSur7dq4Remnant`, and `NRSur3dq8Remnant`—while often yielding low validation errors, can exhibit non-smooth behavior across the parameter space and produce unreliable predictions when extrapolated or in data-sparse regions.

Next, we develop a normalizing-flow-based probabilistic model, `gwModel_kick_prec_flow`, which models the distribution of kick velocities marginalized over all spin orientation angles and conditioned on the binary’s mass ratio and spin magnitudes. Through extensive validation and comparison against the `NRSur7dq4Remnant` and `HLZ` models, we demonstrate that `gwModel_kick_prec_flow` closely matches `NRSur7dq4Remnant` across most of the comparable-mass regime ($q \leq 4$), where the latter is valid. However, `gwModel_kick_prec_flow` predicts noticeably different kick-velocity distributions compared to the `HLZ` model.

Together, for the first time, we provide kick-velocity models for both aligned-spin and precessing binary black holes that remain accurate across the full mass-ratio regime and outperform all existing models. These models enable astrophysical population synthesis and hierarchical merger studies over a wide range of mass ratios without the need to combine models with different validity domains, thereby minimizing potential systematic biases. One current limitation is that our precessing-spin model is probabilistic and marginalizes over spin-orientation angles, as our primary goal here is to support population synthesis and hierarchical merger applications. In future work, we plan to develop an extended version capable of providing point-wise kick predictions for specific precessing-spin configurations. For wider community use, all models presented in this work are publicly available through the `gwModels` package.

We have also performed simplified calculations considering ~ 5000 BBHs in clusters and AGNs, computing their retention probabilities as a function of the host escape velocity using our models `gwModel_kick_prec_flow` (for precessing-spin BBHs) and `gwModel_kick_q200` (for aligned-spin BBHs). We find that when `gwModel_kick_prec_flow` is used, the retention probability increases by up to 10% across several escape-velocity ranges. For aligned-spin systems, the retention probability is slightly lower compared to `NRSur3dq8Remnant`. In both cases, the `HLZ` model yields significantly smaller retention probabilities. We further performed detailed semi-analytic stellar evolution and BBH assembly simulations using the `rapster` code, modeling 1404 different clusters

³ <https://github.com/Kkritos/Rapster>

with varying initial number of stars, metallicity, BH natal spin, and half-mass radius. We again find that the retention probability can change substantially when using `gwModel_kick_prec_flow` instead of the default `HLZ` prescription. Using `gwModel_kick_prec_flow` model in cluster simulations can therefore change hierarchical merger scenarios.

In the future, we plan to extend this work in several directions. First, we aim to incorporate our kick models into different standard star-cluster and population-synthesis frameworks (apart from `rapster`) to study hierarchical mergers in a more systematic and astrophysically consistent manner. Second, we plan to include orbital eccentricity in our kick-velocity models. Although previous studies have suggested that residual eccentricity can enhance recoil velocities based on post-Newtonian and perturbative approximations [15, 33, 66], the limited availability of spinning, eccentric NR simulations currently prevents the construction of accurate data-driven models in this regime. We further aim to distill our flow model into a fully analytic model, which will reduce evaluation costs and facilitate both interpretation and implementability.

ACKNOWLEDGMENTS

We thank Gaurav Khanna for generously providing us the BHPT simulations and Konstantinos Kritos for help with pop-

ulation synthesis codes. We are grateful to the SXS collaboration and RIT NR group for maintaining publicly available catalog of NR simulations which has been used in this study. We thank Shadab Alam and Mahajan Rutvik Ashish for pointing out issues with existing kick models and for motivating this work. We thank Maya Fisbach, Carl Rodriguez, Scott Field, and Vijay Varma for helpful discussions. T.I. is supported in part by the National Science Foundation under Grant No. NSF PHY-2309135 and the Gordon and Betty Moore Foundation Grant No. GBMF7392. Use was made of computational facilities purchased with funds from the National Science Foundation (CNS-1725797) and administered by the Center for Scientific Computing (CSC). The CSC is supported by the California NanoSystems Institute and the Materials Research Science and Engineering Center (MRSEC; NSF DMR 2308708) at UC Santa Barbara. D.W. is supported by NSF Grants No. AST-2307146, No. PHY-2513337, No. PHY-090003, and No. PHY-20043, by NASA Grant No. 21-ATP21-0010, by John Templeton Foundation Grant No. 62840, by the Simons Foundation [MPS-SIP-00001698, E.B.], by the Simons Foundation International [SFI-MPS-BH-00012593-02], and by Italian Ministry of Foreign Affairs and International Cooperation Grant No. PGR01167. This work was carried out at the Advanced Research Computing at Hopkins (ARCH) core facility (<https://www.arch.jhu.edu/>), which is supported by the NSF Grant No. OAC-1920103.

-
- [1] John G. Baker, Joan Centrella, Dae-Il Choi, Michael Koppitz, James R. van Meter, and M. Coleman Miller, “Getting a kick out of numerical relativity,” *Astrophys. J. Lett.* **653**, L93–L96 (2006), [arXiv:astro-ph/0603204](#).
 - [2] John G. Baker, William D. Boggs, Joan Centrella, Bernard J. Kelly, Sean T. McWilliams, M. Coleman Miller, and James R. van Meter, “Modeling kicks from the merger of non-precessing black-hole binaries,” *Astrophys. J.* **668**, 1140–1144 (2007), [arXiv:astro-ph/0702390](#).
 - [3] John G. Baker, William D. Boggs, Joan Centrella, Bernard J. Kelly, Sean T. McWilliams, M. Coleman Miller, and James R. van Meter, “Modeling kicks from the merger of generic black-hole binaries,” *Astrophys. J. Lett.* **682**, L29–L32 (2008), [arXiv:0802.0416 \[astro-ph\]](#).
 - [4] F. Herrmann, I. Hinder, D. Shoemaker, and P. Laguna, “Binary black holes and recoil velocities,” *AIP Conf. Proc.* **873**, 89–93 (2006).
 - [5] Carlos O. Lousto and Yosef Zlochower, “Further insight into gravitational recoil,” *Phys. Rev. D* **77**, 044028 (2008), [arXiv:0708.4048 \[gr-qc\]](#).
 - [6] Frank Herrmann, Ian Hinder, Deirdre Shoemaker, Pablo Laguna, and Richard A. Matzner, “Gravitational recoil from spinning binary black hole mergers,” *Astrophys. J.* **661**, 430–436 (2007), [arXiv:gr-qc/0701143](#).
 - [7] Frank Herrmann, Ian Hinder, Deirdre M. Shoemaker, Pablo Laguna, and Richard A. Matzner, “Binary Black Holes: Spin Dynamics and Gravitational Recoil,” *Phys. Rev. D* **76**, 084032 (2007), [arXiv:0706.2541 \[gr-qc\]](#).
 - [8] Frank Herrmann, Ian Hinder, Deirdre Shoemaker, and Pablo Laguna, “Unequal mass binary black hole plunges and gravitational recoil,” *Class. Quant. Grav.* **24**, S33–S42 (2007).
 - [9] Kelly Holley-Bockelmann, Kayhan Gültekin, Deirdre Shoemaker, and Nico Yunes, “Gravitational Wave Recoil and the Retention of Intermediate Mass Black Holes,” *Astrophys. J.* **686**, 829 (2008), [arXiv:0707.1334 \[astro-ph\]](#).
 - [10] Jose Luis Jaramillo, Rodrigo Panosso Macedo, Philipp Moesta, and Luciano Rezzolla, “Black-hole horizons as probes of black-hole dynamics I: post-merger recoil in head-on collisions,” *Phys. Rev. D* **85**, 084030 (2012), [arXiv:1108.0060 \[gr-qc\]](#).
 - [11] Michael Koppitz, Denis Pollney, Christian Reisswig, Luciano Rezzolla, Jonathan Thornburg, Peter Diener, and Erik Schnetter, “Recoil Velocities from Equal-Mass Binary-Black-Hole Mergers,” *Phys. Rev. Lett.* **99**, 041102 (2007), [arXiv:gr-qc/0701163](#).
 - [12] Carlos O. Lousto and Yosef Zlochower, “Modeling gravitational recoil from precessing highly-spinning unequal-mass black-hole binaries,” *Phys. Rev. D* **79**, 064018 (2009), [arXiv:0805.0159 \[gr-qc\]](#).
 - [13] Carlos O. Lousto and Yosef Zlochower, “Modeling maximum astrophysical gravitational recoil velocities,” *Phys. Rev. D* **83**, 024003 (2011), [arXiv:1011.0593 \[gr-qc\]](#).
 - [14] Jeremy D. Schnittman, Alessandra Buonanno, James R. van Meter, John G. Baker, William D. Boggs, Joan Centrella, Bernard J. Kelly, and Sean T. McWilliams, “Anatomy of the binary black hole recoil: A multipolar analysis,” *Phys. Rev. D* **77**, 044031 (2008), [arXiv:0707.0301 \[gr-qc\]](#).
 - [15] Carlos F. Sopuerta, Nicolas Yunes, and Pablo Laguna, “Gravitational recoil velocities from eccentric binary black hole mergers,” *Astrophys. J. Lett.* **656**, L9–L12 (2007), [arXiv:astro-ph/0611110](#).
 - [16] Denis Pollney *et al.*, “Recoil velocities from equal-mass binary black-hole mergers: A Systematic investigation of spin-

- orbit aligned configurations,” *Phys. Rev. D* **76**, 124002 (2007), arXiv:0707.2559 [gr-qc].
- [17] Luciano Rezzolla, Rodrigo P. Macedo, and Jose Luis Jaramillo, “Understanding the ‘anti-kick’ in the merger of binary black holes,” *Phys. Rev. Lett.* **104**, 221101 (2010), arXiv:1003.0873 [gr-qc].
- [18] Carlos O. Lousto and Yosef Zlochower, “Hangup Kicks: Still Larger Recoils by Partial Spin/Orbit Alignment of Black-Hole Binaries,” *Phys. Rev. Lett.* **107**, 231102 (2011), arXiv:1108.2009 [gr-qc].
- [19] Carlos O. Lousto and Yosef Zlochower, “Nonlinear Gravitational Recoil from the Mergers of Precessing Black-Hole Binaries,” *Phys. Rev. D* **87**, 084027 (2013), arXiv:1211.7099 [gr-qc].
- [20] Carlos O. Lousto, Yosef Zlochower, Massimo Dotti, and Marta Volonteri, “Gravitational Recoil From Accretion-Aligned Black-Hole Binaries,” *Phys. Rev. D* **85**, 084015 (2012), arXiv:1201.1923 [gr-qc].
- [21] Sarah H. Miller and R. A. Matzner, “Multipole Analysis of Kicks in Collision of Binary Black Holes,” *Gen. Rel. Grav.* **41**, 525–539 (2009), arXiv:0807.3028 [gr-qc].
- [22] Wolfgang Tichy and Pedro Marronetti, “Binary black hole mergers: Large kicks for generic spin orientations,” *Phys. Rev. D* **76**, 061502 (2007), arXiv:gr-qc/0703075.
- [23] Yosef Zlochower, Manuela Campanelli, and Carlos O. Lousto, “Modeling Gravitational Recoil Using Numerical Relativity,” *Class. Quant. Grav.* **28**, 114015 (2011), arXiv:1011.2210 [gr-qc].
- [24] James Healy, Carlos O. Lousto, and Yosef Zlochower, “Remnant mass, spin, and recoil from spin aligned black-hole binaries,” *Phys. Rev. D* **90**, 104004 (2014), arXiv:1406.7295 [gr-qc].
- [25] Carlos O. Lousto, Manuela Campanelli, Yosef Zlochower, and Hiroyuki Nakano, “Remnant Masses, Spins and Recoils from the Merger of Generic Black-Hole Binaries,” *Class. Quant. Grav.* **27**, 114006 (2010), arXiv:0904.3541 [gr-qc].
- [26] Hiroyuki Nakano, Manuela Campanelli, Carlos O. Lousto, and Yosef Zlochower, “Perturbative effects of spinning black holes with applications to recoil velocities,” *Class. Quant. Grav.* **28**, 134005 (2011), arXiv:1011.2767 [gr-qc].
- [27] Pranesh A. Sundararajan, Gaurav Khanna, and Scott A. Hughes, “Binary black hole merger gravitational waves and recoil in the large mass ratio limit,” *Phys. Rev. D* **81**, 104009 (2010), arXiv:1003.0485 [gr-qc].
- [28] Tousif Islam, Scott E. Field, and Gaurav Khanna, “Remnant black hole properties from numerical-relativity-informed perturbation theory and implications for waveform modeling,” *Phys. Rev. D* **108**, 064048 (2023), arXiv:2301.07215 [gr-qc].
- [29] Scott A. Hughes, Marc Favata, and Daniel E. Holz, “How black holes get their kicks: Radiation recoil in binary black hole mergers,” in *Conference on Growing Black Holes: Accretion in a Cosmological Context* (2004) arXiv:astro-ph/0408492.
- [30] Richard H. Price, Gaurav Khanna, and Scott A. Hughes, “Black hole binary inspiral and trajectory dominance,” *Phys. Rev. D* **88**, 104004 (2013), arXiv:1306.1159 [gr-qc].
- [31] Richard H. Price, Gaurav Khanna, and Scott A. Hughes, “Systematics of black hole binary inspiral kicks and the slowness approximation,” *Phys. Rev. D* **83**, 124002 (2011), arXiv:1104.0387 [gr-qc].
- [32] Luc Blanchet, Moh’d S. S. Qusailah, and Clifford M. Will, “Gravitational recoil of inspiralling black-hole binaries to second post-Newtonian order,” *Astrophys. J.* **635**, 508 (2005), arXiv:astro-ph/0507692.
- [33] Carlos F. Sopuerta, Nicolas Yunes, and Pablo Laguna, “Gravitational Recoil from Binary Black Hole Mergers: The Close-Limit Approximation,” *Phys. Rev. D* **74**, 124010 (2006), [Erratum: *Phys. Rev. D* **75**, 069903 (2007), Erratum: *Phys. Rev. D* **78**, 049901 (2008)], arXiv:astro-ph/0608600.
- [34] Marc Favata, Scott A. Hughes, and Daniel E. Holz, “How black holes get their kicks: Gravitational radiation recoil revisited,” *Astrophys. J. Lett.* **607**, L5–L8 (2004), arXiv:astro-ph/0402056.
- [35] M. J. Fitchett, “The influence of gravitational wave momentum losses on the centre of mass motion of a Newtonian binary system,” *Mon. Not. Roy. Astron. Soc.* **203**, 1049–1062 (1983).
- [36] M. J. Fitchett and Steven L. Detweiler, “Linear momentum and gravitational-waves - circular orbits around a schwarzschild black-hole,” *Mon. Not. Roy. Astron. Soc.* **211**, 933–942 (1984).
- [37] Alan G. Wiseman, “Coalescing binary systems of compact objects to (post)5/2 Newtonian order. 2. Higher order wave forms and radiation recoil,” *Phys. Rev. D* **46**, 1517–1539 (1992).
- [38] Lawrence E. Kidder, “Coalescing binary systems of compact objects to postNewtonian 5/2 order. 5. Spin effects,” *Phys. Rev. D* **52**, 821–847 (1995), arXiv:gr-qc/9506022.
- [39] Manuela Campanelli, Carlos O. Lousto, Yosef Zlochower, and David Merritt, “Maximum gravitational recoil,” *Phys. Rev. Lett.* **98**, 231102 (2007), arXiv:gr-qc/0702133.
- [40] Bernd Bruegmann, Jose A. Gonzalez, Mark Hannam, Sascha Husa, and Ulrich Sperhake, “Exploring black hole superkicks,” *Phys. Rev. D* **77**, 124047 (2008), arXiv:0707.0135 [gr-qc].
- [41] Manuela Campanelli, Carlos O. Lousto, Yosef Zlochower, and David Merritt, “Large merger recoils and spin flips from generic black-hole binaries,” *Astrophys. J. Lett.* **659**, L5–L8 (2007), arXiv:gr-qc/0701164.
- [42] Dae-Il Choi, Bernard J. Kelly, William D. Boggs, John G. Baker, Joan Centrella, and James van Meter, “Recoiling from a kick in the head-on collision of spinning black holes,” *Phys. Rev. D* **76**, 104026 (2007), arXiv:gr-qc/0702016.
- [43] Sergio Dain, Carlos O. Lousto, and Yosef Zlochower, “Extra-Large Remnant Recoil Velocities and Spins from Near-Extremal-Bowen-York-Spin Black-Hole Binaries,” *Phys. Rev. D* **78**, 024039 (2008), arXiv:0803.0351 [gr-qc].
- [44] Jose A. Gonzalez, Ulrich Sperhake, Bernd Bruegmann, Mark Hannam, and Sascha Husa, “Total recoil: The Maximum kick from nonspinning black-hole binary inspiral,” *Phys. Rev. Lett.* **98**, 091101 (2007), arXiv:gr-qc/0610154.
- [45] J. A. Gonzalez, M. D. Hannam, U. Sperhake, Bernd Bruegmann, and S. Husa, “Supermassive recoil velocities for binary black-hole mergers with antialigned spins,” *Phys. Rev. Lett.* **98**, 231101 (2007), arXiv:gr-qc/0702052.
- [46] James Healy, Frank Herrmann, Ian Hinder, Deirdre M. Shoemaker, Pablo Laguna, and Richard A. Matzner, “Superkicks in Hyperbolic Encounters of Binary Black Holes,” *Phys. Rev. Lett.* **102**, 041101 (2009), arXiv:0807.3292 [gr-qc].
- [47] James Healy and Carlos O. Lousto, “Ultimate Black Hole Recoil: What is the Maximum High-Energy Collision Kick?” *Phys. Rev. Lett.* **131**, 071401 (2023), arXiv:2301.00018 [gr-qc].
- [48] Davide Gerosa and Christopher J. Moore, “Black hole kicks as new gravitational wave observables,” *Phys. Rev. Lett.* **117**, 011101 (2016), arXiv:1606.04226 [gr-qc].
- [49] Angela Borchers, Claire S. Ye, and Maya Fishbach, “Gravitational-wave Kicks Impact the Spins of Black Holes from Hierarchical Mergers,” *Astrophys. J.* **987** (2025), 10.3847/1538-4357/addec6, arXiv:2503.21278 [astro-ph.HE].
- [50] Emanuele Berti, Michael Kesden, and Ulrich Sperhake, “Effects of post-Newtonian Spin Alignment on the Distribution of Black-Hole Recoils,” *Phys. Rev. D* **85**, 124049 (2012), arXiv:1203.2920 [astro-ph.HE].
- [51] Kayhan Gültekin, M. Coleman Miller, and Douglas P. Hamilton, “Growth of intermediate - mass black holes in globular clusters,” *Astrophys. J.* **616**, 221–230 (2004), arXiv:astro-ph/0402532.

- [52] Davide Gerosa, Nicola Giacobbo, and Alberto Vecchio, “High mass but low spin: an exclusion region to rule out hierarchical black-hole mergers as a mechanism to populate the pair-instability mass gap,” *Astrophys. J.* **915**, 56 (2021), [arXiv:2104.11247 \[astro-ph.HE\]](#).
- [53] Davide Gerosa and Emanuele Berti, “Are merging black holes born from stellar collapse or previous mergers?” *Phys. Rev. D* **95**, 124046 (2017), [arXiv:1703.06223 \[gr-qc\]](#).
- [54] Vishal Baibhav, Davide Gerosa, Emanuele Berti, Kaze W. K. Wong, Thomas Helfer, and Matthew Mould, “The mass gap, the spin gap, and the origin of merging binary black holes,” *Phys. Rev. D* **102**, 043002 (2020), [arXiv:2004.00650 \[astro-ph.HE\]](#).
- [55] Yann Bouffanais, Michela Mapelli, Davide Gerosa, Ugo N. Di Carlo, Nicola Giacobbo, Emanuele Berti, and Vishal Baibhav, “Constraining the fraction of binary black holes formed in isolation and young star clusters with gravitational-wave data,” (2019), [10.3847/1538-4357/ab4a79](#), [arXiv:1905.11054 \[astro-ph.HE\]](#).
- [56] Davide Gerosa, Salvatore Vitale, and Emanuele Berti, “Astrophysical implications of GW190412 as a remnant of a previous black-hole merger,” *Phys. Rev. Lett.* **125**, 101103 (2020), [arXiv:2005.04243 \[astro-ph.HE\]](#).
- [57] Michael Boylan-Kolchin, Chung-Pei Ma, and Eliot Quataert, “Core formation in galactic nuclei due to recoiling black holes,” *Astrophys. J. Lett.* **613**, L37 (2004), [arXiv:astro-ph/0407488](#).
- [58] Nader Khonji, Alessia Gualandris, Justin I. Read, and Walter Dehnen, “Core Formation by Binary Scouring and Gravitational Wave Recoil in Massive Elliptical Galaxies,” *Astrophys. J.* **974**, 204 (2024), [arXiv:2408.12537 \[astro-ph.GA\]](#).
- [59] Laura Blecha, Thomas J. Cox, Abraham Loeb, and Lars Hernquist, “Recoiling black holes in merging galaxies: relationship to active galactic nucleus lifetimes, starbursts and the $M_{BH}-\sigma_*$ relation,” *Mon. Not. Roy. Astron. Soc.* **412**, 2154–2182 (2011), [arXiv:1009.4940 \[astro-ph.CO\]](#).
- [60] Davide Gerosa and Emanuele Berti, “Escape speed of stellar clusters from multiple-generation black-hole mergers in the upper mass gap,” *Phys. Rev. D* **100**, 041301 (2019), [arXiv:1906.05295 \[astro-ph.HE\]](#).
- [61] Yosef Zlochower and Carlos O. Lousto, “Modeling the remnant mass, spin, and recoil from unequal-mass, precessing black-hole binaries: The Intermediate Mass Ratio Regime,” *Phys. Rev. D* **92**, 024022 (2015), [Erratum: *Phys. Rev. D* **94**, 029901 (2016)], [arXiv:1503.07536 \[gr-qc\]](#).
- [62] James R. van Meter, M. Coleman Miller, John G. Baker, William D. Boggs, and Bernard J. Kelly, “A General Formula for Black Hole Gravitational Wave Kicks,” *Astrophys. J.* **719**, 1427 (2010), [arXiv:1003.3865 \[astro-ph.HE\]](#).
- [63] Vijay Varma, Scott E. Field, Mark A. Scheel, Jonathan Blackman, Davide Gerosa, Leo C. Stein, Lawrence E. Kidder, and Harald P. Pfeiffer, “Surrogate models for precessing binary black hole simulations with unequal masses,” *Phys. Rev. Research* **1**, 033015 (2019), [arXiv:1905.09300 \[gr-qc\]](#).
- [64] Vijay Varma, Davide Gerosa, Leo C. Stein, François Hébert, and Hao Zhang, “High-accuracy mass, spin, and recoil predictions of generic black-hole merger remnants,” *Phys. Rev. Lett.* **122**, 011101 (2019), [arXiv:1809.09125 \[gr-qc\]](#).
- [65] David Merritt, Milos Milosavljevic, Marc Favata, Scott A. Hughes, and Daniel E. Holz, “Consequences of gravitational radiation recoil,” *Astrophys. J. Lett.* **607**, L9–L12 (2004), [arXiv:astro-ph/0402057](#).
- [66] U. Sperhake, R. Rosca-Mead, D. Gerosa, and E. Berti, “Amplification of superkicks in black-hole binaries through orbital eccentricity,” *Phys. Rev. D* **101**, 024044 (2020), [arXiv:1910.01598 \[gr-qc\]](#).
- [67] Abdul H. Mroue *et al.*, “Catalog of 174 Binary Black Hole Simulations for Gravitational Wave Astronomy,” *Phys. Rev. Lett.* **111**, 241104 (2013), [arXiv:1304.6077 \[gr-qc\]](#).
- [68] Michael Boyle *et al.*, “The SXS Collaboration catalog of binary black hole simulations,” *Class. Quant. Grav.* **36**, 195006 (2019), [arXiv:1904.04831 \[gr-qc\]](#).
- [69] James Healy and Carlos O. Lousto, “Third RIT binary black hole simulations catalog,” *Phys. Rev. D* **102**, 104018 (2020), [arXiv:2007.07910 \[gr-qc\]](#).
- [70] James Healy and Carlos O. Lousto, “Fourth RIT binary black hole simulations catalog: Extension to eccentric orbits,” *Phys. Rev. D* **105**, 124010 (2022), [arXiv:2202.00018 \[gr-qc\]](#).
- [71] Mark A. Scheel *et al.*, “The SXS collaboration’s third catalog of binary black hole simulations,” *Class. Quant. Grav.* **42**, 195017 (2025), [arXiv:2505.13378 \[gr-qc\]](#).
- [72] Gaurav Khanna, “Teukolsky evolution of particle orbits around Kerr black holes in the time domain: elliptic and inclined orbits,” *Phys. Rev. D* **69**, 024016 (2004), [arXiv:gr-qc/0309107](#).
- [73] Lior Burko and Gaurav Khanna, “Accurate time-domain gravitational waveforms for extreme-mass-ratio binaries,” *Europhys. Lett.* **78**, 60005 (2007), [arXiv:gr-qc/0609002](#).
- [74] Pranesh A. Sundararajan, Gaurav Khanna, Scott A. Hughes, and Steve Drasco, “Towards adiabatic waveforms for inspiral into Kerr black holes: II. Dynamical sources and generic orbits,” *Phys. Rev. D* **78**, 024022 (2008), [arXiv:0803.0317 \[gr-qc\]](#).
- [75] Amos Ori and Kip S. Thorne, “The Transition from inspiral to plunge for a compact body in a circular equatorial orbit around a massive, spinning black hole,” *Phys. Rev. D* **62**, 124022 (2000), [arXiv:gr-qc/0003032](#).
- [76] Scott A. Hughes, Anuj Apte, Gaurav Khanna, and Halston Lim, “Learning about black hole binaries from their ringdown spectra,” *Phys. Rev. Lett.* **123**, 161101 (2019), [arXiv:1901.05900 \[gr-qc\]](#).
- [77] Anuj Apte and Scott A. Hughes, “Exciting black hole modes via misaligned coalescences: I. Inspiral, transition, and plunge trajectories using a generalized Ori-Thorne procedure,” *Phys. Rev. D* **100**, 084031 (2019), [arXiv:1901.05901 \[gr-qc\]](#).
- [78] Tousif Islam, Scott E. Field, Scott A. Hughes, Gaurav Khanna, Vijay Varma, Matthew Giesler, Mark A. Scheel, Lawrence E. Kidder, and Harald P. Pfeiffer, “Surrogate model for gravitational wave signals from nonspinning, comparable-to large-mass-ratio black hole binaries built on black hole perturbation theory waveforms calibrated to numerical relativity,” *Phys. Rev. D* **106**, 104025 (2022), [arXiv:2204.01972 \[gr-qc\]](#).
- [79] Tousif Islam, Gaurav Khanna, and Scott E. Field, “Consistency of spin effects between numerical relativity and perturbation theory for inspiraling comparable-mass black hole binaries,” (2025), [arXiv:2510.00531 \[gr-qc\]](#).
- [80] Katie Rink, Ritesh Bachhar, Tousif Islam, Nur E. M. Rifat, Kevin Gonzalez-Quesada, Scott E. Field, Gaurav Khanna, Scott A. Hughes, and Vijay Varma, “Gravitational wave surrogate model for spinning, intermediate mass ratio binaries based on perturbation theory and numerical relativity,” *Phys. Rev. D* **110**, 124069 (2024), [arXiv:2407.18319 \[gr-qc\]](#).
- [81] Maarten van de Meent and Harald P. Pfeiffer, “Intermediate mass-ratio black hole binaries: Applicability of small mass-ratio perturbation theory,” *Phys. Rev. Lett.* **125**, 181101 (2020), [arXiv:2006.12036 \[gr-qc\]](#).
- [82] Barry Wardell, Adam Pound, Niels Warburton, Jeremy Miller, Leanne Durkan, and Alexandre Le Tiec, “Gravitational Waveforms for Compact Binaries from Second-Order Self-Force Theory,” *Phys. Rev. Lett.* **130**, 241402 (2023), [arXiv:2112.12265 \[gr-qc\]](#).
- [83] Tousif Islam, “Interplay between numerical relativity and black hole perturbation theory in the intermediate-mass-ratio regime,”

- Phys. Rev. D **108**, 044013 (2023), arXiv:2306.08771 [gr-qc].
- [84] Tousif Islam and Gaurav Khanna (SXS), “Approximate relation between black hole perturbation theory and numerical relativity,” *Phys. Rev. D* **108**, 124046 (2023), arXiv:2307.03155 [gr-qc].
- [85] Pauli Virtanen, Ralf Gommers, Travis E. Oliphant, Matt Haberland, Tyler Reddy, David Cournapeau, Evgeni Burovski, Pearu Peterson, Warren Weckesser, Jonathan Bright, Stéfan J. van der Walt, Matthew Brett, Joshua Wilson, K. Jarrod Millman, Nikolay Mayorov, Andrew R. J. Nelson, Eric Jones, Robert Kern, Callum Larson, C J Carey, İlhan Polat, Yu Feng, Eric W. Moore, Jake VanderPlas, Denis Laxalde, Josef Perktold, Robert Cimrman, Ian Henriksen, E. A. Quintero, Charles R. Harris, Anne M. Archibald, Antônio H. Ribeiro, Fabian Pedregosa, Paul van Mulbregt, and SciPy 1.0 Contributors, “Scipy 1.0: Fundamental algorithms for scientific computing in python,” *Nature Methods* **17**, 261–272 (2020), arXiv:1907.10121 [cs.MS].
- [86] Conor Durkan, Artur Bekasov, Iain Murray, and George Papamakarios, “Neural spline flows,” (2019), arXiv:1906.04032 [stat.ML].
- [87] R. Abbott *et al.* (LIGO Scientific, Virgo), “Population Properties of Compact Objects from the Second LIGO-Virgo Gravitational-Wave Transient Catalog,” *Astrophys. J. Lett.* **913**, L7 (2021), arXiv:2010.14533 [astro-ph.HE].
- [88] R. Abbott *et al.* (KAGRA, VIRGO, LIGO Scientific), “Population of Merging Compact Binaries Inferred Using Gravitational Waves through GWTC-3,” *Phys. Rev. X* **13**, 011048 (2023), arXiv:2111.03634 [astro-ph.HE].
- [89] Fabio Antonini and Frederic A. Rasio, “Merging black hole binaries in galactic nuclei: implications for advanced-LIGO detections,” *Astrophys. J.* **831**, 187 (2016), arXiv:1606.04889 [astro-ph.HE].
- [90] G. Monari, B. Famaey, I. Carrillo, T. Piffl, M. Steinmetz, R. F. G. Wyse, F. Anders, C. Chiappini, and K. Janßen, “The escape speed curve of the galaxy obtained from gaia dr2 implies a heavy milky way,” *Astronomy & Astrophysics* **616**, L9 (2018).
- [91] Yang Yang *et al.*, “Hierarchical Black Hole Mergers in Active Galactic Nuclei,” *Phys. Rev. Lett.* **123**, 181101 (2019), arXiv:1906.09281 [astro-ph.HE].
- [92] Christina M. Manzano-King, Gabriela Canalizo, and Laura V. Sales, “Agn-driven outflows in dwarf galaxies,” *The Astrophysical Journal* **884**, 54 (2019).
- [93] Konstantinos Kritos, Vladimir Stokov, Vishal Baibhav, and Emanuele Berti, “Dynamical formation of black hole binaries in dense star clusters: Rapid cluster evolution code,” *Phys. Rev. D* **110**, 043023 (2024), arXiv:2210.10055 [astro-ph.HE].
- [94] Eleanor Hamilton *et al.*, “Catalog of precessing black-hole-binary numerical-relativity simulations,” *Phys. Rev. D* **109**, 044032 (2024), arXiv:2303.05419 [gr-qc].

1. Simple analytic expression

As `gwModel_kick_prec_flow` is a data-driven model, we also provide a simple analytic expression, which at this stage is significantly less accurate than `gwModel_kick_prec_flow`. The current version of the model is trained on a set of 80 BAM NR [94] simulations and 62 SXS NR simulations for single-precessing binaries, i.e., configurations with $\phi_1 = \phi_2 = \theta_2 = 0$. Following Refs. [12, 13, 19, 20, 45], we decompose the total

kick velocity as




Figure 12. We show the true and predicted kick velocities for a set of BAM (circles) and SXS (triangles) NR simulations of single-precessing binaries. The predictions are obtained from the simple analytic `gwModel_kick_singleprec` model (blue) and HLZ (orange). The black dashed line denotes the line of equality.

where $V_{\text{kick},\parallel}$ denotes the component of the recoil due to spin precession. We find that $V_{\text{kick},\parallel}$ can be reasonably approximated as:

$$\begin{aligned}
 V_{\parallel} = & \eta^2 |\chi_1| (A_1 \sin \theta_1 - B_1 \sin(2\theta_1)) \\
 & + \eta^2 |\chi_1|^2 (A_2 \sin \theta_1 + B_2 \sin(2\theta_1)) \\
 & + \eta^3 |\chi_1| (A_3 \sin \theta_1 + B_3 \sin(2\theta_1)) \\
 & + \eta^3 |\chi_1|^2 (A_4 \sin \theta_1 + B_4 \sin(2\theta_1)) \\
 & + \eta^4 |\chi_1| (A_5 \sin \theta_1 + B_5 \sin(2\theta_1)) \\
 & + \eta^4 |\chi_1|^2 (A_6 \sin \theta_1 + B_6 \sin(2\theta_1)).
 \end{aligned} \tag{18}$$

The best-fit coefficients for Eq. (18), obtained using the `scipy.curve_fit` [85] routine, are $A_1 = (3.772 \pm 150.262) \times 10^4$, $B_1 = (1.219 \pm 1.354) \times 10^5$, $A_2 = (3.021 \pm 274.358) \times 10^4$, $B_2 = (2.853 \pm 234.152) \times 10^5$, $A_3 = (-1.594 \pm 18.922) \times 10^5$, $B_3 = (1.177 \pm 15.907) \times 10^6$, $A_4 = (1.154 \pm 33.723) \times 10^5$, $B_4 = (-2.894 \pm 26.947) \times 10^6$, $A_5 = (3.835 \pm 50.846) \times 10^5$, $B_5 = (-3.012 \pm 42.894) \times 10^6$, $A_6 = (-1.110 \pm 91.642) \times 10^6$, and $B_6 = (7.595 \pm 71.445) \times 10^6$. We denote this model as `gwModel_kick_singleprec`.

In Figure 12, we show the predictions from `gwModel_kick_singleprec` alongside the corresponding BAM and SXS NR data. As mentioned earlier, this model successfully captures the overall trend but is not as accurate as the data-driven model developed in this work. Furthermore, we find similar level of disagreement between HLZ model and NR data. It is also important to note that, for the NR simulations, we use the initial binary configurations as inputs to our model. In principle, a more consistent treatment of the spin angle θ_1 —defined in an appropriate reference frame—should be employed. We leave this refinement for future work.