Using Black Hole Environments as Laboratories for Testing Accretion and Gravity

by

Boris Georgiev

A thesis
presented to the University of Waterloo
in fulfillment of the
thesis requirement for the degree of
Doctor of Philosophy
in
Physics

Waterloo, Ontario, Canada, 2023

Examining Committee Membership

The following served on the Examining Committee for this thesis. The decision of the Examining Committee is by majority vote.

External Examiner:	Alexander Tchekhovskoy Assistant Professor, Department of Physics and Astronomy, Northwestern University
Supervisors:	Avery Broderick Associate Professor, Department of Physics and Astronomy, University of Waterloo
	Luis Lehner Research Faculty, Perimeter Institute for Theoretical Physics
Internal Members:	James Taylor Associate Professor, Department of Physics and Astronomy, University of Waterloo
	Brian McNamara Professor, Department of Physics and Astronomy, University of Waterloo
Internal-External Member:	Michael Waite Professor, Department of Applied Mathematics, University of Waterloo

Author's Declaration

This thesis consists of material all of which I authored or co-authored: see Statement of Contributions included in the thesis. This is a true copy of the thesis, including any required final revisions, as accepted by my examiners.

I understand that my thesis may be made electronically available to the public.

Statement of Contributions

Work in this thesis also appears in published works, so some of it is written and done by members of the Event Horizon Telescope Collaboration. In particular:

- Paul Tiede performed the imaging simulations in Section 4.5.1, **Hotspot**;
- George Wong performed the imaging simulations in Section 4.5.1, **Elevation**;
- The code used in Subsection 4.3.4 was written by Avery Broderick.
- The text in Subsection 3.2.1 is a modified version of text originally written by Maciek Wielgus;
- The text in Section 1.3 is a modified version of text originally written by George Wong;
- The text in Chapter 3 and Section B.3 contains modified text originally written by Avery Broderick;
- The work summarized in Section B.2 is from publications by the Event Horizon Telescope Collaboration, and was done by a combination of myself, Avery Broderick, Dominic Pesce, and Michi Bauböck.

In addition, none of the GRMHD simulations and images were created by me. Credit goes to many members within the Event Horizon Telescope Collaboration.

Abstract

With the advent of the Event Horizon Telescope (EHT), we have the ability to observe the lensed emission from hot plasmas near event horizons. It contains entangled information about the turbulent magnetohydrodynamic accretion processes and the black hole spacetime. We present studies of both coherent and stochastic variable features measurable from EHT data, and by applying them to libraries of simulations, we describe how they inform on the underlying accretion flow. For a rapidly varying source like Sgr A*, we find that variability is greatest on the largest spatial scales and the longest timescale, and find a universal power-law variability prescription in the visibility domain. For a slowly varying source like M87*, we measure a correlation timescale and a rotation rate, which is inconsistent with the velocity of the fluid causing the emission. The strong predictions from the simulations of the variability measures we create are consistent with current EHT data. We also create a static analytical midplane accretion flow model and use variable simulations to motivate a prescription to incorporate turbulent effects. Our model matches SANE models well and struggles with the more variable MAD models. It can incorporate non-Kerr metrics, can connect the black hole scales to their large-scale environment, and can be incorporated in existing parameter estimation frameworks to fit EHT data.

Acknowledgements

I would like to thank my advisor, Avery Broderick, for outstanding support and mentorship during my graduate studies. Their tireless dedication to science and affinity for leading by example has inspired me to leap forward and explore the most novel ideas. I would also like to thank my advisor, Luis Lehner, for being a grounding presence that has encouraged me to rigorously investigate every aspect of my work.

I would like to thank Paul Tiede, a fellow graduate student, who would always take the time to explain scientific concepts and show me how to work with a variety of software. I would also to thank Dominic Pesce for being patient and helping me build confidence in my writing. To Paul, Dom, and the other students with whom I shared the difficult and rewarding journey to the first black hole images, Britton Jeter, Chunchong (Rufus) Ni, and Hung-Yi Pu: thank you for the immense work and support.

To the Event Horizon Telescope Collaboration, thank you for providing me with the opportunity to be a part of a historic human achievement, and for creating a community of scientists who have taught me much and have allowed me to be of service by teaching and helping them. Specifically, thank you to George Wong, Charles Gammie, Roman Gold, Koushik Chatterjee, Angelo Ricarte, Maciek Wielgus, Ben Prather, Héctor Olivares, Chi-Kwan Chan, and all previously mentioned for excellent discussion and consistent reliability.

I'd like to thank the various students of the University of Waterloo with whom I've overlapped for having just the right balance of physics and life to maintain my social and mental health. I'd also like to thank the University of Waterloo Juggling Club and its members for encouraging discipline and creativity in my work and for giving me energy and motivation.

Dedication

This is dedicated to what could have been but never was.

Table of Contents

List of Figures

List of Tables

Chapter 1

Introduction

Black holes are some of the most extreme objects in the universe. Their inherent lack of a surface means that matter can access vast gravitational potentials, reach highly-relativistic speeds, and create large-scale magnetic fields. These environments are conducive to a multitude of electromagnetic emission mechanisms, whose luminosity can be comparable to and even greater than the light of an entire galaxy while coming from a region the size of a solar system. Thus, the environments near black holes provide the best opportunity to test extreme plasma physics and gravity.

1.1 Black Holes

In General Relativity (GR), matter and energy cause curvature in spacetime described by a metric (Einstein, 1915; Wald, 1984). When the matter is a static spherically-symmetric bounded object, the spacetime in the vacuum surrounding the object is described by the Schwarzschild solution (Schwarzschild, 1916). If the object has enough mass within some radius, it develops an event horizon, the defining feature of a black hole. An event horizon is a null surface that divides the spacetime, such that lightlike trajectories emanating from one side (in this case, the inside of the black hole) cannot cross into the other. Matter and light can still traverse from the outside of the black hole to the inside, and generally, a black hole's mass always increases. From an observational standpoint, we can thus describe the spacetime everywhere with the vacuum Schwarzschild metric, which is entirely categorized by one parameter, the black hole mass M. Its event horizon is located at the Schwarzschild radius,

$$R_{\mathrm{Sch}} = \frac{2GM}{c^2}, \tag{1.1}$$

where G is the universal gravitational constant, and c is the speed of light.

Black holes need not be spherically-symmetric. For the case of an azimuthally-symmetric black hole, the metric is described by the Kerr solution, introducing a second parameter, the angular momentum J (Kerr, 1963; Wald, 1984). The Schwarzschild metric is the Kerr solution with $J = 0$, and there exists a maximum angular momentum for the black hole of $|J| = GM^2/c$. When later constructing physical simulations, we use the dimensionless spin, $a_* = Jc/GM^2$, which in general has both an amplitude and an orientation. Rotation changes the location of the event horizon and introduces an ergoregion where every light or matter particle must co-rotate with the black hole. In this ergoregion, it is possible for matter to extract some of the energy and angular momentum from the black hole via the Penrose process (Penrose & Floyd, 1971). The no-hair theorem in GR makes a powerful prediction that every static azimuthally-symmetric vacuum solution must be the Kerr solution, described by only three parameters, M, J, and the charge Q, regardless of the distribution of matter inside the event horizon[1] (Israel, 1968; Carter, 1971).

Black holes need not be stationary. As a normal star undergoes a core-collapse supernova phase, it can develop an event horizon and become a stellar-mass black hole (see Burrows & Vartanyan, 2021, and references therein). Once all of the material falls into the event horizon, the details of the initial star become irrelevant, and the spacetime outside is described by the stable static Kerr metric. When two black holes collide, the resulting object is also a black hole described by the Kerr metric (Pretorius, 2005; Abbott et al., 2016a). Thus, these seemingly perfect gravitational objects can form, remain, and interact with their environments in an observable way. It has also been proposed that primordial black holes could have formed in the early dense universe (Carr & Kühnel, 2020).

As a stellar-mass black hole resides in a galaxy, it can gravitationally interact with other stars or ambient galactic matter. The matter gains energy as it falls into the black hole and can release it through many electromagnetic emission mechanisms. A black hole can also collide with other black holes, releasing gravitational waves in the process. Both of these effects have been detected, providing strong evidence that these objects are, in fact, black holes (Webster & Murdin, 1972; Bolton, 1972; Remillard & McClintock, 2006; Abbott et al., 2016b). Though the process is not well understood, a black hole's mass can

---

[1] In principle, the electromagnetic charge on a black hole could be nonzero and affect the spacetime. However, for the matter in our universe, electromagnetism is much stronger than gravity, and any significant black hole charge would be equilibrated on timescales much shorter than black hole formation.

increase to $10^{6-10}\,M_\odot$, thus becoming a supermassive black hole (SMBH) and residing in the center of a galaxy (Kormendy & Ho, 2013). A SMBH is much lighter than all of the matter in its host galaxy, but it can energize nearby gas enough to have an enormous effect on the evolution of its galaxy. Evidence for the existence of SMBHs primarily comes from their effects on nearby stars and gas (Lynden-Bell, 1969; Soltan, 1982; Ghez et al., 1998; Genzel et al., 1997; Broderick & Narayan, 2006).

Light rays traveling near a black hole are lensed. For a far-away observer, we can determine the observed image by tracking the light rays back in time and determining from where the light could have originated. For a non-spinning black hole, there is an unstable photon orbit located at a radius of $1.5\,R_{\mathrm{Sch}}$. A light ray starting from the observer at a radius of

$$R_{\mathrm{shadow}} = \sqrt{27}\frac{GM}{c^2 D} \tag{1.2}$$

will slowly approach the photon orbit. Here, D is the distance between the black hole and the observer, and black hole spin changes this formula by about 10%. Light rays inside R_{shadow} fall into the event horizon, and light rays outside extend to some point at infinity. Thus, if the black hole is lit uniformly on all sides far away, the observer would see a "shadow" located at R_{shadow}. For a more complex emission structure, the observer sees some combination of the emission along the light ray's path. For the systems we later introduce, the emission is expected to peak near the event horizon and be optically thin. Thus, the brightness is expected to be largest near the shadow edge, since that is where the light ray has the largest path length through the emission region. Inside of the shadow, the image would have a brightness depression, caused by a combination of lensing and Doppler-boosting effects (Luminet, 1979; Falcke et al., 2000; Narayan et al., 2019).

1.1.1 Black Holes in their Galactic Context: Accretion and Jets

A SMBH in a galactic center environment is surrounded by a reservoir of gas called the circumnuclear disk. The origins and details of this reservoir of gas remain a topic of study, and they are set by the evolution of the galaxy. The gas is typically magnetized, with large ranges in temperature (see Moran 2008, and references therein for an example of a cool accretion flow; see Event Horizon Telescope Collaboration et al. 2019a for an example of a hot accretion flow). The relative motion of the black hole with respect to the galaxy, and thus to this reservoir of gas, is small due to SMBHs living for much longer than settling timescales (Reid et al. 1999; Merritt et al. 2007; see Pesce et al. 2021a for a counterexample). At about the Bondi radius of $10^5\,R_{\mathrm{Sch}}$, the gas enters the black hole's sphere of influence and becomes an accretion flow as it falls in (Bondi, 1952). For example, near the Milky Way's SMBH, the surrounding gas is probably composed of stellar winds emanating from the orbiting Wolf–Rayet stars (Cuadra et al., 2008; Calderón et al., 2020; Ressler et al., 2020).

As the matter falls in, the deep gravitational potential accelerates it and the accretion of fast-moving magnetized gas is transformed into dynamically-important magnetic fields. Furthermore, the gas is compressed as it falls in, densifying and heating it, sometimes up to the virial temperature of $\sim 10^{12}$K (Yuan & Narayan, 2014). Depending on the specifics of the material, it can radiate and cool as it falls in. The result is that without a surface to stop it, the gas near the event horizon can be hot, dense, highly magnetized, and moving at near the speed of light. This environment creates a wide range of possible emission mechanisms, the radiation from which can feed back into the motion of the accretion flow.

It is possible for the radiation to exert a force strong enough to counterbalance gravity, leading to the Eddington luminosity,

$$L_{\text{Edd}} = \frac{4\pi G M m_p c}{\sigma_T}, \tag{1.3}$$

where m_p is the proton mass, and σ_T is the Thomson scattering cross section. As an estimate, the Eddington luminosity is an upper limit followed by many accreting systems, though non-steady or anisotropic radiation can exceed this value. If all of the infalling gas turns into radiation, we get the Eddington mass accretion rate,

$$\dot{M}_{\text{Edd}} = \frac{L_{\text{Edd}}}{c^2}. \tag{1.4}$$

It is practically impossible for a black hole system to efficiently convert all of the accreted material into radiation, and efficiencies can range from 10^{-7} to 10^{-1}. The remaining energy goes into heating the gas, or perhaps in the kinetic energy of outflows.

As a crude estimate, the mass (energy) of accreting matter scales with the density, while emission mechanisms scale with the density squared. Thus, systems with a high mass accretion rate relative to the mass of the central object are more efficient in converting mass energy into radiation.

Since the incoming gas has a net angular momentum and magnetic flux, it is conserved as the gas reaches the event horizon. The rotation, gravity, and gas and magnetic pressure roughly balance together to create an accretion disk, though its thickness can vary substantially. The magnetic flux builds up in the inner flow and can accelerate particles and create relativistic jets of material emanating from the black hole, probably in the direction of the black hole spin or orthogonal to the accretion disk plane. These jets have been observed and extend far beyond the SMBH's host galaxy, and through jets, the SBMH influences the evolution of its galaxy (McNamara & Nulsen, 2012; Kormendy & Ho, 2013). It is unclear how accretion flows create jets and how jets accelerate their material, and observations have thus far placed this process as occurring quite close to the event horizon (Lu et al., 2023). Current theories for jet formation invoke ordered poloidal magnetic fields that extract energy either from the black hole or the accretion disk (Blandford & Znajek, 1977; Blandford & Payne, 1982). Between the jet and the disk (the corona), the gas is likely outflowing, so that not all of the mass and angular momentum of the incoming gas ends up in the black hole (Yuan & Narayan, 2014).

Galaxies with a central region brighter than the stellar emission have an Active Galactic Nucleus (AGN). This extra emission is likely caused by an accreting SMBH and provides the best opportunity to study them. In the unified theory of AGN, the SMBH is surrounded by an accretion disk close by, and by an absorptive dusty torus far away (Urry & Padovani, 1995; Netzer, 2015). Radio-loud AGN are observed in a high- or low-power state, likely due to their SMBH's mass accretion rate, with the cutoff predicted at $\dot{M} \sim 0.01\dot{M}_{\rm Edd}$ (Xie & Yuan, 2012; Yuan & Narayan, 2014). High-power AGN probably have sub-virial temperatures and thin disks as they efficiently transform gravitational energy into light. Low-power AGN do not efficiently radiate and instead transform the gravitational energy into heat, leading to thick puffy disks.

The rest of the AGN classifications can be explained by the inclination of the observer, whether that be near the jet, obscured by the dusty torus, or in the region between. For high-energy AGN, when the observer is not aligned with the jet, the observer sees a quasar, with further differences explained by whether or not the dusty torus obscures the SMBH. For low-energy AGN, at the same inclinations, the observer sees either a broad or narrow line radio galaxy. When the observer is aligned with the jet, it is a blazar, with high-power blazars called Flat-Spectrum Radio Quasars (FSRQs) and low-power blazars called BL Lacs (named after the prototype object BL Lacertae). Most of the sources we will be interested in are blazars or low-power AGN.

The two SMBHs with the largest known angular size on the sky are Sagittarius A* (Sgr A*) and Messier 87* (M87*; Ramakrishnan et al. 2023). The former is located in the Milky Way, and measured through direct observations of orbiting stars, has a mass of $4 \times 10^6 M_\odot$ and a distance of 8 kpc, giving an expected angular size of its shadow of $50\,\mu$as (Do et al., 2019; Gravity Collaboration et al., 2019). Sgr A* has a mass accretion rate of $\lesssim 10^{-6} M_\odot/{\rm yr} \approx 10^{-4}\dot{M}_{\rm Edd}$ (Genzel et al., 2010). Thus, Sgr A* is a low-luminosity AGN, and it has a spectral energy distribution (SED) that peaks near 1 mm (Event Horizon Telescope Collaboration et al., 2022a, and references therein). From increasingly higher-frequency radio observation of Sgr A*, the emission at $\gtrsim 1$ mm is expected to be optically thin and originate from a compact region within 150 μas (15 $R_{\rm Sch}$; Doeleman et al. 2008; Bower et al. 2014; Issaoun et al. 2019). This is also where the frequency dependence of the scattered size of Sgr A* turns over, indicating that we can see through the interstellar scattering screen (Bower et al., 2006; Johnson et al., 2018). Although there is evidence for energetic outflows a few million years ago, no jet has been conclusively detected (Falcke & Markoff, 2000; Yusef-Zadeh et al., 2016; Yang et al., 2022). Sgr A* also has daily X-ray flares and both theoretical models and near-infrared observations hint that they are caused by nonthermal electron hotspots accelerated by magnetic reconnection events (Broderick & Loeb 2006a; GRAVITY Collaboration et al. 2018; Haggard et al. 2019, and references therein).

M87*, on the other hand, is one of the most massive known SMBHs with $\sim 6.5 \times 10^9 M_\odot$ (Liepold et al., 2023; Osorno et al., 2023; Simon et al., 2023) [2]. It is 17 Mpc away, which leads to an angular size of its shadow of 40 μas (Cantiello et al., 2018). The mass accretion rate is $\lesssim 0.1 M_\odot/{\rm yr} \approx 10^{-3} \dot{M}_{\rm Edd}$, so M87* is a low power AGN (Di Matteo et al., 2003). Like Sgr A*, its SED peaks near 1 mm, where it is also expected to be optically thin (EHT MWL Science Working Group et al., 2021). However, M87* has a prominent jet, the first discovered, and it contains a substantial fraction of M87*'s emission (Curtis, 1918; Lu et al., 2023). Radio measurements place the jet as originating at the core, and spots have been traced traveling outward along the jet (Hada et al., 2016).

There are 3 main emission mechanisms in Sgr A* and M87*. The millimeter emission is caused by synchrotron radiation and bremsstrahlung causes high-energy emission. A subset of photons from each of these sources is inverse Compton-scattered to higher frequencies. There may also be emission due to pair annihilation in the evacuated jet region (Anantua et al., 2020). From a hydrodynamic view, Sgr A* and M87* are both in the collisionless regime, although magnetic interactions could effectively decrease the mean free path and anisotropically couple fluid elements. Thus, the electron distribution function need not be thermal. Coloumb collisions are generally insufficient to transfer internal energy from the ions to the electrons, leading to different temperatures for the two populations (Mahadevan & Quataert, 1997; Yuan & Narayan, 2014). For the millimeter images we will later be interested in, we will need to know the density, electron energy distribution, and velocity and magnetic field.

---

[2] Traditionally, the mass of M87* as measured through stellar or gas kinematics differs by about a factor of 2. Due to the many assumptions underlying the estimate from gas kinematics, we choose the higher mass estimate (Jeter & Broderick, 2021).

1.2 The Event Horizon Telescope

The most interesting phenomena in black holes occur near their event horizons. The Event Horizon Telescope (EHT) is a new instrument with the angular resolution to observe (at least) two SMBHs on event-horizon scales (Event Horizon Telescope Collaboration et al., 2019b). At millimeter wavelengths, M87* and Sgr A* are both mostly optically-thin and are bright enough to measure at 0.5-2.5 Jy. Ground-based observations at these frequencies are possible in several bands where atmospheric water vapor does not absorb all the light. However, the diffraction limit of a source is directly proportional to the wavelength of light and inversely proportional to the diameter of the telescope aperture. For 1 mm emission and a $50\,\mu$as black hole shadow, the Rayleigh criterion results in a telescope aperture of 5,000 km, approaching the size of the Earth.

1.2.1 Very-Long Baseline Interferometry

To achieve the necessary angular resolution, the EHT uses a technique known as Very-Long Baseline Interferometry (VLBI; Thompson et al. 2017). In VLBI, an instrument is composed of many radio dishes at different places around the world. Each observatory looks at the same source at the same time and records the signal. A traditional telescope would focus these signals, i.e., change their path length and add them up. In VLBI, the signals are time-shifted to replace the change in path length and then correlated. For any pair of dishes, known as a baseline, the result is a complex number corresponding to the amplitude and phase of the Fourier mode on the sky set by the projected distance and direction between the two dishes. With a large number of dishes, VLBI gives an approximation to the Fourier transform of the on-sky image, which can theoretically be inverted to produce the image.

Practically, instruments do not have enough dishes to fully fill out the Fourier plane. To get more data, an instrument can observe over timescales of a day (or night), and as the Earth rotates[3], the projected distances and directions between all the dishes change, and more of the Fourier plane is filled out. This Earth-aperture synthesis with VLBI traditionally assumes that the source is stationary over a night, an assumption broken by the EHT with Sgr A*. The maximum resolution is set by the longest baseline, and with a high enough signal-to-noise ratio and sufficient coverage, reconstruction methods can super-resolve with some image assumptions.

---

[3] Foucault (1851)

1.2.2 Submillimeter VLBI with the EHT

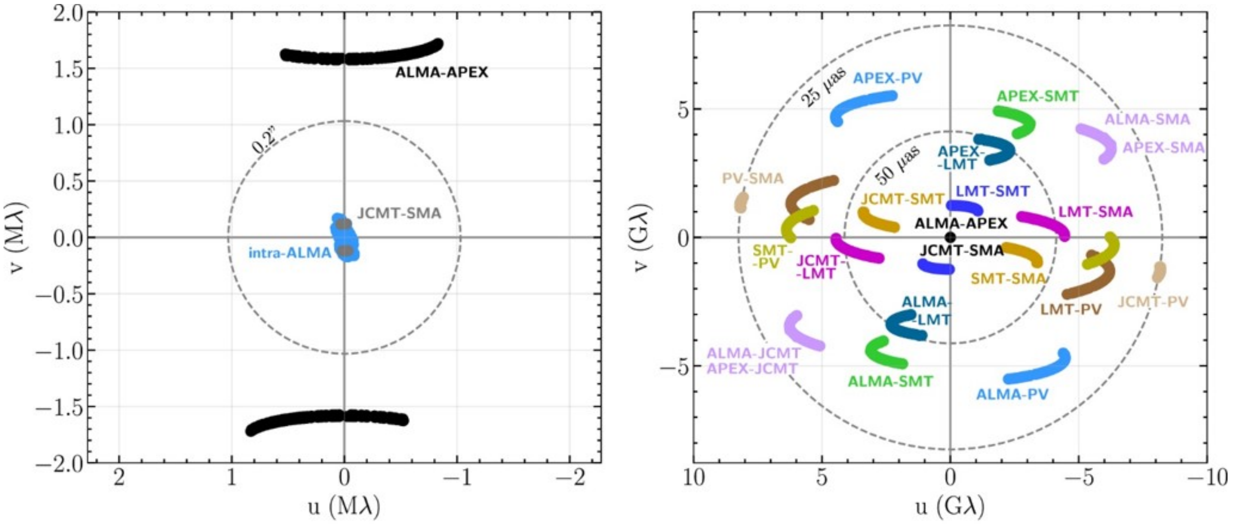

Figure 1.1: Representative EHT baseline coverage during the 2017 observing campaign of M87*. The left panel is a zoom-in of the right panel showing the short baselines (large scales) from co-located observatories. The dashed circles represent the minimum resolving scale. Labels indicate the pair of stations creating each baseline. This figure is taken from Figure 1 of Event Horizon Telescope Collaboration et al. (2019c).

The EHT is a VLBI instrument observing at 230 GHz, with extensions planned for 345 GHz. There are a maximum of 11 stations at 9 geographic locations, though currently available datasets have not included 3 of them (Event Horizon Telescope Collaboration et al., 2019b,d). This results in a typical EHT observation covering only a small portion of the Fourier plane; Figure 1.1 shows an example for M87*. An EHT observation has 10^2 to 10^4 scan-averaged data points. Directly inverting the data results in a "dirty map" and contains many pathologies. Combined with independent uncertainties for each data point, the uncertainty in the resulting image is large, with non-local correlations on small and large scales.

In order to solve the inverse imaging problem, a number of techniques have been used: CLEAN-based deconvolution algorithms DIFMAP (Shepherd, 1997) and `AIPS` (Greisen 2003; also used in Carilli & Thyagarajan 2022), regularized maximum likelihood algorithms `eht-imaging` (Chael et al., 2016, 2019) and `SMILI` (Akiyama et al., 2017a,b), sampling-based Bayesian posterior exploration methods THEMIS, (Broderick et al., 2020b,a), `DMC` (Pesce, 2021), and `Comrade` (Tiede, 2022), a wavelet deconvolution algorithm `DoG-HiT`

(Müller & Lobanov, 2022), a Bayesian deep learning framework (Sun et al., 2022), and a Bayesian variational inference method (Arras et al., 2022). Each of these incorporates some assumption or prior on the allowed image structure, e.g., image positivity, smoothness, flux compactness, or sparsity. Each thus results in a large space of possible images that fit the data, though only for the Bayesian methods will this correspond to a statistical posterior of images. It is important to note that these images are not the data, and much more discriminatory power exists for methods that extract physics directly from the Fourier data.

At EHT frequencies, the primary observational handicap is the atmospheric water vapor in the atmosphere, which changes over timescales of minutes, attenuates the signal, and scrambles the phase delays, equivalent to a complex gain. Since the number of baselines is greater than the number of stations for more than 3 stations, most of these gains can be independently reconstructed for each scan. However, the EHT cannot currently constrain the absolute phase gradient of the Fourier data, equivalent to a shift of the image position on the sky every scan.

Future instruments will substantially expand the spatial and temporal scales that such images probe. Upcoming ground-based projects, e.g., the ngEHT, promise to produce improvements in resolution by pushing toward shorter wavelengths, major improvements in sensitivity, and the capacity to probe variability in these sources on timescales of minutes (see e.g., Doeleman et al., 2019; Raymond et al., 2021). Space-based VLBI missions, building on the demonstrated recent success of RadioAstron (Gómez et al., 2016; Kardashev et al., 2017; Johnson et al., 2021), have the potential to increase imaging resolution by orders of magnitude, both improving the view of EHT targets and expanding the number of horizon-scale sources commensurately (Johnson et al., 2019; Pesce et al., 2021b).

1.2.3 Summary of Current EHT Results

Images of the two horizon-scale EHT targets, M87* and Sgr A* have now been published, revealing in unprecedented detail the astrophysical processes at work in the innermost regions of accretion flows and in the region responsible for launching jets (Event Horizon Telescope Collaboration et al. 2019a,b,d,c,e,f, hereafter M87* Paper I-VI, Event Horizon Telescope Collaboration et al. 2022b,a,c,d,e,f, hereafter Paper I-VI). Figure 1.2 shows the 2017 EHT images of M87* and Sgr A* as reconstructed by the Event Horizon Telescope Collaboration.

The first EHT observations of M87* revealed a bright ring with a central brightness depression. When identified with the lensed event horizon, silhouetted against the luminous

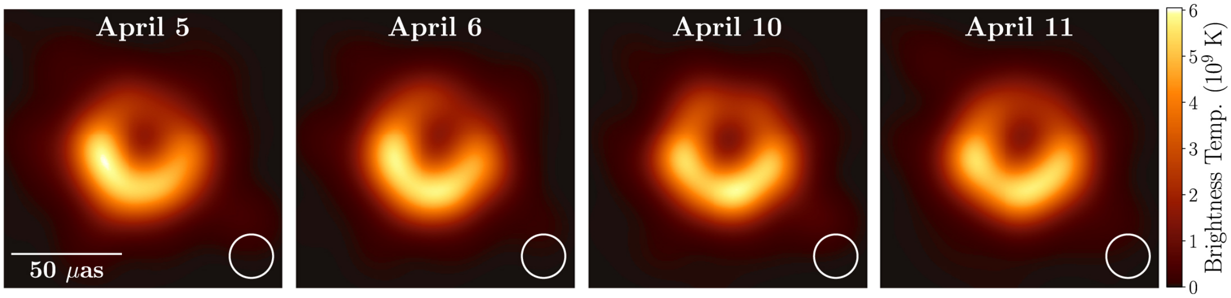

(a) EHT Observations of M87*. Figure taken from M87* Paper IV.

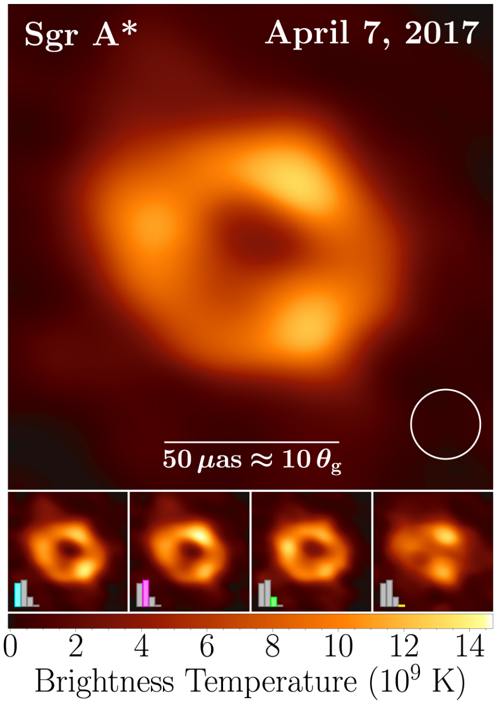

(b) EHT Observations of Sgr A*. Figure taken from Paper I.

Figure 1.2: First EHT observations of the two horizon-scale sources, M87* (a) and Sgr A* (b). Both sources show a bright ring of emission. M87* varies over the course of one week, with the brightest region moving from the left to the bottom of the ring. The four small panels below the image of Sgr A* represent four possible morphologies, with the histograms detailing the number of images reconstructed in Paper III. The size of the beam is shown by the white circle in the bottom right.

plasma in its vicinity, the bright ring had the size predicted by GR for the black hole mass measured by stellar dynamics. The ring had a brightness asymmetry, which informed the direction of motion of the gas. The gas surrounding M87* was hot and magnetized and consistent with fluid simulations. Over the week of observation (of order 10 times the light crossing time), M87* showed a moderate change in the structure of the emission, mainly by a brightness asymmetry that moved clockwise. Bayesian imaging techniques support a connection of the ring and jet, as well as higher evidence for a thin bright photon ring (Broderick et al. 2022a; also, see Lu et al. 2023).

The first EHT observations of Sgr A* also revealed a ring and brightness depression, though with less confidence and without a robust constraint on the emission structure. Sgr A* was found to contain both structural and unresolved variability; this contributes to the uncertainty in the average image, and the work in Chapter 3 was instrumental in separating the average image from the variability (Georgiev et al. 2022; Broderick et al. 2022b; M87* Paper IV). Simulations were mostly consistent with the observations, with the largest discrepancy caused by the variability in the total flux.

Images of both M87* and Sgr A* are grossly consistent with the expectations from theoretical modeling of the near horizon region (Luminet 1979; Falcke et al. 2000; Broderick & Loeb 2006b; Broderick et al. 2009; Broderick & Loeb 2009; Mościbrodzka et al. 2009; Dexter et al. 2009, 2012; Mościbrodzka et al. 2014, 2016; M87* Paper V; Paper V). This is largely a result of two generic elements that shape horizon-scale images: the strong gravitational lensing that imposes the central flux depression, or shadow, independent of the underlying accretion flow; and the relativistic motions within the emission region due to Doppler beaming and shifting (Luminet, 1979; Falcke et al., 2000; Broderick & Loeb, 2009; Narayan et al., 2019). The former is fully described by GR, and thus fully determined in the absence of new physics. In contrast, the latter is dependent on astrophysical details within the emission region, and thus model dependent.

The EHT has also observed several blazars to study the jet structure on 10-100 R_{Sch}. These are 3C 279 (Kim et al., 2020), J1924-2914 (Issaoun et al., 2022), NRAO 520 (Jorstad et al., 2023), and OJ 287 (in preparation). All show a collection of blobs along the path of a jet. These are typically polarized and are moving outward from a central core. In addition, the EHT has observed a forked jet in the radio galaxy Centaurus A (Janssen et al., 2021).

1.3 GRMHD Simulations and Synthetic Images

Astrophysical plasmas are usually modeled as a fluid and their evolution is described by the conservation of mass, energy, and momentum, and by Maxwell's equations. The material around the SMBHs of interest is hot and ionized, and most likely made of hydrogen, so it is modeled as a soup of ions and electrons. At its most basic level, this is done with the coupled Vlasov-Maxwell equations, which track the distribution of particles through phase space. This technique, applied through particle-in-cell codes, is very computationally expensive but can find self-consistent heating mechanisms and electron distribution functions in collisionless plasmas.

The first major assumption is to take moments of the Vlasov equation, yielding the continuity and Euler equations. When the particle distribution function is isotropic and local, one can prescribe an equation of state and thus ignore kinetic viscosity/heating effects. Placing this gas in a relativistic spacetime, which itself is affected by the mass via the GR field equations, leads to GRMHD simulations, which have been applied to a variety of systems (Hawley et al., 1995; Stone & Pringle, 2001; Davis et al., 2010; Tchekhovskoy et al., 2011; McKinney et al., 2012; Narayan et al., 2012; Palenzuela et al., 2013; Mösta et al., 2014; Porth et al., 2014; Paschalidis et al., 2015; Shiokawa et al., 2015; Combi & Siegel, 2023).

Black hole accretion flows are necessarily dynamic and non-laminar; the magnetohydrodynamic (MHD) turbulence that facilitates the transfer of angular momentum responsible for driving material inward is an essential feature of accretion flow models on compact objects (Shakura & Sunyaev, 1973; Balbus & Hawley, 1991). The resulting variability is expected to manifest on a wide variety of temporal and spatial scales, ranging from the cyclotron scale to the orbital scale. An a priori understanding of the properties of fully developed MHD turbulence remains an active area of study (see, e.g., White et al., 2020; Narayan et al., 2021, and references therein). Misaligned accretion flows, i.e., systems in which the angular momentum of the accreting gas is not aligned with the spin of the central black hole, will result in additional variability (Dexter & Fragile, 2013; Chatterjee et al., 2020; Ressler et al., 2020). It is, therefore, nontrivial that state-of-the-art (GRMHD) simulations of the magnetized plasmas near the black hole are in broad agreement with the observed image morphologies (M87* Paper V; Paper V).

1.3.1 GRMHD Simulations of EHT Sources

In order to generate synthetic images for the black hole systems targeted by the EHT, we use fluid simulations of accreting plasma in the Kerr black hole spacetime and, through simulation of emission and absorption mechanisms, the synthetic images that would be seen by a distant observer. The fluid simulations treat the plasma as an ideal, single-temperature fluid that is governed by the equations of ideal GRMHD. A detailed description of these equations can be found in Gammie et al. (2003), and we write them where applicable in Chapter 2. Each chapter uses slightly different assumptions, parameter spaces, and codes for the GRMHD fluid simulations and images.

Our ideal GRMHD simulations make several assumptions. First, we assume that the plasma can be treated as a fluid, even though the mean free path of particles in the flows is often much longer than the characteristic length scales in the systems. Second, we assume that the plasma has infinite conductivity, such that the magnetic field lines are "frozen" into the fluid.

In reality, kinetic processes in the (near-)collisionless plasma may increase the effective particle collision rate and allow deviations from the ideal fluid picture that can modulate typical variability timescales (see e.g., Foucart et al., 2016). Deviations from the infinite conductivity ideal fluid approximation (including the effects of viscous or resistive dissipation and heat conduction) may alter the thermodynamics of the flow (see Foucart et al. 2017 for a discussion; see Chandra et al. 2017, Most & Noronha 2021, and Most et al. 2021 for descriptions of viscosity in GRMHD; see, e.g., Ripperda et al. 2020 for an exploration of the effect of finite resistivity, especially in the context of plasmoid-mediated reconnection powering flares). It is particularly important to resolve the spatial scales on which the aforementioned dissipation mechanisms act (e.g., Ripperda et al. 2021). Variability timescales governed by turbulent dissipation or reconnection may, however, be very different in collisionless plasmas, and be strongly affected by non-thermal electron distribution functions, requiring a first-principles relativistic kinetic approach (e.g., Guo et al. 2014; Sironi & Spitkovsky 2014; Kunz et al. 2014; Werner et al. 2016; Zhdankin et al. 2017; Comisso & Sironi 2018; Bransgrove et al. 2021; Nättilä & Beloborodov 2021). All of these effects could cause significant changes in average structure or variability not captured in the simulations used in this work.

We also assume that the radiation is not dynamically important to the flow so that it can be treated entirely in a post-processing step. The non-radiative assumption is a valid approximation for our systems of interest, Sgr A* and M87*, where accretion is radiatively inefficient and $\dot{M} \ll \dot{M}_{\rm Edd}$; here $\dot{M}$ is the accretion rate (Yuan & Narayan, 2014).

We assume an ideal gas equation of state for the plasma with adiabatic index, $\Gamma_{\rm ad}$[4]. The system evolves on a stationary background spacetime that is unaffected by the accretion flow, whose total energy content is orders of magnitude lower than the central supermassive black hole. In this work, we choose the spin axis to point in the y-direction in the image and v-direction in the Fourier domain.

Concretely, the numerical simulations are initialized with an axisymmetric hydrodynamically stable Fishbone–Moncrief torus (Fishbone & Moncrief, 1976), that is parameterized by the inner radius of the disk, r_{in} and the radius at pressure maximum, $r_{\max}$. The torus is chosen to rotate either in the same ($a_* > 0$) or opposite ($a_* < 0$) direction as the black hole. An axisymmetric electromagnetic vector potential is used as the initial condition for the magnetic field (see Wong et al. 2022 for more details), and the thermal energy of the fluid is perturbed to seed development of instabilities such as the magneto-rotational instability (MRI; Balbus & Hawley, 1991) and kickstart accretion. After $\gtrsim 5,000\,GM/c^3$, the inner region reaches a quasi-steady state so that: the initial condition of the simulation has a negligible effect on the final state, and the simulation can reproduce the relatively steady-state light curve observed over the 2017 EHT observations of Sgr A*, which exhibits variability on the scales of 10% (Paper II; Paper V). There is evidence that Sgr A* may have recently been in a period of increased accretion (see e.g., Montero-Castaño et al., 2009), but such more complex feeding scenarios are beyond the scope of this work. Far from the black hole, the simulation is dominated by its initial condition, which supplies material to the inner turbulent flow. Since the torus is initialized with finite mass, in long-duration studies, it is possible that an appreciable fraction of the disk mass may accrete onto the black hole. This net decline may introduce an "artificial" systematic downward trend in long timeseries generated from the simulation.

The strength and structure of the magnetic field during the late-time, quasi-equilibrium phase dictate the magnetic flux, $\Phi_{\rm mag}$ threading the event horizon. Each simulation we consider equilibrates into one of two qualitatively different modes of accretion, each of which has a characteristic value of the non-dimensionalized magnetic flux, $\phi \sim \Phi_{\rm mag}/\sqrt{\dot{M}}$. When $\phi \sim \phi_c$, where $\phi_c \sim 15$ (Tchekhovskoy et al. 2011 but using the unit convention of Porth et al. 2019) is the critical value at which the local magnetic pressure near the event horizon balances the inward fluid ram pressure, we obtain a magnetically arrested disk (MAD) (Bisnovatyi-Kogan & Ruzmaikin, 1974; Igumenshchev et al., 2003; Narayan et al., 2003). In contrast, when $\phi \ll \phi_c$, the system produces a standard and normal evolution (SANE) disk (Narayan et al., 2012; Sądowski et al., 2013).

---

[4] The symbol γ is used for the adiabatic index in Chapter 2.

MAD

In ideal MHD, the magnetic flux is advected at the fluid velocity. If the material entering the black hole system is magnetized and the material ends up inside the black hole, then magnetix flux accumulates on the horizon. MAD accretion takes this effect to the extreme, where enough magnetic flux builds up to inhibit accretion. In MAD accretion, the magnetic flux is highly ordered with a large magnetosphere around the black hole. Away from the midplane, the magnetic pressure is sufficient to prevent the material from accreting. Near the midplane, accretion is possible, but it is choppy and characterized by transient ejection events where evacuated bubbles of magnetic flux that have accumulated on the horizon are expelled by the system. The strong magnetic fields near the horizon thus squish the disk, forcing the infalling plasma to accrete in thin, disordered accretion filaments or strands.

The jets in MAD simulations extract energy from the black hole via the Blandford-Znajek mechanism (Tchekhovskoy et al., 2011). Thus, MADs efficiently transform the infalling material into outflows, and make certain predictions about the relationtship between mass accretion rate and jet power on larger scales (Event Horizon Telescope Collaboration et al., 2022e).

SANE

In contrast, when the advected magnetic flux is low or randomly oriented, the steady state is characterized as a SANE disk with less ordered magnetic fields, though still dynamically important. SANE disks are turbulent and angular momentum transport is thought to be governed by the MRI. Due to the more disordered magnetic field structure, SANE disks are thick and have steady accretion. Unlike MADs, SANEs can have more drastic vertical electron temperature profiles, as the lower variability means less mxing between the midplane ant the corona. The jets in SANEs can still extract energy from the black hole, but are less efficient than MADs.

1.3.2 Synthetic Images of Black Holes

The GRMHD fluid simulations are used to generate images through general relativistic radiative transfer (GRRT), which proceeds in two steps. First, photon trajectories are computed backward from the camera into the simulation domain by solving the geodesic equation. Then, the covariant radiative transfer equation is solved forward to the camera along the precomputed photon geodesics, accounting for emission and absorption from the local plasma. The majority of results in this work consider images generated under the "fastlight" assumption, in which we approximate the fluid as unchanging during the entire GRRT process so that any image corresponds to a single (Kerr–Schild) timeslice.

The GRRT procedure takes as input:

- the fluid and electromagnetic data (mass density, fluid internal energy, fluid velocity, magnetic field) on a particular Kerr–Schild timeslice,
- a prescription for translating local fluid data into emission and absorption transfer coefficients at the appropriate observing frequency,
- the mass and spin of the central black hole, and
- camera parameters (including distance to source D and inclination angle i).

Here the inclination is defined as the angle between the line of sight to the observer and the black hole spin axis. In principle, we could introduce an azimuthal position angle of the camera, but there is not a unique choice for accretion flows aligned with the black hole spin.

The equations of non-radiative GRMHD are invariant under rescalings of length and time with the black hole mass, M, and rescalings of the fluid density, internal energy, and square of magnetic field strength with some density scale. The angular scale of the image, $GM/(c^2 D)$, is fixed by providing the distance to the source, D.

The expected near-collisionless nature of the plasma in the galactic center accretion system likely produces a flow in which the electron and ion components are best described as having different temperatures, $T_e \neq T_i$. Since our fluid simulations only track the total energy of the fluid, it is necessary to prescribe the temperature of the electrons, which are responsible for the emission. We use the electron temperature prescription described in Paper V, which is based on the model of Mościbrodzka et al. (2016), in which the ratio of ion to electron temperatures is a function of the local plasma $\beta = P_{\rm gas}/P_{\rm mag}$, according to

$$\frac{T_i}{T_e} = R_{\rm high} \frac{\beta^2}{1+\beta^2} + \frac{1}{1+\beta^2}. \tag{1.5}$$

Heuristically, this model allows the two species to have approximately equal temperatures in the low-density, highly magnetized jet region and produces increasingly cool electrons in the bulk disk as the free parameter $R_{\rm high}$ is increased. In this work, we assume that the electrons are well described by a thermal distribution, which is parameterized by a single (temperature) variable.

Numerical Implementation

In the present work, we consider fluid simulations performed by the following GRMHD codes: KHARMA and `iharm3d` (Prather et al., 2021), `KORAL` (Sądowski et al., 2013, 2014), and `BHAC` (Porth et al., 2017). These codes solve a set of conservation laws, where the conserved quantites correspond to the mass, energy-momentum, and magnetic fluxes. In addition, one of the electromagnetic equations is a spatial constraint. The fluid variables are evolved on a fixed grid, where the time coordinate is chosen depending on the static background metric used. Each code makes its own choices about numerical floors, grid geometry, and numerical schemes. See Porth et al. 2019 for a comparison of contemporary GRMHD codes.

The synthetic images of the GRMHD "snapshots" are generated by performing general relativistic ray tracing. We consider images or "frames" produced using the `ipole` and `BHOSS` GRRT codes; Gold et al. 2020 presents a comparison of contemporary GRRT codes.

1.4 Comparing Phenomenological and Physical Models to EHT data

With the analysis of EHT data being mostly limited to forward modeling, there is a growing need for models that incorporate prior knowledge of the expected physical processes constructed in a way that can be robustly compared to the EHT data. These models range from full GRMHD simulations to geometric models, each with their own assumptions and systematics.

1.4.1 GRMHD Models

GRMHD models begin from first principles, make simplifying assumptions, and simulate the emitting material as a fluid. Separate models perform the ray-tracing and radiative transfer to create reasonably self-consistent movies of black holes. These movies are turbulent with stochastic variability and it is not feasible to exactly match a simulation with an observation. Thus, data comparisons are done statistically, using the simulations' prediction for some type of variance. This has been done for EHT observations, with average image scoring for M87* (M87* Paper V; M87* Paper VI) and noise modeling for Sgr A* (Paper IV; Paper V). Alternatively, one could compare some quantity measured both from a library of simulations and from the data. For example, Paper V used an m-ring to ask whether the simulated and observed data had similar radii, widths, and asymmetries.

However, GRMHD simulations are computationally expensive and only sparsely sample the possible space of black holes and fluid parameters. With existing GRMHD simulation libraries, it is possible to measure and compare the effects of the black hole spin and the accretion state (MAD or SANE). However, it is difficult to test the robustness of results to assumptions, as measuring the effects of, e.g., resistivity, radiative forces, disk misalignment, and outer boundary conditions requires new fluid simulations. Comparatively, the radiative transfer is computationally cheaper, so exploring different types of electron thermodynamics, camera position, or emission frequencies can be done using standard Bayesian sampling methods (or grid searches as in M87* Paper V and Paper V). Thus, improving the capability of GRMHD simulations to extract physical information from EHT data entails either more computational ability or more sophisticated geometric and semi-analytical approximations. Existing data can already place strong constraints on models, though they have only explored a small parameter space.

1.4.2 Geometric Models

The conceptually simplest models are geometric, e.g., a sum of gaussians, rings, or disks, motivated by features seen either in data or simulations. These models have the advantage of being robustly measurable, but are difficult to interpret, and often require also being applied to GRMHD simulations to connect geometric results to physical results. Many implementations of these types of models exist in data fitting frameworks, such as THEMIS, `eht-imaging`, `DMC`, and `Comrade`, and these models have successfully been used to interpret EHT data (M87* Paper VI; Paper IV; Lockhart & Gralla 2022; Saurabh & Nampalliwar 2023). In some sense, images are models with a large number of parameters and strong priors and contain similar interpretational difficulties. Placing physically-motivated priors increases the interpretability of geometric models, under the assumption that the source is indeed described by the now semi-physical model. For example, motivated by the presence of a photon ring in GR and in many simulations, Broderick et al. (2022a) added a thin bright ring to an image model, greatly increasing the precision of the parameter estimation and increasing evidence for the presence of a photon ring. Variability in geometric models can be incorporated through time-variable parameters, and this technique is used to measure jet speeds in 3C 279 and OJ 287 (Kim et al., 2020).

Improving geometric models entails describing model features such that they encompass a wide range of possible physical effects expected in accretion flows. For this, libraries of

GRMHD simulations and images are useful as extraction and validation tools. Examples of such models include principal component analysis (Medeiros et al., 2023), and gaussian random fields as stochastic variability (Lee & Gammie, 2021). We contribute to these types of models in Chapter 3 and Chapter 4.

1.5 Semi-Analytical and Phenomenological Accretion Flow Modelling

Due to the stochastic nature of turbulent variability, it is not necessary to model every aspect of an accretion flow to produce a comparison to EHT data. Semi-analytical models make simplifying assumptions about the structure and variability of the flow in order to greatly reduce the computational cost and to explore a wider range of possible physical assumptions. Ray-tracing and radiative transfer are computationally cheap enough to be done normally, though they can be approximated as well.

1.5.1 Bondi Accretion

A common approximation of accretion flows is static, spherically-symmetric accretion of an ideal fluid, with no radiation or viscosity (Bondi, 1952; Frank et al., 2002). The resulting Bondi–Hoyle–Lyttleton accretion requires that the accretion flows conserves mass and momentum and adiabatically heats the gas as it is compressed. The accretion flow has three constants of motion,

$$\begin{cases} -4\pi r^2 \rho v, & \text{Conservation of mass} \\ \frac{p}{\rho^\gamma}, & \text{Conservation of energy} \\ \frac{v^2}{2} + \frac{\gamma}{\gamma-1}\frac{p}{\rho} - \frac{GM}{r}, & \text{Conservation of momentum} \end{cases} \tag{1.6}$$

where r is the radius, p is the gas pressure, ρ is the density, v is the radial velocity, and γ is the adiabatic index. Bondi accretion comes in three types depending on the sound speed of the gas ($c_s = \sqrt{\gamma p/\rho}$). When the incoming gas velocity is low, the gas always remains subsonic and approaches zero as the gas reaches the central object. When the incoming gas velocity is high, the gas reaches a pressure barrier where it becomes supersonic and either turns around or shocks to a subsonic solution. The third type of solution, or transonic, passes through a critical point at the sonic radius with a gas velocity equal to the sound speed. This point happens at a radius of

$$r_{\text{sonic}} = \frac{GM}{2c_s^2}. \tag{1.7}$$

This transonic solution is taken as the Bondi accretion solution. A similar solution exists for a wind.

In Bondi accretion, the mass accretion rate, and density, pressure, and velocity profiles are entirely determined by the central object's mass and the density and pressure at infinity. All of these quantites increase in magnitude towards the black hole with a power-law dependence on radius. Usefully,

$$\dot{M} \sim \left.\frac{G^2M^2\rho}{c_s^3}\right|_{r=\infty}, \tag{1.8}$$

with leading factors of order unity.

Bondi accretion provides order-of-magnitude estimates for how much mass can be supplied by a black hole's environment. Notably, it scales as M^2, whereas the radiation-limited Eddington limit scales as M. For supermassive black holes, the environment roughly exists outside of $\sim 10^5 R_{\text{Sch}}$, and thus Bondi accretion is typically used as a sub-grid model for galactic/cosmological simulations.

However, accretion on black holes is not spherical, as we can observe jets. As such, this type of model is insufficient to explain structure in EHT images but can provide reasonable estimates for ancillary observations (Quataert et al., 1999; Di Matteo et al., 2003; EHT MWL Science Working Group et al., 2021).

1.5.2 Angular Momentum Transport

Adding rotation complicates matters substantially. For vertically-symmetric, axisymmetric accretion, the midplane accretion flow has the following constants:

$$\begin{cases} \dfrac{p}{\rho^\gamma}, & \text{Conservation of energy} \\ r^2\Omega, & \text{Conservation of angular momentum} \\ \frac{v^2}{2} + \frac{\gamma}{\gamma-1}\frac{p}{\rho} - \frac{GM}{r} + \frac{\left(r^2\Omega\right)^2}{2r^2}, & \text{Conservation of radial momentum} \end{cases} \tag{1.9}$$

where Ω is the angular velocity. From the last expression, we see the creation of an effective potential barrier. The r^{-2} dependence of the angular momentum term is always positive and larger than the r^{-1} dependence of the gravitational term for small enough radius. Thus, material with angular momentum cannot reach the central infinitessimal object in this paradigm. Note that the finite size of a black hole and general relativistic effects do allow gas with small angular momentum to reach the horizon (see Chapter 2), but the gas from the reservoir easily has a larger value.

For matter to accrete, there must be a mechanism for outward angular momentum transport. Molecular viscosity is insufficient; the mean free path in the flow is

$$\frac{\lambda}{r} \approx \left(\frac{T}{10^{12}K}\right)^2 \left(\frac{v}{r\Omega_K}\right) \left(\frac{H}{r}\right) \left(\frac{r}{R_{\rm Sch}}\right)^{1/2} \left(\frac{\dot{M}}{\dot{M}_{\rm Edd}}\right)^{-1}, \tag{1.10}$$

where Ω_K is the Keplerian angular velocity and H is the disk height. Thus, the Reynolds number of the virialized fluid is often greater than 10^9, such that the gas would not accrete. The lack of viscosity is due to the long mean free path – the same problem that makes the fluid approximation inadequate. The same solution applies – magnetic fields in these plasmas are dynamically relevant, and can presumably create an effective fluid viscosity wherever the magnetic processes are saturated. However, three-dimensional turbulence is typically dominated by the largest spatial scales, and thus the form of the effective viscosity will depend on the global properties of the accretion flow.

1.5.3 The α-Viscosity Prescription

It is possible to introduce a form of angular momentum transport, and thus viscous heating, in a way that allows the fluid equations to have analytical solutions (see Yuan & Narayan 2014, and references therein). The fundamental assumption is that, in the saturated turbulent accretion flows, the shear is driven by large-scale features (such as vorticies, or MRI spatial modes). The largest of these features are of the size of the accretion disk height, and information travels through them at the sound speed. More specifically, the kinematic viscosity,

$$\nu = \alpha c_s H, \tag{1.11}$$

where c_s is the soundspeed, H is the disk height, and α is a proportionality constant containing all of the information about viscosity (Shakura & Sunyaev, 1973). Importantly, Equation 1.11 is not self-consistently found from GRMHD simulations, though it has had success in explaining observations [5].

---

[5] That is, $\alpha = 0$, or no viscosity and no angular momentum transport, is a poorer explanation than positive $\alpha < 1$.

A secondary major assumption is included to better model the gross features of accretion flows. As the gas infalls, it is possible that the effective viscosity heats mainly the ions in the gas, which do not efficiently transfer the energy to the radiating electrons. Similarly to the viscosity it is possible to parametrize this with a constant,

$$f \equiv \frac{\text{Advected energy}}{\text{Viscous Heating}} = 1 - \frac{\text{Radiated Energy}}{\text{Viscous Heating}}. \tag{1.12}$$

The solutions to this model (to first order in α) are

$$\begin{cases} \dot{M} \sim r^s \\ \rho \sim \alpha^{-1} \dot{M} M^{-1/2} r^{-3/2} \\ p \sim \alpha^{-1} \dot{M} M^{1/2} r^{-5/2} \\ v \sim \alpha M^{1/2} r^{-1/2} \\ \Omega \sim M^{1/2} r^{-3/2} \\ H \sim r \\ \text{Viscous Heating} \sim \dot{M} M r^{-4} \end{cases}. \tag{1.13}$$

The mass accretion rate can be a power law in radius, representing a loss of mass due to a vertical wind, and still keep the self-similar nature of the solutions. The proportionality constants are functions of α, f, and the power-law of the mass accretion rate s. Depending on the value of these constants, the flow is characterized as an advection-dominated accretion flow (ADAF), slim disk, or a thin disk.

Thin Disk

When $f \approx 0$ and $s = 0$, the solutions correspond a thin disk of the type introduced by Shakura & Sunyaev (1973). The gas in these disks is nearly Keplerian with small radial velocities. Since the gas takes so long to arrive at the black hole, it can radiate almost all of its internal energy, thus cooling the disk. The cool disk means low pressures, which leads to small disk heights. However, the density is high, and these disks are optically-thick.

A relativistic result of this thin disk has been obtained by Novikov & Thorne (1973). For a black hole, the Keplerian angular velocity does not always increase near the black hole; there exists an innermost stable circular orbit. Inside of this region, the angular momentum in this model is constant, corresponding to a no-torque condition at the inner boundary of the disk. Physically, this means that the relativistic effects cause the gas to rapidly infall, creating an evacuated region near the black hole. Outside of this boundary, the gas still infalls slowly.

The viscosity assumption relies upon an equilibrium of the different ways in which energy can be partitioned. Then, regardless of the source of the viscosity, it is local and proportional to that seen by the gas pressure. In these thin disks, this is plausible since the largest non-local effects are limited by a small vertical height, and the gas has many orbits to equilibrate microphysical effects. However, this need not be (and will not be) the case for hotter, thicker accretion flows.

This thin disk model has had success explaining X-ray binaries (see, e.g., Li et al. 2005), particularly their SEDs peaking at X-ray frequencies. It fares less well for AGN, such as the sources we wish to describe (Kokubo, 2018).

Slim Disks

The assumptions underlying thin disks are only valid when the gas can radiate faster than viscous heating. However, for increasingly large mass accretion rates near the Eddington limit, the advected energy grows faster than the radiated energy and the gas can no longer remain radiatively efficient. The accretion flow then transitions to a slim disk. In an appropriate limit, slim disks converge to the earlier thin disk model.

A slim disk is thicker, characterized by faster infall velocities, and does not have an evacuated region inside of the innermost stable circular orbit. Slim disks have super-Eddington mass accretion rates, and can even have super-Eddington luminosities, where the radiation escapes vertically and does not inhibit accretion. As the gas cannot efficiently radiate and $0 < f < 1$, the inner region of the accretion flow is more luminous than the outside.

Thin and slim disks make different predictions for the behavior of SEDs of AGN. However, their applicability remins uncertain (Castelló-Mor et al., 2017). Slim disks are not good approximations to Sgr A* or M87*, as these systems have a much smaller accretion rate.

Advection-Dominated Accretion Flows

An ADAF has $f \approx 1$ and a low mass accretion rate. In these systems, nearly all of the energy is advected into the black hole. Fundamentally, this happens because the low mass accretion rates lead to low densities and thus energy cannot be efficiently transferred from the viscously-heated ions to the radiating electrons. Thus, a two-temperature flow is essentially required. As the gas infalls and becomes hot, the pressure increases, and ADAFs are thick or quasi-spherical, and sub-Keplerian. Like slim disks, their emission peaks near the black hole.

ADAFs are a decent model for AGN where their SED shows a peak near millimeter wavelengths (Yuan & Narayan, 2014). These require a hot temperature, explained by the inefficient radiation, and magnetic fields to produce the synchrotron emission. However, classical ADAFs do not track the structure of the magnetic fields. SANE, but not MAD, simulations do show magnetic fields that are random and mainly respond to the fluid flow, so there is some confidence that the assumptions of ADAFs are not broken in the accretion flow. However, the α-viscosity prescription is rarely satisfied in simulations, and more complicated structures exist in the corona and jet (Narayan et al., 2012).

1.5.4 Semi-Analytical Models for EHT Sources

Semi-analytical also models motivate accretion flows with parametrized fluid variables, such as ones where each fluid quantity follows a radial power law (Broderick et al., 2009; Pu & Broderick, 2018). It is computationally feasible to directly fit these semi-analytical models along with a radiative transfer routine in, e.g., the Bayesian parameter estimation framework THEMIS. Similarly, a Kerr axisymmetric equatorial emission model has been implemented for fitting to EHT observations (Palumbo et al., 2022).

Similar semi-analytical models exist for the black hole jet. In the force-free paradigm, the jet is assumed to be magnetically dominated, resulting in an analytical model for the fluid (Tchekhovskoy et al., 2008; Broderick & Loeb, 2009). This paradigm has been extended to incorporate GR effects (Pu & Takahashi, 2020).

Semi-analytical models are not limited to the average emission. The daily flares in Sgr A* and the moving features in the jet of M87* and other blazars hint at moving spots of gas, with either a higher density or a different electron distribution function. These have been modeled as shearing midplane hotspots (Tiede et al., 2020), infalling arc-shaped gas clouds (Moriyama et al., 2019), shearing hotspots embedded in a jet (Jeter et al., 2020), and fully reconstructed with a neural network approach (Levis et al., 2022). When applied to EHT data, these variable features promise high-precision measurements of the spacetime, the accretion flow, and the determination of the type of jet-launching mechanisms.

None of these semi-analytical models are self-consistent, in that they do not create their main effects from known microphysical processes. Rather, they are approximations of features in EHT data we wish to describe. Their parameters dictate the large-scale behavior of complex processes, which data can then efficiently inform. The assumptions underlying each semi-analytical model can be tested with, e.g., GRMHD simulations, which we contribute to in Chapter 2. Targeted simulations can further inform extensions to semi-analytical models to include effects like jet loading mechanisms, magnetic reconnection events, kinetic viscosity, and impacts of galactic boundary conditions. Thus, improving the ability of EHT data to inform physics via semi-analytical models entails a cycle of creating increasingly more sophisticated models, validation with and motivation from simulations, and comparisons against data.

1.6 Thesis Summary

In this thesis, we use SMBHs and their environments to study accretion and gravity, with a focus on EHT data. In Chapter 2 and Appendix A, we construct an extendable static GRMHD midplane model and compare it with a library of GRMHD simulations. We find significant agreement with SANE models and motivate a variability prescription based on special solutions with zero radial velocity. In Chapter 3 and Appendix B, we study the variability present in a library of GRMHD images of Sgr A*. We find a universal power-law prescription for the variability and create measurements by which EHT data can discriminate accretion flow types. In Chapter 4, we study coherent variability in GRMHD images of M87*, and create a set of metrics to distinguish accretion flows. Surprisingly, we find that the features moving in synthetic EHT images do not correspond to gas moving at the fluid velocity. Finally, we summarize and conclude in Chapter 5.

Combined, these three projects are all in service of improving the models necessary to interpret horizon-scale EHT observations. Our semi-analytical model has a degree of validity in that it can encompass many of the expected turbulent effects. It can also be connected to large scales, incorporate non-GR spacetimes, and can efficiently explore a vast parameter space. When compared to future EHT data, a simulation-consistent semi-analytical model can provide information about the underlying spacetime and perform tests of GR. It can also answer ongoing questions about the role of turbulence in accretion, which can further be applied to an array of astrophysical systems. Finally, it can be used to provide a connection of the black hole, the galactic feeding environment, and the large-scale jet, and as a sub-grid model for galactic simulations. The observables mined from libraries of GRMHD images directly inform the interpretability of appropriate geometric models. They are also crucial to understanding the variability seen in EHT data. This data then answers questions about the inertial cascade of astrophysical turbulence, and how it depends on plasma properties, such as temperature, ionization state, and mean-free path length. Frameworks in which to fit geometric and semi-analytical models already exist, so our model improvements can be straightforwardly included and tested against EHT observations. We leave their implementation and validation to future work.

Chapter 2

A GRMHD-motivated Variability Decomposition for Semi-analytical Accretion Flows

2.1 Introduction

Data and images coming from the EHT contain a mix of information. In order to separate gravity, emission mechanisms, and accretion flow turbulence, we need physical simulations that include all of those effects. However, this is plagued by high computational cost, and so simulations typically cover only one realization of the possible accretion flow. Most of the information contained is random turbulence about some average accretion state.

Semi-analytical models attempt to simulate the average fluid state, and perhaps some simplified variability, as these are the quantities expected to be measured robustly in data. These "static" models[1] are much faster to compute and can be applied to study a vast parameter space of possible accretion flows, and can be physically extrapolated to test gravity. Furthermore, they can cover many scales, connecting the black hole to its environment. However, the details of the effects of variability on the average solution are typically significant. Semi-analytical accretion flows that use the α-viscosity prescription, such as ADAFs, do not approximate time-variable fluid simulations well, and require better models of the variability. No simple variability model exists that can explain the saturated turbulence state and be consistent with simulations.

---

[1] In a static model, the density, pressure, velocity, and magnetic field are constant in time, but the material is still moving with the fluid velocity.

Code	Flux	a_*	$\Gamma_{\rm ad}$	$r_{\rm in}$	$r_{\rm max}$	Resolution	Cadence
KHARMA	MAD	0, ±0.5, ±0.94	4/3	20	41	288×128×128	5
KHARMA	SANE	0, ±0.5, ±0.94	4/3	10	20	288×128×128	5

Table 2.1: GRMHD fluid simulation parameters. $r_{\rm in}$ and $r_{\rm max}$ are in units of GM/c^2, and the cadence is in units of GM/c^3.

We use ten GRMHD simulations and assume the variability they contain is indicative of accretion flows. We describe them in Section 2.2 and measure the importance of their variability on the average solution in Appendix A. We find that the impact of variability is moderate in SANE models and significant in MAD models, with the midplane seeing the least impact from variability. Limiting ourselves to the midplane for simplicity, we construct an analytic static midplane model and build intuition by showing its behavior in several regimes in Section 2.3. In particular, we note that solutions that exist in one part of the accretion flow do not extend everywhere and that permitted regions are bounded by critical solutions with zero radial velocity. To incorporate variability, we study one simulation from each of the two accretion flow types in Section 2.4 and construct an approximation to the midplane profiles, composed of an inner and outer critical solution. When applied to all the simulations, it is successful for SANE models but struggles with the overpowering variability present in MAD models. We then interpret the variability assumptions and discuss the implications toward a more complete variability prescription.

2.2 Description of the GRMHD Simulation Library

In this chapter, we are interested in finding semi-analytical models that contain a sufficient prescription for the steady-state turbulence. We use GRMHD simulations as a test set and as motivation for the ideas presented in the (semi-)analytical models. The simulations used in this chapter are also used in **??**, and their numerical details are given in Table 2.1. A summary of GRMHD simulations and their limitations is given in Section 1.3. In this section, we outline the GRMHD equations, which we will later use to derive a static model. We also describe the salient features in the simulations.

2.2.1 The Ideal GRMHD Equations

The equations of ideal GRMHD model a fluid in a background spacetime with a system of partial differential equations (PDEs). For a full derivation and description, see Gammie et al. (2003). The fluid primitive variables are the rest mass density, ρ, the gas pressure, p, the 4-velocity u^α, and the magnetic field 4-vector b^α. GRMHD assumes that the energy-momentum tensor is given by

$$T^{\alpha\beta} = \left(\rho + u_E + p + b^2\right) u^\alpha u^\beta + \left(p + \frac{b^2}{2}\right) g^{\alpha\beta} - b^\alpha b^\beta, \tag{2.1}$$

where $g_{\alpha\beta}$ is the metric, and $u_E = p/(\gamma - 1)$ is the internal energy for an ideal gas. The set of equations governing the conservation of energy and momentum are

$$\nabla_\alpha T^{\alpha\beta} = 0, \tag{2.2}$$

and the conservation of mass is given by

$$\nabla_\alpha \left(\rho u^\alpha\right) = 0. \tag{2.3}$$

Since we are in the ideal non-resistive limit, Ohm's law enforces that the electric field in the frame of the plasma vanishes. Thus, the magnetic field is governed by the source-free part of Maxwell's equation,

$$\nabla_\alpha M^{\alpha\beta} = 0, \tag{2.4}$$

where the dual of the electromagnetic field tensor,

$$M^{\alpha\beta} = u^\alpha b^\beta - b^\alpha u^\beta. \tag{2.5}$$

Additionally, the fluid velocity and magnetic field must satisfy $u^\alpha u_\alpha = -1$ and $b^\alpha u_\alpha = 0$.

A useful combination of the GRMHD equations that we will employ later is

$$\frac{\gamma - 1}{p}\left[u_\beta \nabla_\alpha T^{\alpha\beta} + b_\beta \nabla_\alpha M^{\alpha\beta} + \frac{\rho + u_E + p}{\rho} \nabla_\alpha (\rho u^\alpha)\right] = u^\alpha \partial_\alpha \ln\left(\frac{p}{\rho^\gamma}\right) = 0. \tag{2.6}$$

2.2.2 Kerr Metric

In GR, the unique solution for a spinning black hole is given by the Kerr metric. To avoid coordinate discontinuities at the event horizon, the GRHMD simulations use Kerr-Schild coordinates, that is, $ds^2 = g_{\alpha\beta} x^\alpha x^\beta$ with $x^\alpha = (t, r, \theta, \phi)$ and

$$g_{\alpha\beta} = \frac{1}{\Sigma}\begin{bmatrix} -\left(\Sigma - 2r\right) & 2r & 0 & -2ar\sin^2\theta \\ 2r & \Sigma + 2r & 0 & -a\left(\Sigma + 2r\right)\sin^2\theta \\ 0 & 0 & \Sigma^2 & 0 \\ -2ar\sin^2\theta & -a\left(\Sigma + 2r\right)\sin^2\theta & 0 & \left[\left(r^2 + a^2\right)\Sigma + 2a^2 r\sin^2\theta\right]\sin^2\theta \end{bmatrix}, \tag{2.7}$$

where $\Sigma = r^2 + a^2\cos^2\theta$. The area element is $\sqrt{-|g|} = \Sigma\sin\theta$. Where we later plot cross-sections in the r-θ plane, we cylindrical Kerr-Schild coordinates, labeled R, Z.

2.2.3 Salient Features of the GRMHD Simulations

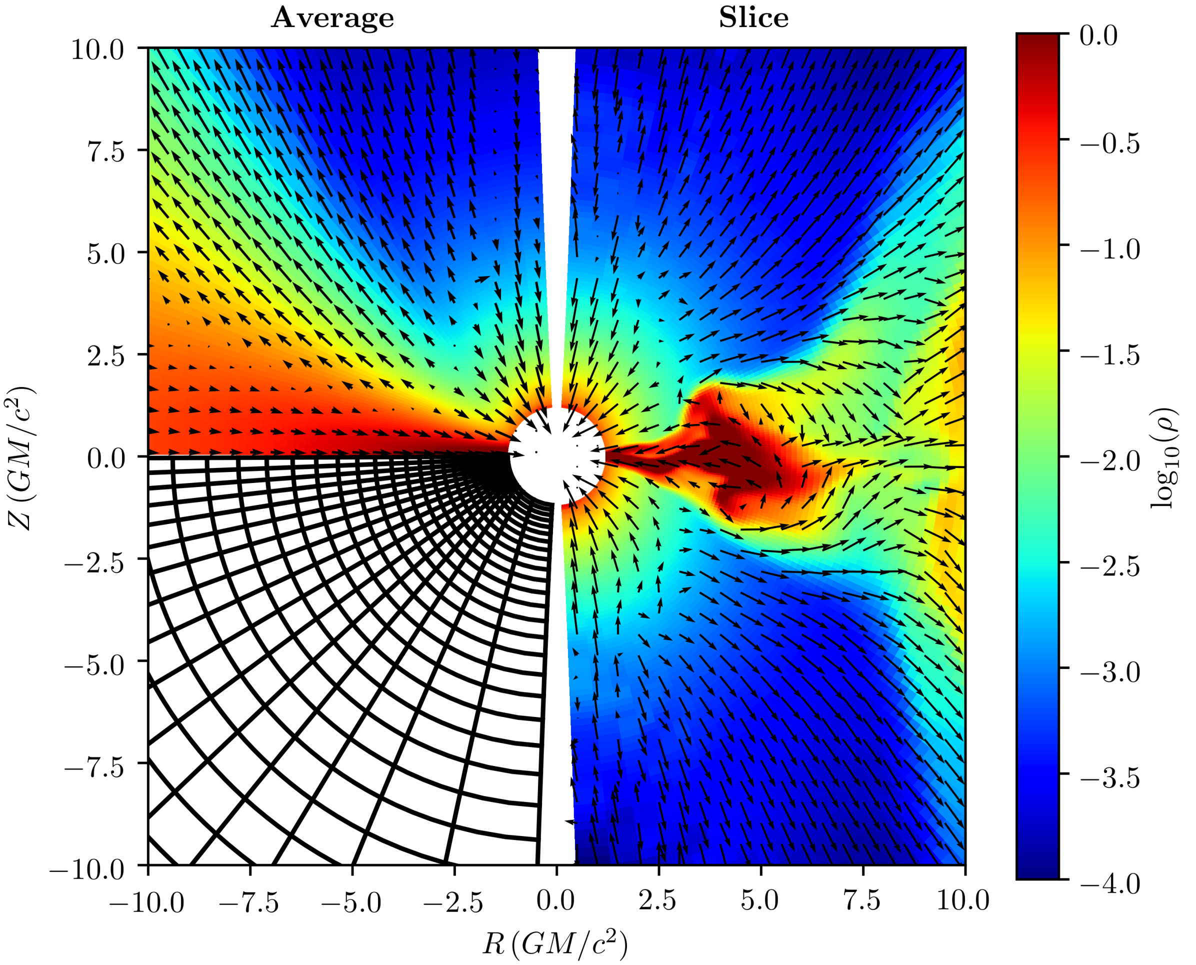

Figure 2.1: Indicative structure of a cross-section of a MAD, $a_* = 0$ simulation. The density is shown in the heat map, and the arrows represent the direction of the velocity in the R-Z plane with the length of the arrow roughly measuring the logarithm of its magnitude. The top left region shows the time- and azimuthally-averaged quantities, while the right shows the instantaneous values at $25,000\,GM/c^3$ and some azimuth. The bottom left region shows the simulation grid with every 4×4 points shown.

We use 10 GRMHD simulations spanning the two accretion states, MAD and SANE, and five spin values, $a_* \in \{-0.94, -0.5, 0, 0.5, 0.94\}$. The simulations are run for $30,000\,GM/c^3$

and consist of dumps of code variables saved every 5 GM/c^3 to a precision of about one part in 10^6. The GRMHD dumps are in a computational coordinate, and we first convert all quantities to Kerr-Schild coordinates before averaging. This makes the two analytical equations satisfied exactly at each frame and location. Additionally, there is an expected draining of the initial mass reservoir leading to a direct long-timescale trend in density, pressure, and magnetic field strength. Thus, we limit averaging and analysis to the last 1,000 frames to avoid the draining effect. This leads to a sampling error that has radial and poloidal structure.

The length and time units are in terms of the black hole mass. The units of the magnetic field, pressure, and density can all be rescaled without affecting the GRMHD equations, so an arbitrary scaling is chosen. Velocities are in units of the speed of light, c.

Figure 2.1 shows a collection of general features of the simulations as seen in the MAD, $a_* = 0$ simulation. The bottom left region shows the grid, which has finer spacing in the midplane and near the black hole. The top left shows a time- and azimuthal-average of the density and velocity. The right shows a slice in time and azimuth of the same quantities.

The simulations all average to a vertically-symmetric state where the midplane is dense with a scale height comparable to the radius. MAD models are typically thinner than SANEs due to stronger magnetic pressure. The gas in the disk accretes inward, with a heavily simulation-dependent outflow in the corona. The gas rotates faster close to the black hole, with a complex vertical dependence. The magnetic field in the midplane is dynamically important but changes changes directions often. It becomes stronger, more ordered, and poloidal near the poles as the magnetic flux accretes with the plasma.

2.2.4 Mean decomposition

The simulations exhibit hydrodynamical-like turbulence, with eddies and traveling waves. The MADs in particular have flaring episodes, where built-up magnetic flux is expelled, pausing accretion. A primary goal of this work is to find a static model that can approximate the effects of these variable effects.

One simple approach is to identify this static model with the time- and azimuthal-average of the simulations. Similar to an ADAF or thin disk model, we could then have static fluid equations with some extra terms included. Averaging the GRMHD equations yields time- and azimuthally- independent equations, but these are not fluid equations. Due to the nonlinear nature of the fluid differential equations, the averages of products do not necessarily equal products of averages.

In Appendix A, we calculate the effects of averaging the GRMHD equations using the solutions from GRMHD simulations and identify in which pieces the magntiude of variability becomes significant. We find that variability is least important in the midplane, especially for SANE models. We will now use this variability in the GRMHD simulations to understand how to amend a static accretion model.

2.3 Static Analytic Midplane Model

The analysis in Appendix A hints that a static model could be a reasonable approximation to the fluid near the midplane, at least for some GRMHD simulations. Towards that comparison, in this section, we rewrite the GRMHD equations in terms of auxiliary variables for which the midplane profiles can be integrated analytically. We then explore the behavior of the static midplane model in GRMHD-motivated regimes focusing particularly on regions that have zero radial velocity. We will later show that midplane profiles of GRMHD simulations are heavily influenced by these regions. We also explore a toy variability model which strongly suggests that accretion flows with zero radial velocity are important limits for variable simulations.

2.3.1 A Midplane GRMHD Accretion Model

Variable Change

The GRMHD equations are functions of the primitive variables, $\{\rho, p, u^\alpha, b^\alpha\}$. We introduce a set of auxiliary variables $\{\rho, S, \mathcal{T}_\alpha, \mathcal{M}^\alpha\}$ and, with the exception of ρ, label these as $\vec{A}$. We will show that these auxiliary variables are constant in radius. The forward conversion is:

$$\mathcal{M}^\alpha \equiv \frac{M^{r\alpha}}{\rho u^r} = \frac{u^r b^\alpha - b^r u^\alpha}{\rho u^r}, \tag{2.8}$$

$$\mathcal{T}_\alpha \equiv \frac{\mathcal{H} + b^2}{\rho} u_\alpha - \frac{b^r}{\rho u^r} b_\alpha, \tag{2.9}$$

$$S \equiv \frac{p}{\rho^\gamma}, \tag{2.10}$$

where $\mathcal{H} = \rho + \gamma p/(\gamma - 1)$. We have chosen the radial direction as the one of interest (by dividing by u^r), but the following relations hold for any direction. Note that $\mathcal{M}^r = 0$ by construction and $\mathcal{T}_r$ can be set by the velocity normalization condition. Here, we introduce Latin indices i, j to represent t, θ, ϕ.

For the inverse relations, we have

$$p = S\rho^\gamma \tag{2.11}$$

and

$$u_\alpha = \frac{\rho}{u^r}\frac{u^r\mathcal{T}_\alpha + b^r\mathcal{M}_\alpha}{\mathcal{H}+\rho^2\mathcal{M}^2} = \frac{\rho}{\mathcal{H}}\frac{\mathcal{H}\mathcal{T}_\alpha + \mathcal{M}_\alpha\rho^2\mathcal{M}\cdot\mathcal{T}}{\mathcal{H}+\rho^2\mathcal{M}^2}, \tag{2.12}$$

where $\mathcal{M}\cdot\mathcal{T} = \mathcal{M}^\alpha\mathcal{T}_\alpha = \mathcal{M}^i\mathcal{T}_i$, and similarly $\mathcal{M}^2 = g_{ij}\mathcal{M}^i\mathcal{M}^j$. We can alternatively find u_r by

$$u^{r2} = \left(g^{ir}u_i\right)^2 - g^{rr}\left(g^{ij}u_iu_j + 1\right) \tag{2.13}$$

and

$$u_r = \frac{u^r - g^{ri}u_i}{g^{rr}}, \tag{2.14}$$

thus removing all dependence on $\mathcal{T}_r$. Requiring $u^{r2} > 0$ constrains the possible space of $\vec{A}$, a point which will be crucial in later analysis. The magnetic fields are

$$b^\alpha = \rho\mathcal{M}^\alpha + \frac{b^r}{u^r}u^\alpha = \rho\mathcal{M}^\alpha + \frac{\rho^2\mathcal{M}\cdot\mathcal{T}}{\mathcal{H}}u^\alpha. \tag{2.15}$$

It is also useful to note that

$$b^2 = \rho^2\mathcal{M}^2 + \left(\frac{b^r}{u^r}\right)^2. \tag{2.16}$$

With these substitutions, we can write

$$M^{\alpha\beta} = \rho\left(u^\alpha\mathcal{M}^\beta - \mathcal{M}^\alpha u^\beta\right) \tag{2.17}$$

and

$$T^\alpha_\beta = \rho u^\alpha\mathcal{T}_\beta - \rho\mathcal{M}^\alpha b_\beta + \left(p + \frac{b^2}{2}\right)\delta^\alpha_\beta. \tag{2.18}$$

The r-direction of the electromagnetic equations becomes

$$0 = \partial_\alpha\left(\sqrt{-|g|}M^{\alpha r}\right) = -\partial_\alpha\left(\sqrt{-|g|}\rho\mathcal{M}^\alpha u^r\right). \tag{2.19}$$

There are no Christoffel symbols as $M^{\alpha\beta}$ is an anti-symmetric tensor. The rest of the electromagnetic equations become

$$0 = \partial_\alpha\left(\sqrt{-|g|}\rho u^\alpha\mathcal{M}^\beta - \sqrt{-|g|}\rho\mathcal{M}^\alpha u^\beta\right), \tag{2.20}$$

Variable	Expression	Physical Quantity
$\mathcal{S}$	p/ρ^γ	Specific entropy
$\mathcal{T}_t$	$T^r_t/(\rho u^r)$	Specific energy
$\mathcal{T}_\phi$	$T^r_\phi/(\rho u^r)$	Specific angular momentum
$\mathcal{M}^t$	$(u^r b^t - u^t b^r)/(\rho u^r)$	Specific temporal magnetic flux
$\mathcal{M}^\phi$	$(u^r b^\phi - u^\phi b^r)/(\rho u^r)$	Specific azimuthal magnetic flux

Table 2.2: Collection of the auxiliary variable symbols and descriptions.

which can combine with the continuity equation and Equation 2.19 to become

$$u^\alpha \partial_\alpha \mathcal{M}^\beta = u^r \mathcal{M}^\alpha \partial_\alpha \left(\frac{u^\beta}{u^r} \right). \tag{2.21}$$

Similarly, the energy-momentum equations become

$$u^\alpha \partial_\alpha \mathcal{T}_\beta = u^r \mathcal{M}^\alpha \partial_\alpha \left(\frac{b_\beta}{u^r} \right) - \frac{1}{\rho} \partial_\beta \left(p + \frac{b^2}{2} \right) + \Gamma^\gamma_{\alpha\beta} \left(\rho u^\alpha \mathcal{T}_\gamma - \rho \mathcal{M}^\alpha b_\gamma \right). \tag{2.22}$$

For our purposes, rather than writing the r- or θ-direction of this equation, we instead take the entropy equation

$$u^\alpha \partial_\alpha S = 0. \tag{2.23}$$

Midplane Equations

We now impose that the metric is independent of t, ϕ and that both the metric and the accretion flow are vertically symmetric at the midplane. That is, we assume $\mathcal{T}_\theta = \mathcal{M}^\theta = u^\theta = b^\theta = 0$. Under this assumption, we can write

$$u^r \partial_r \begin{pmatrix} S \\ \mathcal{T}_t \\ \mathcal{T}_\phi \\ \mathcal{M}^t \\ \mathcal{M}^\phi \end{pmatrix} = -u^i \partial_i \begin{pmatrix} S \\ \mathcal{T}_t \\ \mathcal{T}_\phi \\ \mathcal{M}^t \\ \mathcal{M}^\phi \end{pmatrix} + u^r \mathcal{M}^i \partial_i \begin{pmatrix} 0 \\ b_t/u^r \\ b_\phi/u^r \\ u^t/u^r \\ u^\phi/u^r \end{pmatrix} - \frac{1}{\rho} \begin{pmatrix} 0 \\ \partial_t (p + b^2/2) \\ \partial_\phi (p + b^2/2) \\ 0 \\ 0 \end{pmatrix}. \tag{2.24}$$

The derivatives in the right-hand side are in t, ϕ only (since $\mathcal{M}^r = 0$). For a static model, all of the terms in the right-hand side are 0, and the solution is

$$u^r \partial_r \vec{A} = 0, \tag{2.25}$$

and thus the auxiliary variables are constant in radius. This does however require a density profile, as the continuity equation requires a boundary condition on $\partial_\theta u^\theta$, or an amount of mass leaving the midplane. For what follows, we instead assume the boundary condition is an input to the model given as a fixed $\rho(r)$.

The second possible solution is that $u^r = 0$, which defines a special relationship between the auxiliary variables, the density, and r. We label these special values "critical". Static solutions can thus be piecewise, changing between zero radial velocity and constant auxiliary variables. Table 2.2 shows the auxiliary variable symbols and their physical meaning.

2.3.2 Static Model Solutions

In this section, we briefly note some properties of the static GRMHD midplane model to build intuition for the ideas in Section 2.4. In particular, we explore the behavior leading to zero radial velocity, and it will become the basis for the variability prescription. The GMRHD simulations typically have weak pressure and magnetic fields in the midplane, so we focus on these limits. We assume the Kerr metric in Kerr-schild coordinates as written in Subsection 2.2.2 for the following descriptions, and give all auxiliary variable values in units of $G = c = 1$.

No Magnetic Field and No Pressure

In the limit of no magnetic field, and no pressure, we are reduced to the geodesic equations, with $u_\phi = \mathcal{T}_\phi$ and $u_t = \mathcal{T}_t$. Most interesting is the equation for the radial velocity, which becomes

$$u^{r2} = -\frac{a_*^2 + (r-2)r}{r^2} - \frac{r-2}{r^3}\mathcal{T}_\phi^2 + \frac{4a_*}{r^3}\mathcal{T}_\phi\mathcal{T}_t + \frac{a_*^2(r+2)+r^3}{r^3}\mathcal{T}_t^2. \tag{2.26}$$

Figure 2.2 shows plots of the radial velocity and its dependence on the auxiliary constants in two regimes for $a_* = 0$. The following behaviors also hold for nonzero spin, but the exact formulas are spin-dependent.

The left panel shows u^r where $\mathcal{T}_\phi = 0$, or no angular momentum. For $\mathcal{T}_t < -1$, there exists a solution everywhere and $u^r \to \sqrt{\mathcal{T}_t^2 - 1}$ as $r \to \infty$. However, when $\mathcal{T}_t > -1$, the solution only exists for $r < 2/(1 - \mathcal{T}_t^2)$. At greater radii, the gas is given an energy of $-u_t < 1$, but even non-moving gas would have more energy than this.

The middle panel has a specific angular momentum of $\mathcal{T}_\phi = 5$. Similarly to the first case, when $\mathcal{T}_t < -1$, the solution exists everywhere. However, at larger $\mathcal{T}_t > -1$, the solution

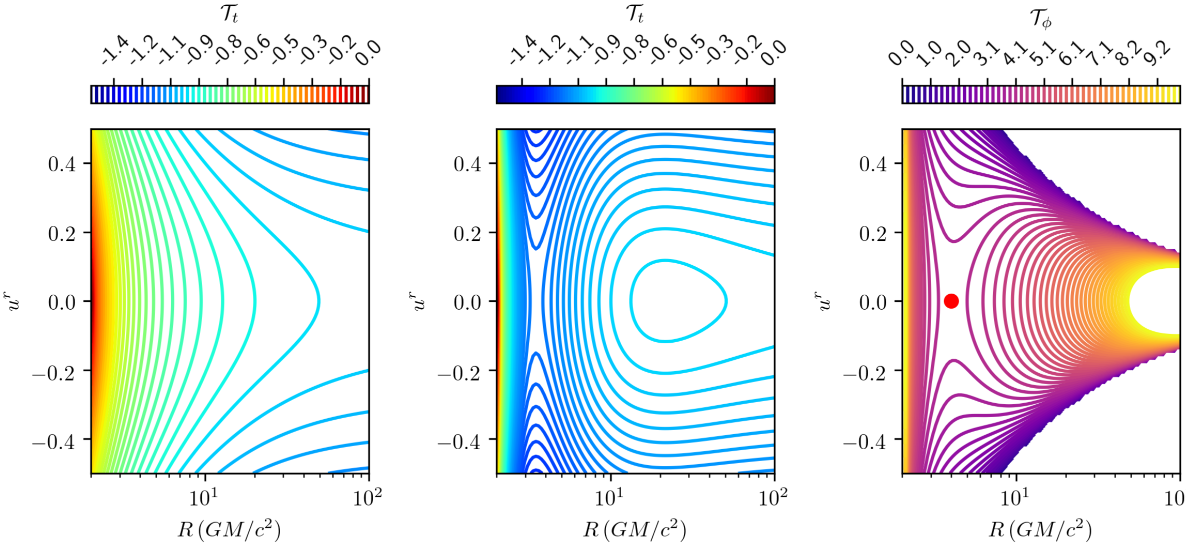

Figure 2.2: Radial midplane velocity in static model where $S = \mathcal{M}^t = \mathcal{M}^\phi = 0$. The left panel has $\mathcal{T}_\phi = 0$ or no angular momentum, and the middle panel has $\mathcal{T}_\phi = 5$. The right panel instead has $\mathcal{T}_t = -1$, or zero velocity at infinite radius, and varies $\mathcal{T}_\phi$. The red dot shows a critical location that separates the space of solutions.

splits into two allowed regions, the outer of which contains a stable circular orbit. The two allowed regions meet at a critical point, thus allowing a solution where gas inflows close to the black hole, but outflows far away. The first unallowed region is at large radii, like in the previous case. The second (inner) unallowed region is a consequence of the gas having too much angular momentum to fall into the black hole. For zero spin, if $|\mathcal{T}_\phi| < \sqrt{12}$, the strength of the black hole potential is always strong enough to pull the gas in, the stable circular orbit disappears, and no second unallowed region forms.

The third panel shows the case where $\mathcal{T}_t = -1$, and thus the gas has $u^r = 0$ at infinity. This case is most relevant for the GRMHD simulations. When the gas comes in with an angular momentum less than some critical value $|\mathcal{T}_\phi| < \mathcal{T}_{\phi,\mathrm{crit}}(r_\mathrm{crit})$, the gas reaches the black hole. More angular momentum than this causes the gas to stop at some radius (or turn around and outflow). This critical value is

$$\mathcal{T}_{\phi,\mathrm{crit}}(r_\mathrm{crit}) = \pm 2\left(\sqrt{a_* + 1} + 1\right), \tag{2.27}$$

and it happens at $r_\mathrm{crit} = \mp a_* + 2\sqrt{1 - a_*} + 2$. The location of this critical value is marked with a red dot.

No Magnetic Field with Pressure

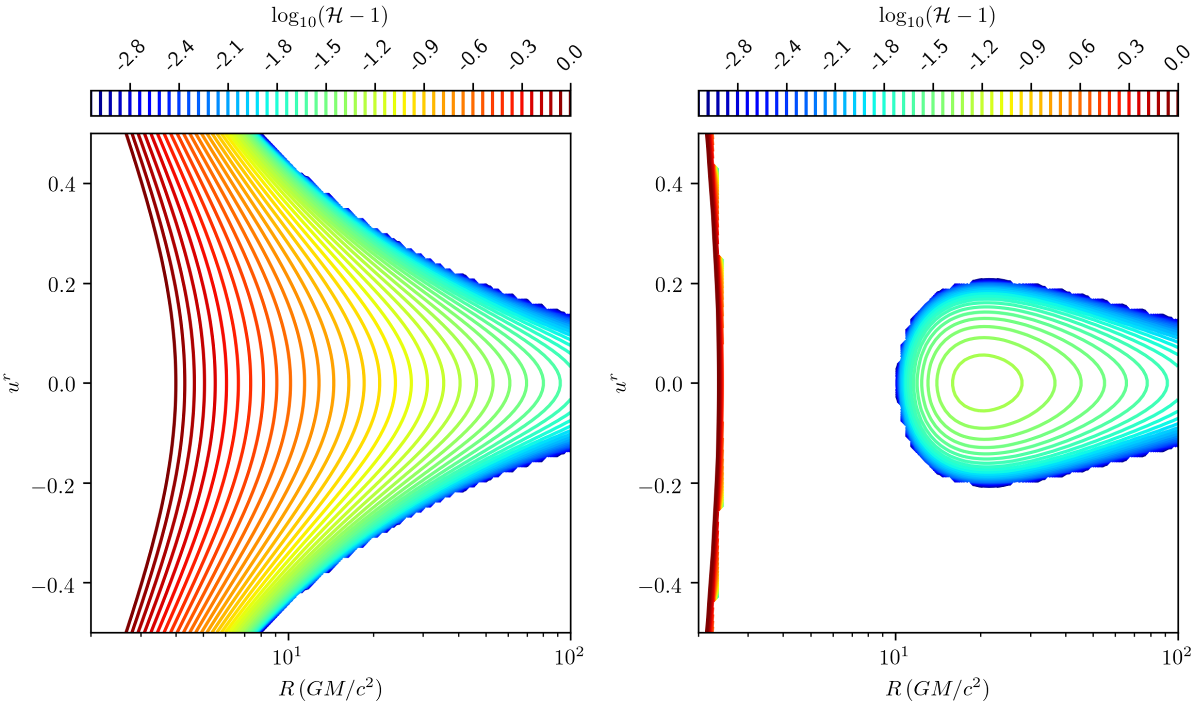

Figure 2.3: Radial midplane velocity in the static model where $\mathcal{M}^t = \mathcal{M}^\phi = 0$ and $\mathcal{T}_t = -1$. The left panel has $\mathcal{T}_\phi = 0$ and the right panel has $\mathcal{T}_\phi = 5$.

We repeat the same analysis as in Section 2.3.2, and confine ourselves to the limit of no magnetic field, $a_* = 0$, and $\mathcal{T}_t = -1$. In this regime,

$$u^{r2} = -\frac{r-2}{r^3(\mathcal{H}/\rho)^2}\mathcal{T}_\phi^2 + \frac{2}{r} - 1 + \frac{1}{(\mathcal{H}/\rho)^2}, \tag{2.28}$$

with $\mathcal{H} = \rho + \gamma p/(\gamma - 1)$ and $\gamma = 4/3$. For what follows, we set ρ as a constant in r, but the results from radial profiles we will later consider from GRMHD simulations qualitatively show the same behavior.

Figure 2.3 shows the dependence of the radial velocity and allowed regions on the specific angular momentum and the entropy (via the enthalpy $\mathcal{H}$). The left panel has zero angular momentum. Similar to the effects of varying the energy, high entropy gas is confined to the region near the black hole. The right panel has higher angular momentum, and, as previously, a second (inner) unallowed region forms for all values of $\mathcal{H}$ that have $\rho > 0$.

Critical Curves

Of interest will be the surface in $(r, \rho, \vec{A})$ that has $u^r = 0$, as this sets the boundary of allowed solutions. We later refer to these as critical solutions. As a general rule, S_{crit} and $\mathcal{T}_{t,\text{crit}}$ decrease with radius, and $\mathcal{T}_{\phi,\text{crit}}$ has a more complicated behavior we explore in the next section.

2.3.3 Variability Approximation

Characteristic Argument

Solving Equation 2.24 is of similar complexity to solving the fully variable GRMHD equations. However, we are only interested in a time- and azimuthally- averaged solution. Since this equation does not have any dependence on t, ϕ, one could imagine writing it in conservation form,

$$\partial_r \vec{A} = \partial_t \vec{f_t}\left(r, \vec{A}\right) + \partial_\phi \vec{f_\phi}\left(r, \vec{A}\right), \tag{2.29}$$

where $\vec{A}$ are the auxiliary variables, and $\vec{f_t}$ and $\vec{f_\phi}$ chosen such that their derivatives with respect to $\vec{A}$ give Equation 2.24. For bounded $\vec{A}$, averaging this equation over a large time leads to $\partial_r\langle\vec{A}\rangle = 0$. However, this formalism is invalid since the auxiliary variables are discontinuous, as solutions to many hyperbolic PDEs naturally are, and thus $\partial_r\langle\vec{A}\rangle \neq 0$. This hints that we can approximate the contributions of the variable behavior entirely by some ansatz for these discontinuities.

Locally, solutions to hyperbolic PDEs are transported along characteristics. For example, an equation of the form

$$\partial_t \boldsymbol{u} + \boldsymbol{\Lambda}^i \partial_i \boldsymbol{u} = 0 \tag{2.30}$$

has

$$\left(\boldsymbol{w}_n \cdot \boldsymbol{u}\right)\left(t + dt, x^i + \lambda^i_n dt\right) = \left(\boldsymbol{w}_n \cdot \boldsymbol{u}\right)(t, x^i), \tag{2.31}$$

where $\boldsymbol{w}_n$ are the eigenvectors of $\boldsymbol{\Lambda}^i$, and the speeds are the eigenvalues, λ^i_n. Each combination of variables is locally transported with its own characteristic speed. Since $\boldsymbol{\Lambda}^i$ depends on $\boldsymbol{u}$, the characteristics evolve with the variables leading to two types of discontinuities. When characteristics cross, the solution is over-constrained and the two possible solutions are separated by an evolving discontinuity known as a shock. When a region has no characteristics leading to it, the solution is under-constrained and the region is filled with a solution that has all of its characteristics emanating from a single point (known as a rarefaction wave). The specifics of these cases are chosen by the original conservation form of the GRMHD equations and followed in the GRMHD simulations (see Toro 2009 for an exhaustive overview).

There is a third type of discontinuity specific to our problem. When we evolve a solution in time, it can be shown that the GRMHD equations have finite characteristic speeds for nonzero density. In fact, these are exactly the entropic, sonic, Alfvén, and fast and slow magnetosonic wave speeds found in a perturbative analysis. When we change to evolve in radius instead of time, the characteristic speeds are inverted. Where they were zero before, they are infinite now.

To approximate the effects of variability, we assume that

1. At every radius, there is a random distribution in t, ϕ of the fluid auxiliary variables, $\vec{A}$.
2. In the absence of discontinuities, the $\vec{A}(t, r, \phi)$ can be connected via characteristics to a linear combination of $\vec{A}(t, r_{\text{in,out}}, \phi)$ at some inner or outer radius. This leaves the distribution of $\vec{A}$ unchanged.
3. Shocks and rarefaction waves do not significantly change this distribution from one radius to another. The next section shows support for this when the timescale of variability is greater than the infall time, but this assumption will ultimately fail.

Under these assumptions, the primary contribution to the variability will be the discontinuities arising from characteristics that have speed 0. Since the t, θ, and ϕ directions are determined by the Killing fields and vertical symmetry in the midplane, this defines the r-direction, and thus zero speed.

Evolution of a Solitary Auxiliary Variable

In this section, we show that when only one of the auxiliary variables is allowed to vary, the global solutions converge to a combination of a critical and constant solution. For simplicity, we take $\mathcal{M}^t = \mathcal{M}^\phi = S = 0$, $\mathcal{T}_t = -1$, zero spin, and allow $\mathcal{T}_\phi$ to only be a function of t, r. In this limit, the metric in Kerr-Schild coordinates is

$$ds^2 = -\left(1 - \frac{2}{r}\right) dt^2 + \left(1 + \frac{2}{r}\right) dr^2 + \frac{4}{r} dt dr + d\Omega^2. \tag{2.32}$$

The GRMHD equations become that of geodesics, and thus $\mathcal{T}_\phi = u_\phi$. We have

$$u^\alpha = \left\{ \frac{1}{r-2}\left(r - 2\sqrt{\frac{2}{r} - \frac{(r-2)\mathcal{T}_\phi^2}{r^3}} \right), -\sqrt{\frac{2r^2 - (r-2)\mathcal{T}_\phi^2}{r^3}}, 0, \frac{\mathcal{T}_\phi}{r^2} \right\}, \tag{2.33}$$

where we pick the convention for an ingoing flow. The only non-trivial fluid equation is

$$u^r \partial_r \mathcal{T}_\phi + u^t \partial_t \mathcal{T}_\phi = 0, \tag{2.34}$$

where we take $\mathcal{T}_\phi$ as the integrating variable for numerical simplicity. For this toy example, we take the shock speed to be [2]

$$v_{\text{shock}} = \frac{1}{2}\left(\left.\frac{u^r}{u^t}\right|_{\text{left}} + \left.\frac{u^r}{u^t}\right|_{\text{right}} \right). \tag{2.35}$$

We can now numerically solve Equation 2.34. We do this with a one-dimensional reconstruct-evolve-average (REA) solver validated on the inviscid Burger's equation.

Static Boundary Conditions

Here, we perform two tests with static boundary conditions to describe the main types of behavior in the dynamical system. For the first test, labeled case 1, we initialize $\mathcal{T}_\phi(t = 0) = 3$ and set an outer boundary condition of $\mathcal{T}_\phi(r = 10) = 4.5$. For the second test, labeled case 2, we initialize $\mathcal{T}_\phi(t = 0)$ with the infinite time limit of Case 1 and set an outer boundary condition of $\mathcal{T}_\phi(r = 10) = 3.5$.

Figure 2.4 shows the solutions for cases 1 and 2. Each case converges to the initial state of the other at infinite times. Also shown in the left panel are characteristics originating from $(t, r) = (0, 10)$ for $\mathcal{T}_\phi = 3.5$ (solid), $\mathcal{T}_\phi = 4$ (dashed), and $\mathcal{T}_\phi = 4.5$ (dotted). These divide the solution space into 4 regions.

In case 1, the gas inserted at the outer radius is infalling slower than the initial gas. This creates a rarefaction wave carrying gas with $3.5 < \mathcal{T}_\phi < 4.5$ with each piece of gas moving along its own characteristic. For gas in the first (from left to right) region, the

---

[2] The shock speed is determined by the original form of the GRMHD equations. However, by changing variables, we have assumed derivative continuity, thus altering the shock speed. This is exacerbated by the assumptions of the toy model and $\rho = 0$ in particular, so we use an average speed as an approximation. Our choice does not change the conclusions of this toy example.

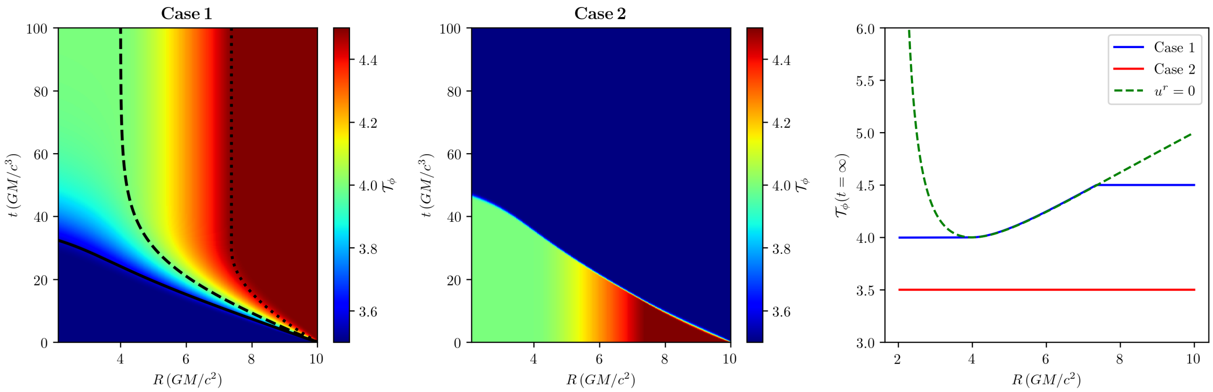

Figure 2.4: Example solutions to the GRMHD midplane equations where only $\mathcal{T}_\phi(t, r)$ can vary. The left two panels show the time-evolved solutions with different boundary and initial conditions. The right panel shows the value of $\mathcal{T}_\phi$ as $t \to \infty$ for the two cases. The black lines in the left panel are specially-chosen characteristics that divide the solution into 4 regions.

material only sees the initial infalling gas and thus has $\mathcal{T}_\phi = 3.5$. The second region is the portion of the rarefaction wave that eventually ends up in the black hole. This region has $u^r < 0$ everywhere, or more specifically, $\mathcal{T}_\phi < \mathcal{T}_{\phi,\text{crit}}(r_\text{crit}) = 4$. Had we not set the entropy S, and therefore the gas pressure, to zero, this region would be slowly evacuated of matter as all the characteristics enter the black hole. For $\mathcal{T}_\phi > 4$ (the third region), each of the characteristics turns around at

$$2r_\text{turn}^2 = (r_\text{turn} - 2)\mathcal{T}_\phi^2 \tag{2.36}$$

for Schwarzschild. As they attempt to turn around, they are immediately hit (i.e. shocked) by gas slightly further out and kept in place. The further away gas is itself kept in place by gas at a larger radius all the way out to the boundary condition keeping the whole system stable. Since this gas originated from rarefaction, with the introduction of pressure, this gas would also decrease in density over time as it is expanded to cover more volume. The fourth region is similar to the third region, except that all the gas has $\mathcal{T}_\phi = 4.5$ enforced by the boundary condition.

The third panel in Figure 2.4 shows the result of case 1 as $t \to \infty$. For $r \gtrsim 7.38$, the gas is entirely dominated by the boundary condition, and the auxiliary variables are constant. Between $4 < r \lesssim 7.38$, the gas converges to have $u^r = 0$, analytically calculable with Equation 2.36. For $r < 4$, the turning radius for $\mathcal{T}_{\phi,\text{crit}}$, the only infalling gas left has $\mathcal{T}_\phi = \mathcal{T}_{\phi,\text{crit}}(r_\text{crit}) = 4$. Note that since we imposed $u^r \leq 0$, the final state of the inner gas is reached regardless of the initial conditions for sufficiently large incoming angular momentum.

Case 2 is simpler. When gas with a faster infalling speed hits the final state of case 1, it shocks and replaces the existing gas. Since this case has an outer boundary condition less than $\mathcal{T}_{\phi,\mathrm{crit}}$, its final state is a constant $\mathcal{T}_\phi$.

These two tests indicate that the boundary conditions each try to push the radial profile either toward a constant value for $\mathcal{T}_\phi(r_{\mathrm{out}}) \leq \mathcal{T}_{\phi,\mathrm{crit}}(r_{\mathrm{crit}})$ or towards the three-section curve as seen in case 1 for $\mathcal{T}_\phi(r_{\mathrm{out}}) > \mathcal{T}_{\phi,\mathrm{crit}}(r_{\mathrm{crit}})$. Only close to the outer radius does the infinite-time solution depend on the specific value of the boundary condition.

Variable Boundary Conditions

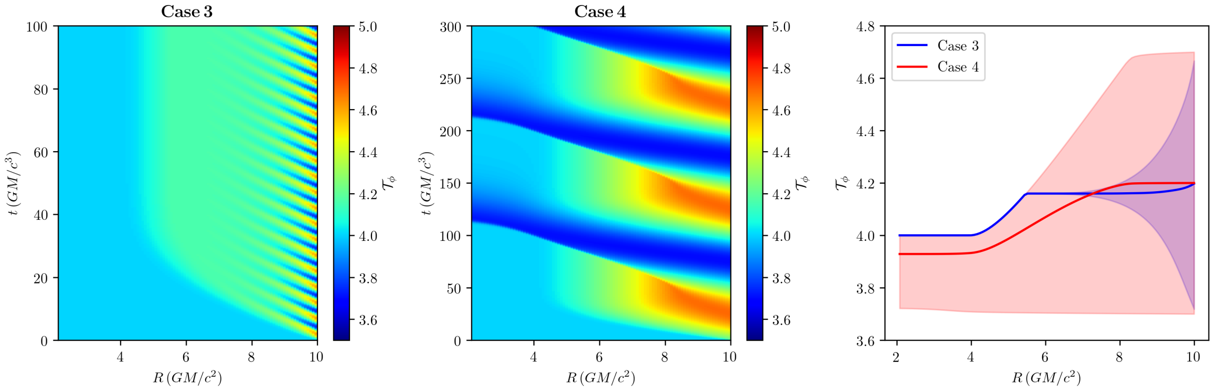

Figure 2.5: Similar to Figure 2.4 with a sinusoidally-variable outer boundary condition. The initial value is set to $\mathcal{T}_\phi(t=0)=4$ for both cases. The right panel shows the average values as well as a band covering all covered values. The average begins at $t = 50\,GM/c^3$ for case 3 and $t = 100\,GM/c^3$ for case 4, and both cases are averaged over an integer number of wavelengths of the outer boundary sinusoid.

In the GRMHD simulations, the outer boundary is variable, and the solution to which this model converges may change for a variable boundary condition. We perform two more tests (cases 3 and 4) to test the effect of short- and long-timescale variability.

Figure 2.5 shows the solutions corresponding to cases 3 and 4. Both cases have their initial condition set to the critical value (i.e., $\mathcal{T}_\phi = 4$), and this initial gas quickly falls into the black hole. Case 3 has

$$\mathcal{T}_\phi(r=10) = 4.2 + 0.5\sin\left(2\pi\frac{tc^3}{5\,GM}\right) \tag{2.37}$$

, and case 4 has

$$\mathcal{T}_\phi(r=10) = 4.2 + 0.5\sin\left(2\pi\frac{tc^3}{100\,GM}\right). \tag{2.38}$$

In case 3, the fast-moving gas overtakes the slower gas at around $r \approx 7\,GM/c^2$ before both can fall into the black hole. This results in a static solution for $r \lesssim 7\,GM/c^2$. In case 4, the gas has time to fall into the black hole before the boundary conditions appreciably change, and thus every radius feels the effects of every boundary condition.

The GRMHD simulations are, at minimum, variable on the timescale of an orbit, though often for much longer. For this Schwarzschild example, this timescale is hundreds of GM/c^3, greater than the infall time, and so we expect GRMHD simulations to behave more like case 4.

Note on Multiple Variables

If only one of the auxiliary variables is allowed to vary at a time, then each equation can be written as

$$\partial_t\vec{A} + \frac{u^r}{u^t}\partial_r\vec{A} = 0, \tag{2.39}$$

with the exception of

$$\partial_t\mathcal{M}^t + u^r\left[u^t - u^r\mathcal{M}^t\frac{\partial}{\partial\mathcal{M}^t}\left(\frac{u^t}{u^r}\right)\right]^{-1}\partial_r\mathcal{M}^t = 0. \tag{2.40}$$

All of these have a zero characteristic speed when $u^r = 0$. Thus, when each auxiliary variable is treated independently, the gas behaves like in the previous section. As it falls in, it either maintains constant auxiliary variables or, if that would lead to an unallowed solution, continues along a critical curve with $u^r = 0$.

Of course, when adding in the effects of interactions, the distribution of variability at a particular radius can evolve differently as information travels throughout the flow. To understand these interacting effects, we must now turn toward the GRMHD simulations.

2.4 Comparison with GRMHD Simulations

We now apply the intuition from Subsection 2.3.1 to compare the model to the GRMHD simulations in the midplane. Specifically, we compare to the median midplane profiles of the auxiliary variables and take appropriate quantiles when discussing variability. We will first focus on the SANE models, using $a_* = 0$ as a motivating test case. We'll first show that the effects of variability are closely related to the critical solutions with zero radial velocity. From this, we construct an approximation that matches the profiles of all the SANE simulations with fewer than a doubling of the number of parameters. These approximations will not work as well for the MAD models. We then discuss the physical interpretation of this approximation and highlight the successes and failures of the ideas behind the model.

2.4.1 SANE Models

Zero Spin

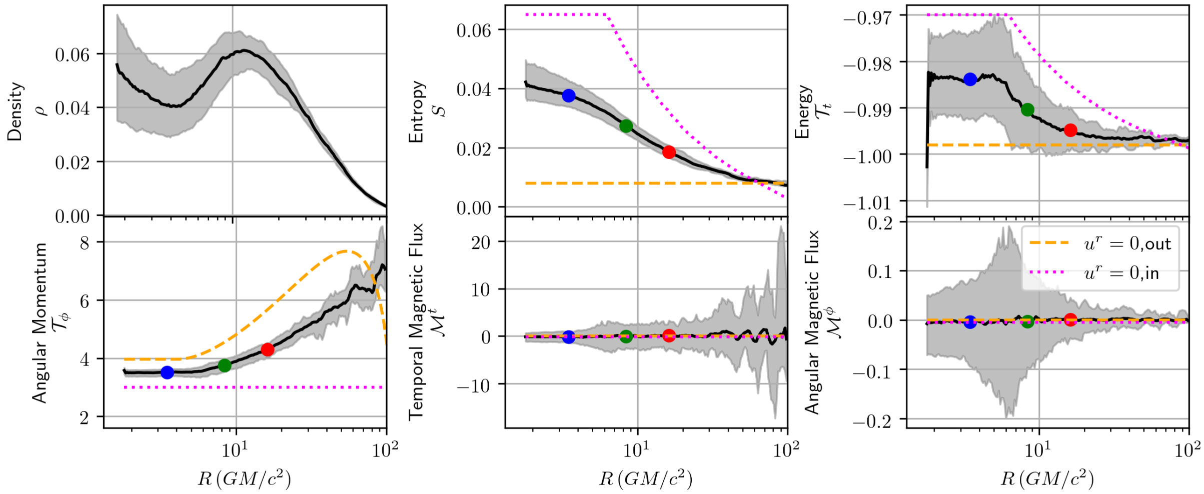

Figure 2.6: Midplane radial profiles of the density and auxiliary variables for a SANE, $a_* = 0$ simulation. Solid lines are the median in the time range $25,000$-$30,000\,GM/c^3$, and the shaded region shows the inner 68% percentile. The three dots indicate the locations expanded in detail in Figure 2.7. The dashed and dotted line represent two possible critical solutions with specific parameters listed in Table 2.3.

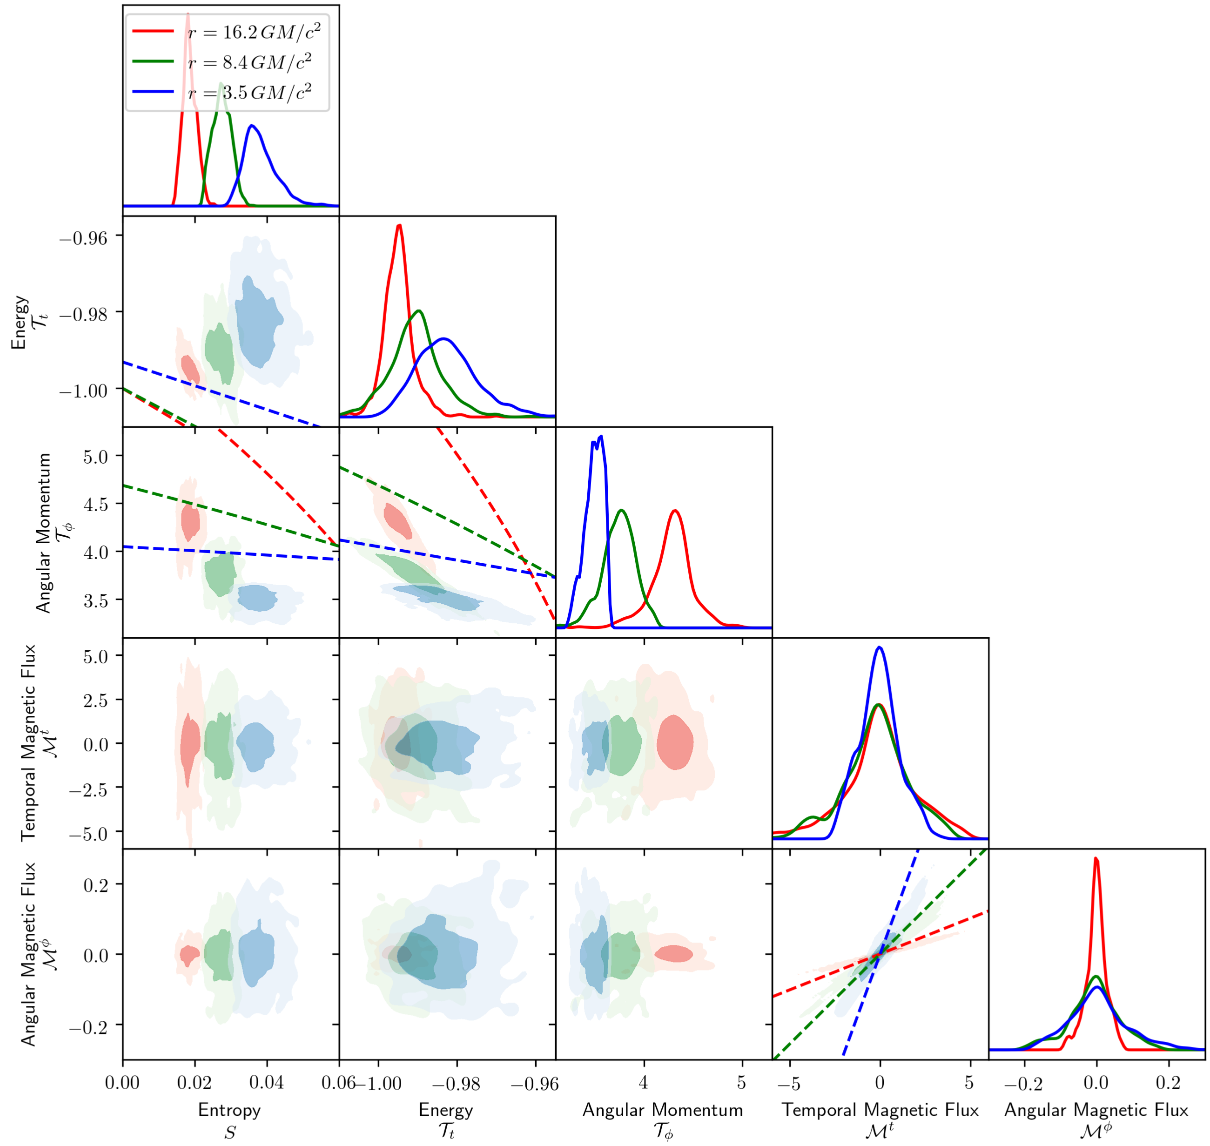

Figure 2.7: Corner plot of the auxiliary variables for the SANE, $a_* = 0$ simulation in the midplane for three radii. Dashed lines show the contour of $u^r = 0$ with only the two corresponding auxiliary variables changing, and the rest set to $S = \mathcal{M}^t = \mathcal{M}^\phi = 0$, $\mathcal{T}_t = -1$, and $\mathcal{T}_\phi = \mathcal{T}_{\phi,\mathrm{crit}}$, appropriately.

Figure 2.6 shows the median and variability in the midplane radial profiles of the density and auxiliary variables for a SANE, $a_* = 0$ simulation. Were these generated from a static simulation, we would expect any static density profile and constant values for the five auxiliary variables. This is not the case; as the gas falls in, $\mathcal{T}_\phi$ varies by a factor of 2, and $\mathcal{T}_t$ and S increase, though the latter two have only a minor effect on the overall flow. Physically, this means that the gas heats up more than through adiabatic compression, and there must is a mechanism that reduces the angular momentum and increase the radial momentum. The magnetic field varies much more than its median value and also does not significantly impact the overall flow. Still, the variability in the simulation is small, such that calculating the primitive variables p, u^α, b^α from the median auxiliary variables matches well with the median primitive variables to within the ability of the simulation to measure these quantities.

The analysis in Section 2.3.3 suggests that if we extend the static model so that only the angular momentum $\mathcal{T}_\phi$ varies in time, the overall fluid solution would be constant auxiliary variables and a piecewise critical solution to $\mathcal{T}_\phi$. The orange dashed line shows this critical solution for $S = 0.008$. The curving down at large radii is due to the specifics of the density profile and is likely caused by the simulation not equilibrating to a steady state at large radii. Although this solution grossly matches the primary observed behavior, it fails to quantitatively match $\mathcal{T}_\phi$, and does not explain the change in S or $\mathcal{T}_t$.

Figure 2.7 shows a triangle plot of the distribution (in time) of the auxiliary variables for the SANE, $a_* = 0$ simulation. These distributions were constructed at a chosen midplane value of (r, θ, ϕ) and sampled in the range $t = [25,000, 30,000]\, GM/c^3$. The static GRMHD solutions we derived had these five auxiliary variables constant in radius, and follow critical values ($u^r = 0$) where the solution no longer exists. To that end, in some of the panels, we show the critical values that lead to $u^r = 0$, varying only the two associated auxiliary values. The rest are set to the default outer solution, namely zero magnetic field and pressure, and energy $\mathcal{T}_t = -1$ such that $u^r \to 0$ as $r \to \infty$. For the default $\mathcal{T}_{\phi,\mathrm{crit}}$, we take the piecewise solution through Equation 2.27,

$$\mathcal{T}_{\phi,\mathrm{crit}} = \begin{cases} \frac{\sqrt{2}\sqrt{r}\sqrt{a^2+(r-2)r}-2a}{r-2}, & r \geq -a - 2\sqrt{1-a} + 2 \\ 2\left(1 - \sqrt{1-a}\right), & r < -a - 2\sqrt{1-a} + 2 \end{cases}. \tag{2.41}$$

If we approximate the flow with auxiliary variables that don't interact (as in Section 2.3.3), the distributions should be truncated as the gas falls and more of the distribution creates zero radial velocity. This is not the case, but the correlations between auxiliary variables and their change with radius follow similar behavior to the critical values. This is most evident where changes in $\mathcal{T}_\phi$ are concerned, as these are responsible for the dominating change in the structure of the flow. The other notable correlations are in the magnetic fields. If we take $\mathcal{T}_\phi = \mathcal{T}_{\phi,\text{crit}}$, $\mathcal{T}_t = -1$, and $S = 0$, then any value along

$$\mathcal{M}^\phi = \frac{2a + (r-2)\mathcal{T}_{\phi,\text{crit}}}{a^2(r+2) - 2a\mathcal{T}_{\phi,\text{crit}} + r^3}\mathcal{M}^t \tag{2.42}$$

will still yield $u^r = 0$ for $r \geq r_{\text{crit}}$. This correlation corresponds to where $b^2 = 0$. Since this roughly matches the correlation seen in the simulation, it implies that the magnetic fields are varying in such a way that their effect on the momentum of the flow (in particular, through u^r) is negligible.

Motivated by the dependence of the flow structure on the critical solutions with $u^r = 0$, we propose that the variability can be split into two important effects. The first is a global effect which sets the evolution of the distributions through allowed regions in the space of auxiliary variables that have $u^{r2} > 0$. As fluctuations push close to $u^r = 0$, the gas tries to stay still, which is unstable. The second effect is local and somehow mixes the distributions of nearby elements through the wave interactions we have so far ignored. As these waves propagate, they move information through the flow until some stable steady state is reached. We leave further understanding of this to future work and approximate its effect in the next section.

GRMHD Approximation

We introduce the approximation of the median GRMHD SANE radial profiles as a sum of two critical solutions. The motivating idea is that an outer boundary condition feeds in gas which pushes the solution towards one critical solution. This has:

$$\begin{aligned} S &= S_{\text{out}} \approx 0, \\ \mathcal{T}_t &= \mathcal{T}_{t,\text{out}} \approx -1, \\ \mathcal{T}_\phi &= \mathcal{T}_{\phi,\text{crit}}, \\ \mathcal{M}^t &= \mathcal{M}^t_{\text{out}} \approx 0, \\ \mathcal{M}^\phi &= \mathcal{M}^\phi_{\text{out}} \approx 0. \end{aligned}$$

Here, $\mathcal{T}_{\phi,\text{crit}}$ behaves as described in Section 2.3.3 and approximately equals Equation 2.41. We need not set its outer value, since the angular momentum can keep increasing well past the steady-state regime of the simulations. We also introduce a second solution sourced

Spin	S_{in}	S_{out}	$\mathcal{T}_{t,\text{in}}$	$\mathcal{T}_{t,\text{out}}$	$\mathcal{T}_{\phi,\text{in}}$	$\mathcal{T}_{\phi,\text{out}}$	$\mathcal{M}^t_{\text{in}}$	$\mathcal{M}^\phi_{\text{out}}$	$\mathcal{M}^\phi_{\text{in}}$	$\mathcal{M}^\phi_{\text{out}}$
-0.94	0.045	0.008	-0.973	-0.998	3.7	>7	-0.02	0	-0.001	0
-0.5	0.055	0.008	-0.97	-0.998	3.4	>7	-0.05	0	-0.002	0
0	0.065	0.008	-0.97	-0.998	3	>7	-0.15	0	-0.005	0
0.5	0.095	0.008	-0.961	-0.998	2.7	>7	0.04	0	0.001	0
0.94	0.2	0.008	-0.93	-0.998	1.8	>7	-0.03	0	0.001	0

Table 2.3: SANE simulation fitting values.

by some saturated state near the black hole, which attempts to push the solution towards a different critical solution:

$$\begin{aligned} S &= S_{\text{crit}}, \\ \mathcal{T}_t &= \mathcal{T}_{t,\text{crit}}, \\ \mathcal{T}_\phi &= \mathcal{T}_{\phi,\text{in}}, \\ \mathcal{M}^t &= \mathcal{M}^t_{\text{in}}, \\ \mathcal{M}^\phi &= \mathcal{M}^\phi_{\text{in}}. \end{aligned}$$

Here, the critical values start at some constant $S_{\text{in}}, \mathcal{T}_{t,\text{in}}$ near the black hole and stay constant as the radius increases. When $u^r = 0$, we have a critical surface to move on (in $r, \vec{A}$), and only one solution to choose from. We arbitrarily choose it such that S_{crit} varies linearly from S_{in} to 0 as $\mathcal{T}_{t,\text{crit}}$ varies linearly from $\mathcal{T}_{t,\text{in}}$ to -1.

These two solutions are both shown in Figure 2.6 with the outer solution in orange (dashed) and the inner solution in magenta (dotted). To get the final combined solution, we take the average of the two,

$$\vec{A} = \frac{\vec{A}_{\text{out}} + \vec{A}_{\text{in}}}{2}. \tag{2.43}$$

Comparison with All SANE Simulations

Figure 2.8 compares our model (dashed) to the GRMHD simulations (solid). It agrees well at all radii. In particular, there are breaks in the median radial profiles of the entropy, energy, and angular momentum (S, $\mathcal{T}_t$, and $\mathcal{T}_{\phi,\text{out}}$, respectively) and our midplane approximation gets the correct location. It mostly matches the radial dependence on either side of the break, except for the entropy, where we have limited ourselves to a constant value at small radii. The simulations' entropies do have a break, but they increase as the gas

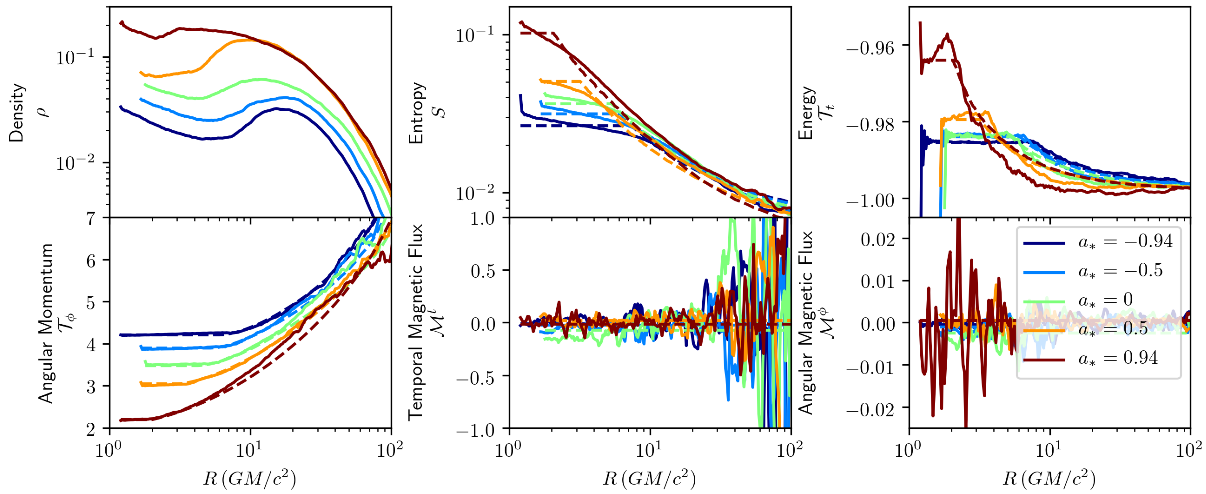

Figure 2.8: Median radial midplane profiles for the GRMHD SANE simulations (solid) and our GRMHD approximation (dashed). Our model quantitatively matches the locations of the breaks in the entropy, energy, and angular momentum, as well as their behaviors on either side. The median magnetic fields in the simulations are very noisy and consistent with zero.

falls into the black hole. The magnetic fields in the simulations are small and variable, so measuring a typical value from the simulations is difficult. Similarly, the radial profiles in our model are minimally affected by the particular value of the magnetic field parameters.

The values used to generate the model curves are listed in Table 2.3. There are nominally 10 parameters. All those associated with the magnetic field can be consistent with zero. The value for $\mathcal{T}_\phi$ need not be set, so we extend the angular momentum to the full range of the solution. The outer values for the entropy and energy have been set to match the simulations better, but we can still get good fits with $S = 0$ and $\mathcal{T}_t = -1$. Thus, we have a 5-parameter fit (plus the given simulation density profile) that reproduces the three nonzero GRMHD simulation auxiliary variables and has the correct spin dependence.

2.4.2 MAD Models

Zero Spin

Figure 2.9 and Figure 2.10 show the profiles and the corner plot for the MAD, $a_* = 0$ simulation. Similar to the SANE model, the main effect is a decrease in angular momentum

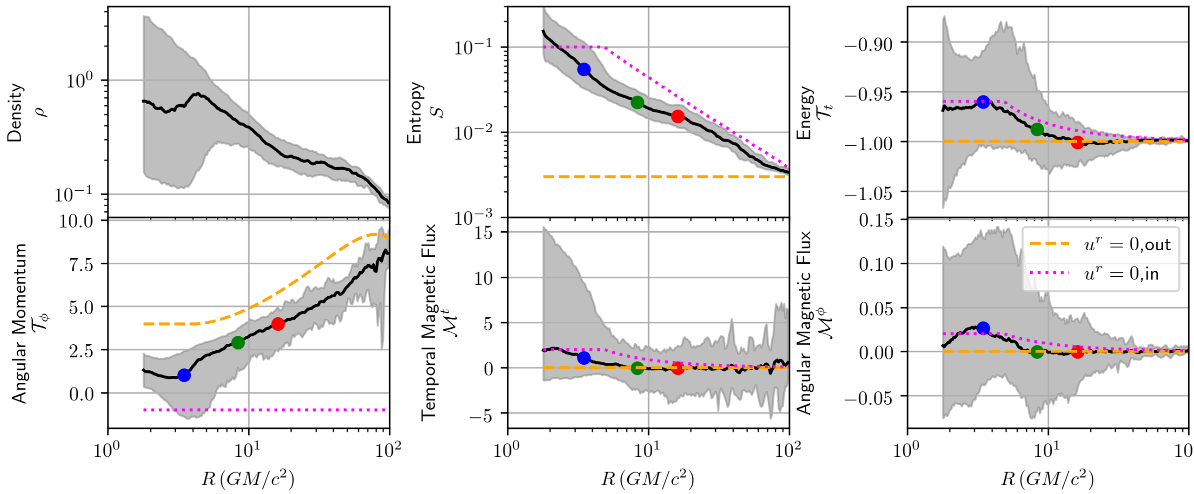

Figure 2.9: Same as Figure 2.6, but for the MAD, $a_* = 0$ simulation.

as gas falls into the black hole, as well as slight increases in the entropy and energy. The auxiliary variables are still influenced by the curves which have $u^r = 0$.

The key difference with SANE models is the higher degree of variability. Near the black hole, there exist large fluctuations with small $\mathcal{T}_\phi$, small density, large entropy, and large $\mathcal{T}_t$. These also have a greater magnitude of the magnetic field, and this behavior is seen in all MAD simulations. This large variability presents a problem in describing static profiles, as the primitive variables constructed from the median auxiliary variables do not correspond well to the median primitive variables. This is primarily due to the large variance in $\mathcal{T}_\phi$ which makes $\langle u^\alpha \rangle \langle u_\alpha \rangle \neq -1$ and instead contain significant radial structure. Unlike in the SANE simulations, to properly characterize the flow, we need the co-variances of the auxiliary variables along with the median values.

All MAD Simulations

Figure 2.11 shows the median midplane profiles of the auxiliary variables for the five MAD simulations. Also shown in bands is the magnitude of the variability, which is particularly important for the energy and angular momentum. The spin dependence in the MAD simulations roughly matches the SANEs, as well as the decrease in angular momentum and increase in entropy and energy as the gas infalls. Though less visible, there is evidence for similar breaks in the auxiliary variables in the two accretion flow types.

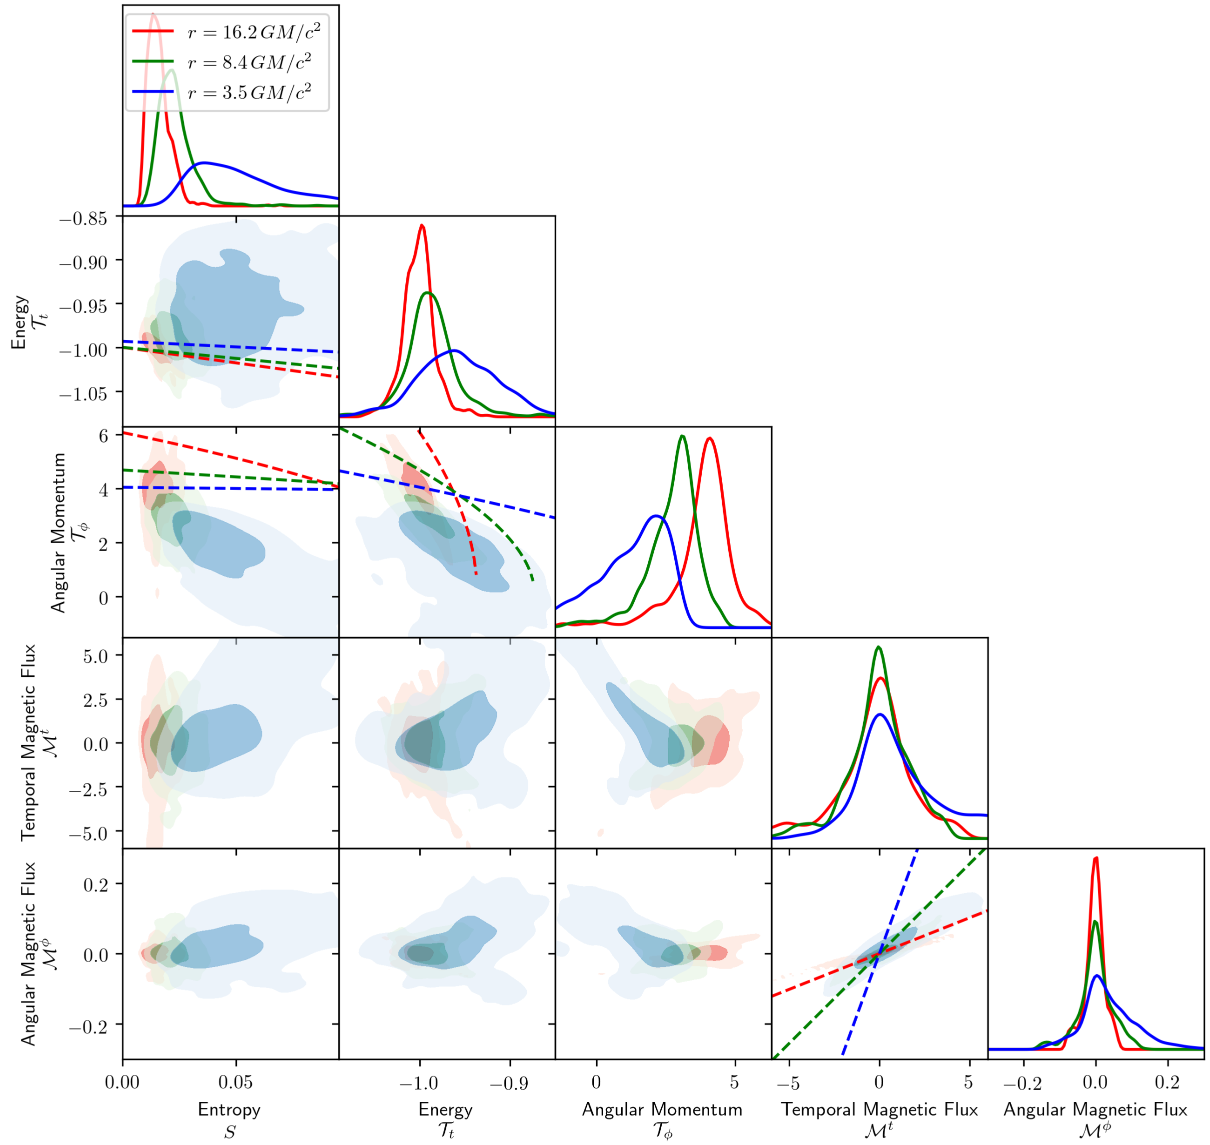

Figure 2.10: Same as Figure 2.9, but for the MAD, $a_* = 0$ simulation.

The main difference is in the angular momentum, which does not settle to a constant value near the black hole but rather becomes more variable. Similarly, the energy does not become constant and increases by a greater amount. Outside of the break in the entropy, MAD and SANE profiles follow roughly the same radial power law; however,

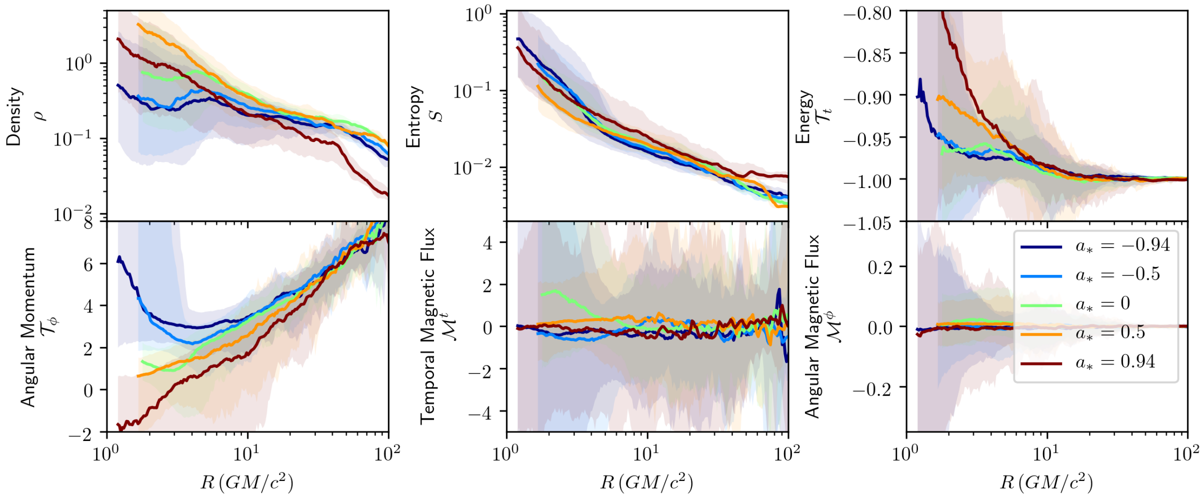

Figure 2.11: Median radial midplane profiles for the GRMHD MAD simulations (solid lines) and the inner 68% region. The MAD simulations are more variable than the SANEs, particularly for the energy and angular momentum.

while the SANE entropy flattens near the black hole, the MAD entropy becomes steeper. Surprisingly, the magnetic field appears to have the same level of variability (at least, far from the black hole), despite the more ordered magnetic structure in MAD models. However, they differ significantly, as the average magnitude of the magnetic field in MAD simulations is larger and strongly correlated with other auxiliary variables, especially near the black hole. All of these effects combined imply the existence of a significant turbulent effect near the black hole that heats up and massively accelerates the gas.

The approximation that worked for SANEs fails for the MAD models. The radial dependence of the most dynamically-relevant quantity, the angular momentum $\mathcal{T}_\phi$, does not behave like $\mathcal{T}_{\phi,\mathrm{crit}}/2$ at large radii, but more like $\mathcal{T}_{\phi,\mathrm{crit}}$. If there is some sourced turbulence near the black hole like in the SANE models, its effects do not reach in the same way to large radii. Furthermore, the angular momentum switches sign near the black hole, thus approaching another critical surface at approximately $-\mathcal{T}_{\phi,\mathrm{crit}}$ (though this value varies slightly with spin). Many of the other auxiliary variables approach multiple surfaces with $u^r = 0$ due to the higher variability, especially when combined with the larger variability in the density.

2.4.3 Successes and Failures of the Static Model

We now have a physically-motivated model that, once calibrated to the behavior of SANE simulations, accurately matches the behavior for different spins. At a minimum, it is an interpolation machine, one that can extrapolate to some non-Kerr metrics and to large radii. Moreover, it implies a specific non-linear way in which the saturated turbulence affects the median midplane profiles, and can be extended to the flow away from the midplane.

Outward Information Travel

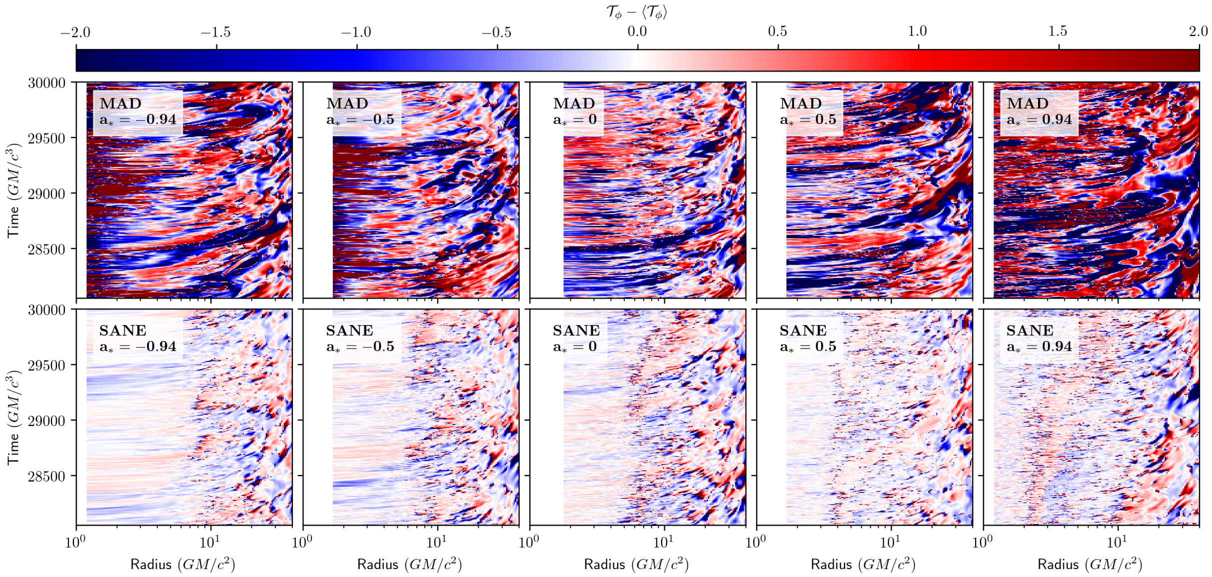

Figure 2.12: Median-subtracted angular momentum in the simulation midplanes. The variability in the MAD models is larger, and all simulations show structures that appear to be moving outwards at all radii.

Figure 2.12 shows the median-subtracted angular momentum at a slice in the midplane and some constant azimuthal angle for all the simulations. As also found in Section 2.4.2, the variability in the MAD models is larger than in the SANEs. Moreover, the variability is not random; there are clear structures that appear to be moving outward. They typically move faster than the sonic, Alfvén, or radial speeds, and more closely match the azimuthal velocity. This does not necessarily mean that anything in the system is moving outward.

There could be, for example, spiral features in the flow that when rotated past a fixed azimuth, appear as a radially outgoing gas. Nevertheless, there is constantly information traveling outward from the event horizon.

At the beginning of the simulation, the gas starts far from the black hole and infalls, and thus information travels inward. In fact, the median radial velocity is typically negative everywhere the simulation has reached a steady state. Our model assumed that both the gas near the horizon and at some outer radius were each trying to bring the flow to a particular solution (constant auxiliary variable or zero radial velocity). The simulations have support for this assumption.

Radial Characteristics

In Section 2.3.3, where the auxiliary variables did not interact, each boundary condition was associated with a connected region within which ${u^r}^2 > 0$. Figure 2.13 assumes a simulation's distribution of auxiliary variables at some radius and tracks the maximum radius in either direction that solution could exist for.

Near the black hole, all the gas within about r_{crit} is infalling and can make it to the horizon. However, none of the gas outside of about $r \gtrsim 10\,GM/c^3$ can. Further from the black hole, the range of radii that one location influences is somewhat small. Thus, there must be some local effect that evolves the distribution, and it must be present at all radii, perhaps with the exception of those very close to the horizon.

In a flat Cartesian spacetime, we would still expect this effect to exist. Mathematically, it takes the form of the shocks and rarefaction waves we ignored. Physically, it corresponds to steady-state turbulence. There exist many prescriptions for turbulence in many regimes, and there is evidence in the GRMHD simulations that these prescriptions could be applied on top of a global background critical solution with $u^r = 0$ that incorporates the large-scale gravitational and rotational effects.

Angular Momentum Variability in MAD models

Figure 2.14 shows $\mathcal{T}_\phi$ for two MAD models as well as the two critical solutions with $u^r = 0$. These are given by the first line of Equation 2.41. In the SANE models, we used the idea that the gas near the horizon created a particular fluctuation that then pushed the entire gas toward an inner critical solution. For the MADs, a similar idea holds, but at small radii, there is a larger distribution of possible auxiliary variable values.

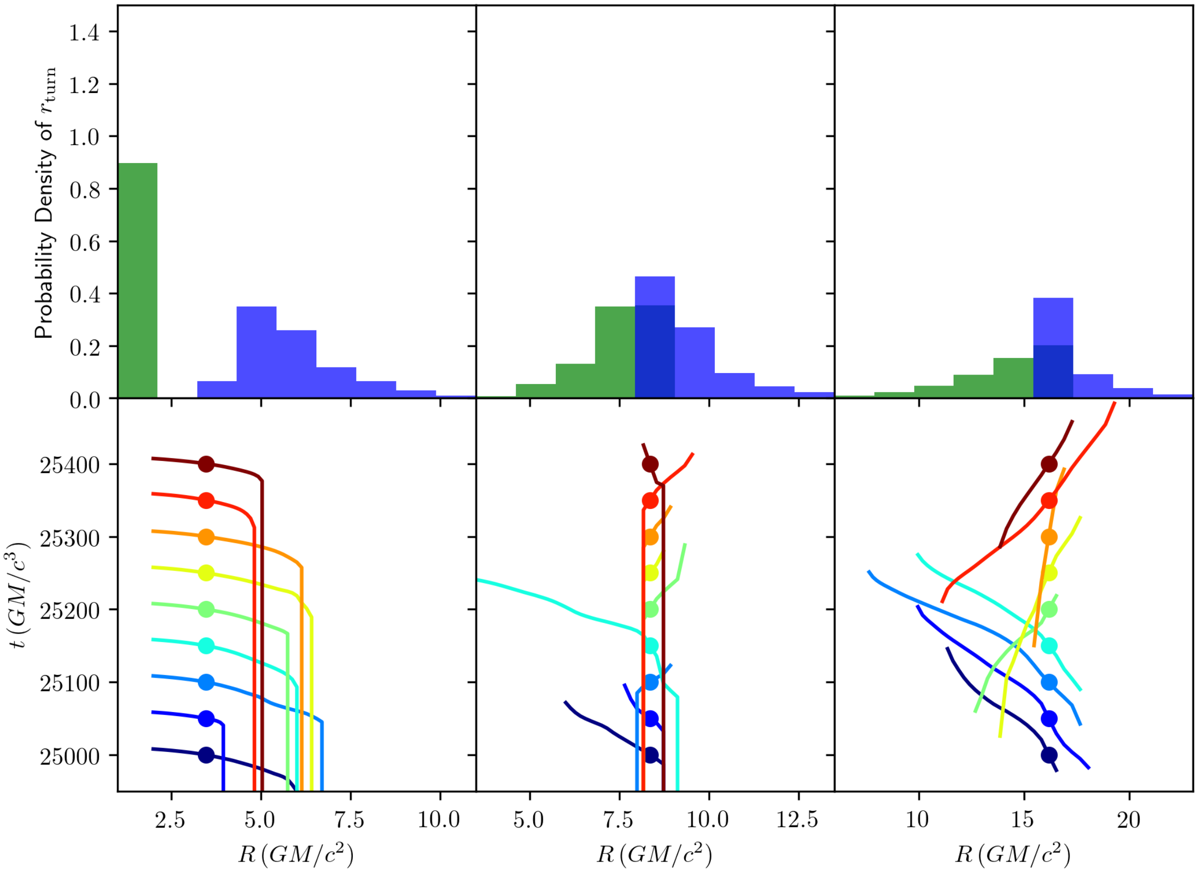

Figure 2.13: Probability density of the truncating radius where a static solution exists for a SANE, $a_* = 0$ simulation. Bottom panels show example characteristics moving at the u^r/u^t for the instantaneous auxiliary variables and density. The filled circles represent the starting radii and times. Top panels show the maximum (blue) and minimum (green) distribution of the truncating radius taken over all values from $25,000$ to $30,000\,GM/c^3$. The three columns correspond to three different starting locations.

Near the black hole, the two critical $\mathcal{T}_\phi$ seem to partially bound the variability in the GRMHD simulations. This could either be gas exploring every possible angular momentum before it falls in or fluctuations created near the horizon truncating when they cannot travel further outward. The variability present in MADs but not in SANEs is not random but rather fills out a region in the space of auxiliary variables between critical surfaces.

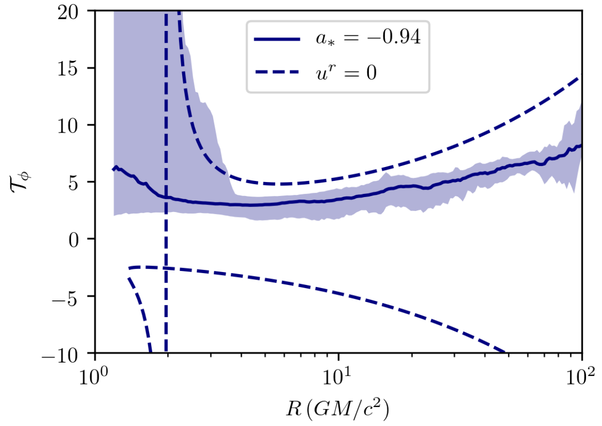

Figure 2.14: Comparison of the variability in the angular momentum of two MAD GRMHD simulations with the critical value that causes $u^r = 0$. The solid lines and shaded regions match Figure 2.11. For the critical curves, $\mathcal{T}_t = -1$, and $S = \mathcal{M}^t = \mathcal{M}^\phi = 0$. The variability in the simulation profiles follows the critical solution at small radii.

Density and the Vertical Structure

In constructing the midplane model, we essentially assumed a boundary condition on the vertical derivative of all quantities. This was not a complete set, and we chose to set the density to the GRMHD profile. By doing so, we conflated variability in the density with variability in the auxiliary variables. Especially for the MAD models, this remains a problem. We can solve this by solving the static GRMHD equations as a function of r, θ with proper outside boundary conditions. We leave this extension for future work. However, it is interesting to note that each boundary condition will correspond to a region of an allowed solution. The method by which the static model approximated the SANE simulations implies that a similar average of an inner and outer critical solution could well approximate the vertical structure of SANE models.

Furthermore, the density profiles in most simulations converge to a curious 3-section function. The density peaks both at the black hole and at some middle radius. We connect this effect to the ideas for the formation of critical solutions in Section 2.3.3. There, the mass would increase with radius, since the rarefaction waves stretched out the gas as it fell in. Similarly, if we interpret the information traveling outward as also rarefying before converging to a critical solution, then we would expect a decrease in density with radius. These effects would meet near r_{crit}, where the angular momentum was low enough so all the gas could fall in. Indeed, for the SANE simulations, the minimum occurs near this value with the appropriate spin dependence.

Arbitrary Choices in the Approximation

We made two significant arbitrary choices in the approximation to SANE models. For the inner critical solution, we had two auxiliary variables to vary (S and $\mathcal{T}_t$) along the critical surface such that they lead to their default values at infinite radius. Without an understanding of the mechanism by which the entropy and energy evolve to a critical solution, we cannot repeat the analysis of Section 2.3.3. However, if we had, we could then write a relationship between S_{crit} and $\mathcal{T}_{t,\text{crit}}$ independent of radius. By choosing them to be linearly related, we have essentially assumed that S_{in} is close to zero and $\mathcal{T}_{t,\text{in}}$ is close to -1, and linearly expanded. From the variability observed in the SANE simulation, this approximation holds well.

The second choice is to take an equally-weighed average of the inner and outer critical solutions. For every radius, this corresponds to a point in the space of auxiliary variables between two critical surfaces, and there does not seem to be anything special about this location (e.g., u^r does not reach a maximum value there). We choose an average only because it approximates the simulations well and a first-principles approach to the evolution of the distribution of auxiliary values should settle on a better choice.

2.5 Conclusions

In this chapter, we studied the variability in GRMHD simulations and constructed a model to approximate the median midplane profiles. For this, we used 10 simulations spanning the two accretion flow types, MAD and SANE, and 5 spin values.

Since the GRMHD equations can be decomposed into an average and a variability piece, we first directly computed the contributions of the variability on the average solution. For SANE simulations, we find that the impact of the covariance of variability fluctuations is small in the midplane and disk, only becoming significant in the corona and jet. MAD simulations, on the other hand, are more variable, though the midplane remains the location where variability effects on the GRMHD equations are minimized. However, the effects on the primitive variables of interest are not as negligible, and we further test this by constructing a midplane model.

In the midplane, for a vertically symmetric static axisymmetric solution, the GRMHD equations can be integrated analytically for a given density profile. The radial constants, which we dub the auxiliary variables, correspond to the specific entropy, energy, angular momentum, and temporal and azimuthal magnetic flux. However, a midplane solution does not necessarily cover the entire accreting region, and we highlight three cases relevant in the simulations. When the angular momentum is large, the incoming gas cannot reach the black hole. When the energy or entropy is too large, gas near the black hole cannot have originated from infinity. That is, the gas would need to be heated more than gravity and adiabatic compression would dictate.

We then look at the temporal evolution of the model when only the angular momentum can vary at a time. The solution converges to a piecewise solution, where the gas has a constant angular momentum wherever the solution exists, and follows a critical surface with zero radial velocity otherwise. If the incoming gas is itself variable, the solution still follows either a constant or a critical solution, but now depends on the distribution of variability being input. This motivates us to introduce a variability model where the midplane solution is some combination of static auxiliary variables and critical solutions.

We first test this model with the GRMHD SANE simulations. Due to their low variability, the median auxiliary variables and density provide a nearly complete description of the midplane flow. Furthermore, the radial dependence of the full distribution of auxiliary variables is heavily influenced by the critical surfaces. Upon applying our static model, we find that we can approximate the SANE profiles as averages of two critical solutions. The first is an outer solution that has a high angular momentum, and nearly zero entropy, magnetic fields, and velocity, and is characterized by a decrease in angular momentum as the gas falls in, until it settles at some metric-dependent value. The second is an inner solution, and has some saturated values for all the auxiliary variables, while the entropy and energy get smaller as the radius increases. This approximation required several arbitrary assumptions, so we make no physical interpretation of the saturated auxiliary values.

The approximation works well for the five SANE simulations, and can reproduce the complex radial structure better than the classical ADAF model. Although it nominally has 10 parameters, about half of those can be set by reasonable constraints on the type of accretion flow. The auxiliary variable profiles are characterized by a number of breaks and multiple types of radial dependencies. Our model consistently gets the location of the breaks, and after setting the magnitude of the fluid variables, our model matches the simulations' radial dependencies. Since our model assumes only a vertically-symmetric, axisymmetric, static metric, it is poised to make strong predictions for GRMHD SANE profiles in non-Kerr spacetimes.

Since the MAD models are more variable, it is not straightforward to characterize their profiles with just a median value, and they require at least some measure of the magnitude of the variability. Nevertheless, like the SANE models, the distribution of auxiliary variables is clearly influenced by the critical solutions. Unfortunately, the variability is too great, and we cannot decompose the profiles into an average of two solutions. To better approximate the MAD models, we need a self-consistent way of tracking the distribution of the variables, and not just their medians.

To fully be able to reduce variable simulations to static ones, future work is needed on the evolution of the distribution of variability as bounded by the critical solutions. Our model can be extended to two dimensions, and there is evidence that the variability distributions will become even more important in the corona region.

Chapter 3

A Universal Power Law Prescription for Variability from Synthetic Images of Black Hole Accretion Flows

3.1 Introduction

The event-horizon sources the EHT observes are expected to be variable on timescales of their light-crossing time, which is about a minute for Sgr A*. This variability is a consequence of the turbulent magnetohydrodynamic features present on all spatial scales and is consistently observed in simulated images of black holes. Thus, it is now relevant for EHT observations to develop the theoretical expectations for the short-timescale ($\lesssim GM/c^3$) and small-spatial scale ($\lesssim GM/c^2$) aspects of these images.

For Sgr A* in particular, the light curve is significantly variable and has been well characterized at near-infrared and radio wavelengths (Wielgus et al. 2022, and references therein; Witzel et al. 2012; Gravity Collaboration et al. 2020; Witzel et al. 2021; Paper II). Below a break that lies between 2 and 8 hrs, the light curve variability is consistent with a red noise process, having a power-law spectrum that is dominated by the longest timescales. At 1.3 mm wavelengths, the typical degree of variability is observed to be less than 10% during the 2017 EHT observing campaign (Wielgus et al., 2022). The EHT observations of Sgr A* imply a similar degree of variability in the source structure (Paper IV).

The variability properties of Sgr A* are grossly consistent with the expectations from GRMHD simulations (Paper V). These simulations exhibit coherent and stochastic structural fluctuations on timescales spanning all of those accessible to the simulation domain; for Sgr A*, this ranges from seconds to many days.

The presence of substantial horizon-scale variability in EHT targets has dramatic implications for the analysis of EHT VLBI data. Sgr A* and, to a lesser degree, M87* necessarily violate the stationarity assumptions underlying Earth-aperture synthesis. In the case of the former, this presents a significant impediment to imaging and model comparison (Broderick et al. 2022b; Paper III; Paper IV). At the same time, direct measurements of the structural variability – i.e., variability caused by the evolving source structure – provide a novel constraint on models of the underlying astrophysical processes responsible for accretion (Wielgus et al., 2020).

In this chapter, we explore the properties of the structural variability observed within simulated images from the large library of GRMHD simulations produced to interpret Sgr A* (Paper V). These simulations span a wide range of black hole and accretion flow properties: including spin, inclination, flow type, and microphysical electron heating prescription. We do this primarily via the construction of statistical measures of the structural variability within the images. We find a surprising degree of uniformity, despite the extreme variation among image morphologies.

This work in this chapter is also presented in Georgiev et al. (2022). We begin in Section 3.2 with a description of the GRMHD simulations. Definitions of the statistical measures employed can be found in Section 3.3. Application to the GRMHD library from Paper V can be found in Section 3.4 and the dependence on the accretion flow properties is discussed in Section 3.5. Finally, conclusions may be found in Section 3.6.

In order to span a large parameter space, the GRMHD simulations have made significant assumptions, some of which we can test with targeted simulations. We study longer timescale simulations and the effects of secular evolution in Subsection B.1.3. In Subsection B.1.2, we study the effects of grid resolution, and we briefly discuss differences between fastlight and the "slowlight" alternative in Subsection B.1.1. We compare the results from two separate fluid simulation codes in Subsection B.1.4. We show applications of our results to EHT data as shown in Paper IV and Paper V in Section B.2. In addition, we comment on the applicability of the analysis in light of the current inability of the EHT to measure the absolute position of sources in Section B.3. Section B.4 contains a collection of average images for many of the simulations used.

3.2 Horizon-scale Observations and Simulations of Sgr A*

The immediate application of this work is to observations of the supermassive black hole at the galactic center, Sgr A*. In this section, we describe the current leading-edge observations of Sgr A* using the EHT and identify the necessary outputs any physical model must contain. Then, we describe a library of GRMHD simulations, from which the synthetic EHT observables are created.

3.2.1 EHT Observables

The EHT is a global VLBI array of radio telescopes. Each pair of concurrently observing telescopes forms a baseline and measures the complex Fourier transform of the on-sky intensity distribution $I(x, y, t)$ – visibility $\mathcal{V}(u, v, t)$. Hence, at time t, the measured visibility is

$$\mathcal{V}(u, v, t) = \iint dxdye^{-2\pi i(ux+vy)} I(x, y, t). \tag{3.1}$$

In 2017, the EHT array consisted of eight stations at six geographical sites (Paper II). However, not all stations can observe Sgr A* simultaneously, and hence the available data correspond to only 13 non-redundant baselines. This results in a sparse sampling in the Fourier domain. The Fourier coordinates (u, v) probed by each baseline are time-dependent themselves, as an effect of the Earth's rotation. In a standard approach, the source is assumed to be static on the timescale of the observations ($\sim$ several hours), thus allowing the use of the Earth's rotation for the aperture synthesis (Thompson et al., 2017). Sgr A* constitutes an unusual case in which this condition is violated, as the angular resolution of the EHT (about 25 μas, Paper III) corresponds to a linear scale of only 2.5 Schwarzschild radii, given Sgr A*'s mass and distance (Do et al., 2019; Gravity Collaboration et al., 2019). This translates to a light crossing time of 100 s, or 5 GM/c^3, and sets the shortest timescale of structural variability to which EHT observations are sensitive. The longest timescales probed during an observing night correspond to several hours, or $\sim 10^3$ GM/c^3. Therefore, the source is expected to vary over the timescale of the observations.

The 2017 EHT VLBI baselines probe spatial scales varying from 0.5 Gλ to 8.7 Gλ, corresponding to angular scales of about $(25 - 400)\,\mu$as (Paper III). Additionally, two of the EHT sites, ALMA and SMA, are themselves connected element interferometric arrays, probing linear scales of $\sim 10^6$ Schwarzschild radii. While the phased-up signal is used for the VLBI EHT observations (Goddi et al. 2019; Paper II), the internal correlations between the individual telescopes can be used to generate the total flux light curves, representing the time-dependent total flux of the compact Sgr A* source (Wielgus et al., 2022), unresolved at spatial scales of ~ 100 kλ. These observations effectively correspond to $\mathcal{V}(0, 0, t)$ for a compact Sgr A* source. We neglect any large-scale structure like the Sgr A* "minispiral", which is overresolved on baselines longer than ~ 100 kλ and filtered out from the observations (any large-scale source structure is naturally missing in the simulated images, which only represent the event horizon scale compact source and are limited by a correspondingly narrow field of view). The variability of the light curve thus corresponds to the scaling of the integrated flux density of the source, rather than to a structural variability related to spatial redistribution of brightness in the resolved image.

In this work, we characterize the spatiotemporal variability of the visibility domain representation of a broad collection of GRMHD simulations. We attempt to characterize the typical deviation of $\mathcal{V}(u, v, t)$ from the average across the whole simulation, to provide a physically motivated prior for the measured source variability probed by the EHT (Broderick et al., 2022b). In order to facilitate the structural variability analysis of the resolved source, we consider the properties of simulated images normalized by the corresponding light curves (instantaneous total flux values).

3.2.2 Modeling the Galactic Center

Following Paper V, we set the mass of the black hole to be $M = 4.14 \times 10^6 M_\odot$ where $M_\odot$ is the solar mass, and the distance to be $D = 8.127\,\mathrm{kpc}$ (Do et al., 2019; Gravity Collaboration et al., 2019). The field of view (FOV) of the images is chosen large enough to contain the majority of the flux (Paper V). The density scale is chosen so that the average simulated observed compact flux at 230 GHz matches the EHT observed value of 2.4 Jy (Wielgus et al., 2022)[1]. For small enough mass accretion rates, the gas is optically-thin and the image intensity is proportional to the density scale. At 2.4 Jy, Sgr A* is marginally optically-thin, so these GRMHD images can loosely approximate other optically-thin black hole systems, such as M87* with an appropriate mass and flux rescaling.

A full list of fluid simulations used in this work can be found in Table 3.1. We use simulations created by three algorithms, HARM (KHARMA and `iharm3d`), `KORAL`, and `BHAC`.

---

[1] Due to approximations made in fitting for the density scale, the total flux in the simulations varies by about 20 %.

Label	Code	Flux	a_*	$\Gamma_{\rm ad}$	$r_{\rm in}$	$r_{\rm max}$	Resolution	Cadence
A	KHARMA	MAD	0, ±0.5, ±0.94	4/3	20	41	288×128×128	5
A	KHARMA	SANE	0, ±0.5, ±0.94	4/3	10	20	288×128×128	5
B	`iharm3d`	MAD	0.94	13/9	20	41	384×192×192	0.5
C	`KORAL`	MAD	0, ±0.3, ±0.5, ±0.7, ±0.9	13/9	20	42-43	288×192×144	10-50
D	`BHAC`	MAD	0, ±0.5, ±0.94	4/3	20	40	384×192×192	10
D	`BHAC`	SANE	0, ±0.5, ±0.94	4/3	6	12	512×192×192	10
E	`BHAC`	MAD	0.94	4/3	20	40	96^3, 128^3, 192^3	10

Table 3.1: GRMHD fluid simulation parameters. The radii are in units of the gravitational radius, GM/c^2, and the cadence in units of GM/c^3. The resolution study in model set E is composed of three simulations where, e.g., $96^3 = 96 \times 96 \times 96$.

Label	Code	$R_{\rm high}$	i (deg)	Avg. Flux (Jy)	FOV	Res.	Start	End	Cadence
A	`ipole`	1, 10, 40, 160	10-170	2.08-2.47	200	0.5	15,000	20,000	5
A	`ipole`	1, 10, 40, 160	10-170	2.06-2.48	200	0.5	20,000	25,000	5
A	`ipole`	1, 10, 40, 160	10-170	2.13-2.47	200	0.5	25,000	30,000	5
$\rm B_s$	`ipole`	40	60	3.3	160	1	5,520	9,249.5	0.5
B	`ipole`	40	60	3.3	160	1	5,357	9,357	0.5
C	`ipole`	20	10-170	2.24 - 2.43	200	0.5	10,000	100,000-113,500	10-100
D	`BHOSS`	1, 10, 40, 160	10-90	2.4	200	0.5	20,000	25,000	10
D	`BHOSS`	1, 10, 40, 160	10-90	2.4	200	0.5	25,000	30,000	10
E	`BHOSS`	10	30	1.93-2.00	160	0.5	12,000	15,000	10

Table 3.2: GRMHD image properties. The labels match Table 3.1. The inclinations are spaced in 10-degree increments and all start, end, and cadence times are in units of GM/c^3. $\rm B_s$ is imaged using slowlight. The field of view and resolutions are in μas.

HARM

`iharm3d` is an implementation of the HARM algorithm outlined in Gammie et al. 2003, and KHARMA is a GPU port of `iharm3d` closely following the original codebase and functionality (also see Prather et al. 2021). The GRMHD models produced by KHARMA that are considered for this work include MAD and SANE states, each spanning five spins, $a_* = 0,\ \pm 1/2,\ \pm 15/16$ (hereafter written as 0, ±0.5, and ±0.94). These simulations form model set 'A'. A comprehensive study of these models can be found in Dhruv 2023.

The GRMHD snapshots produced by these codes are imaged using the publicly available GRRT code `ipole` (Mościbrodzka & Gammie, 2018). A comprehensive discussion of the steps and approximations involved in the generation of these synthetic images from HARM output using `ipole` can be found in Wong et al. 2022. The images span nine inclinations from 10 to 170 degrees, and four values of $R_{\rm high}$ from 1 to 160. From each simulation, three windows each of length $5,000\ GM/c^3$ are chosen and treated as separate simulations with different density scales.

We also consider a high cadence MAD simulation ($a_* = 15/16$) carried out with `iharm3d`, for the study of short timescales and the effects of the fastlight approximation (cf. model set 'B').

KORAL

The GRMHD models used from `KORAL` are all in the MAD state with spins of $a_* = 0,\ \pm 0.3,\ \pm 0.5,\ \pm 0.7,\ \pm 0.9$. These simulations were run for $100,000\ GM/c^3$ to study the effects of long timescale trends. They form model set 'C' and have been studied in Narayan et al. 2021.

These simulations were also imaged using `ipole` (Mościbrodzka & Gammie, 2018), using a cadence of $100\ GM/c^3$ between times $10,000\ GM/c^3$ and $\sim 100,000\ GM/c^3$, with some segments of some simulations further imaged at a lower cadence. Using the electron temperature prescription of Mościbrodzka et al. (2016), only $R_{\rm high} = 20$ is sampled. We sample inclinations of $i \in \{10, 30, 50, 70, 90, 110, 130, 150, 170\}$. For each model set, we set the density scale, $\mathcal{M}$, to produce an average total flux of 2.4 Jy across the time window, consistent with the Sgr A* total flux measured by Wielgus et al. (2022) during the EHT 2017 campaign. Since these simulations are run for such a long time, the mass accretion rate drops noticeably as the torus drains. To counteract this, we allow the density scaling $\mathcal{M}$ to evolve slowly with time by fitting for a and b in $\ln \mathcal{M} = a + b(tc^3/GM)$ instead of the constant value of $\mathcal{M}$ used for all other codes in this work. This procedure effectively assumes that the effects of the draining torus are limited only to the total flux light curve; this assumption is somewhat tested in Subsection B.1.3.

BHAC

The BlackHoleAccretionCode (`BHAC`) is a multidimensional GRMHD code to solve the equations of GRMHD in arbitrary spacetimes (Porth et al., 2017; Olivares et al., 2019). GRMHD simulations of magnetized accretion flows onto black holes have been performed in both MAD and SANE states with five different black hole spin, $a_* = 0$, $\pm 1/2$, $\pm 15/16$. Imaging of these GRMHD simulations is performed by the GRRT code (`BHOSS`) (Younsi et al., 2012, 2020). Only inclinations less than 90 degrees are considered as the simulations are approximately equatorial symmetric. These simulations and images correspond to model set 'D'.

We have also performed a resolution test by performing the same fluid simulation with three increasingly lower resolutions. This simulation is in the MAD state with $a_* = 15/16$ and forms model set 'E'.

3.2.3 Salient Features of the GRMHD Simulations

Figure 3.1 shows an example output from the simulating procedure for two simulations of model set A on each row. The left three panels show three consecutive frames and the right panel shows an average over all frames in this simulation. The differences in the two simulations chosen are indicative of the wide range of morphology and variability present among simulations.

Source emission is limited to within several times the gravitational length scale. Images typically show asymmetric ring-like emission, with an asymmetric flux distribution extending to larger scales. Variability in this extended region is mainly composed of coherent moving features. On the short timescales shown, these features move less, i.e., the variability on these timescales is lower.

3.3 Power Spectral Density Characterization of the Variability

For each simulation in the GRMHD library, the evolving emission structure is described by a series of "frames," each of which is uniquely determined at time t by its intensity at

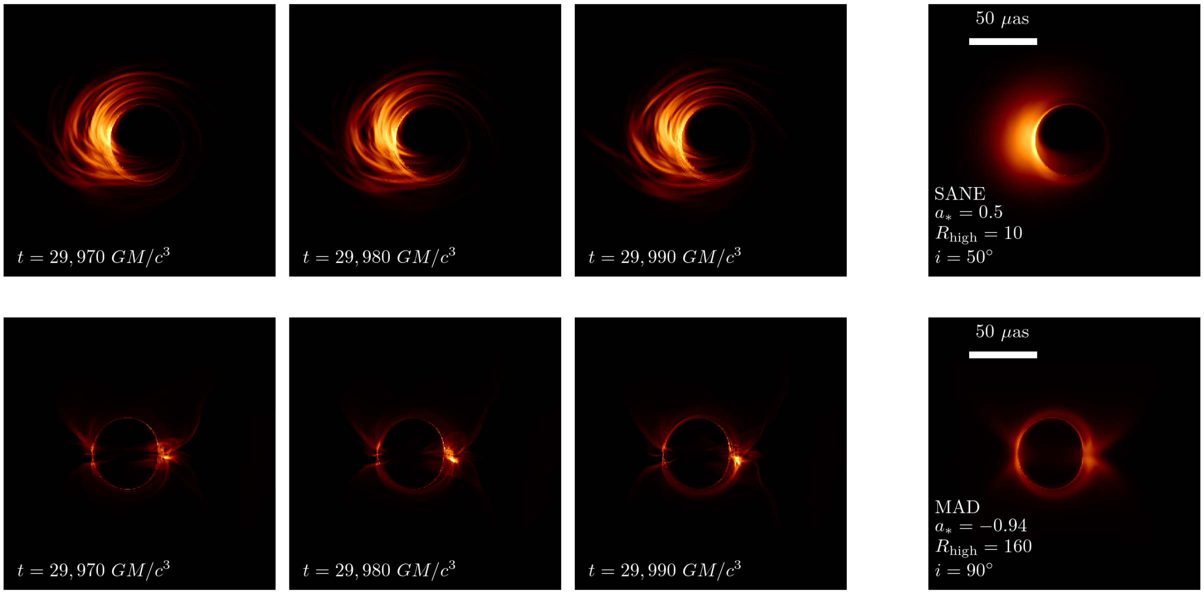

Figure 3.1: Example consecutive frames (left 3 columns) and average image (rightmost column) of a two models from model set A with different structural variability. The top row show a SANE model with $a_* = 0.5$, $R_{\rm high} = 10$, and $i = 50$ degrees and the bottom panel shows a MAD model with $a_* = -0.94$, $R_{\rm high} = 160$, and $i = 90$ degrees. The average is over the third window (25,000 to 30,000 GM/c^3) of each simulation. Images in each row share a common colorscale.

every pixel location (x, y). The full time series of frames from a single simulation then describes a dynamic "image" $I(x, y, t)$ of the evolving emission structure. Our goal is to characterize the variability in the images contained within the GRMHD library. In this paper, we work primarily with both spatial and temporal power spectral density (PSD) representations of the source structure. In this section, we define our PSD construction formalism and describe some general behavior. As we shall see in Section 3.4, we will be dealing with red-noise processes, so we also introduce filtering and averaging procedures to produce accurate statistical representations in the temporal and spatial scales of interest.

Parameter	Units	Name	Description
$I(x,y,t)$	$\mathrm{Jy\,sr^{-1}}$	image	intensity at position (x,y) at time t
$\bar{I}_T(x,y,t)$	$\mathrm{Jy\,sr^{-1}}$	time-smoothed image	intensity at position (x,y) smoothed over a timescale T around time t
$L(t)$	Jy	light curve	spatially-integrated image intensity of $I(x,y,t)$ at time t
$\bar{L}_T(t)$	Jy	time-smoothed light curve	spatially-integrated image intensity of $\bar{I}_T(x,y,t)$ at time t
$\mathcal{I}(x,y,t)$	$\mathrm{sr^{-1}}$	light-curve-normalized image	$I(x,y,t)/L(t)$
$\Phi(u,v,\omega)$	$\mathrm{Jy^2\,Hz^{-2}}$	spatiotemporal PSD	squared amplitude of the three-dimensional Fourier transform of $I(x,y,t)$
$P_T(u,v,t)$	$\mathrm{Jy^2}$	mean-subtracted PSD	PSD of the mean-subtracted image, $I(x,y,t)-\bar{I}_T(x,y,t)$
$\langle P_T\rangle(u,v)$	$\mathrm{Jy^2}$	average mean-subtracted PSD	$P_T(u,v,t)$ averaged in time
$\langle\hat{P}_T\rangle(u,v)$	unitless	average residual PSD	PSD of the residual image, $\mathcal{I}(x,y,t)-\bar{\mathcal{I}}_T(x,y,t)$, averaged in time

Table 3.3: Description of terminology and symbols used in this chapter.

3.3.1 Variability PSD Definitions

The spatiotemporal PSD of a set of images is given by the complex square of the three-dimensional Fourier transform of $I(x,y,t)$,

$$\Phi(u,v,\omega) = \left|\mathcal{F}_{xyt}\left[I(x,y,t)\right]\right|^2 \equiv \mathcal{F}^*_{xyt}\left[I(x,y,t)\right]\mathcal{F}_{xyt}\left[I(x,y,t)\right], \tag{3.2}$$

where

$$\mathcal{F}_{xyt}(f) \equiv \iiint f(x,y,t)e^{-2\pi i(ux+vy+\omega t)}dxdydt, \tag{3.3}$$

is a three-dimensional spatial Fourier transform from (x,y,t) to (u,v,ω), and the asterisk denotes complex conjugation. We will often work with the "spatial PSD" of the image, in which the Fourier transform is only performed on the two-dimensional spatial portion of $I(x,y,t)$,

$$\mathcal{F}_{xy}(f) \equiv \iint f(x,y,t)e^{-2\pi i(ux+vy)}dxdy. \tag{3.4}$$

In this paper, we will often refer to the spatial PSD as simply the PSD, and we will specify otherwise when referring to the full spatiotemporal PSD.

In the absence of variability, the PSD computed at any time t would be identical to the PSD computed at any other time. However, variability in the source introduces variability into the PSD as well, which we can characterize in terms of its variance with respect to an average,

$$\langle P_T(u,v,t)\rangle = \left\langle\left|\mathcal{F}_{xy}\left[I(x,y,t)-\bar{I}_T(x,y,t)\right]\right|^2\right\rangle. \tag{3.5}$$

where the angled bracket notation,

$$\langle f(t)\rangle \equiv \frac{1}{\tau}\int_0^{\tau} f(t)dt, \tag{3.6}$$

represents an average in time over the total duration τ of the simulation. In Equation 3.5, $\bar{I}_T(x,y,t)$ is a time-averaged version of the image, defined using a Gaussian smoothing kernel,

$$\bar{I}_T(x,y,t) = \int I(x,y,t')\frac{1}{\sqrt{2\pi}T}\exp\left[-\frac{(t-t')^2}{2T^2}\right]dt'. \tag{3.7}$$

Here, T is the smoothing timescale and we introduce the bar and subscript notation $\bar{I}_T$ to denote Gaussian smoothing over this timescale. As $T\to\infty$, $\langle P_T\rangle$ measures the variance of the amplitude spectrum at every spatial frequency (u,v). For finite T, contributions to this variance from timescales longer than T are suppressed. For the red-noise processes relevant to this work, $\langle P_T\rangle$ will thus be dominated by the variability on timescale T (for more discussion and application, see Subsection 3.3.3 and Subsection 3.4.2).

We note that $\langle P_T\rangle$ is in general not a complete description of the source variability. Because physical variability originates from continuous processes, we expect the Fourier intensities at (u,v,t) and at $(u+\Delta u, v+\Delta v, t+\Delta t)$ to be strongly correlated for small values of Δu, Δv, and Δt. We also expect the strength of this correlation to decrease as these separations increase. $\langle P_T\rangle$ is just this correlation evaluated at $\Delta u=\Delta v=\Delta t=0$.

In the context of astrophysical observations, variations on the largest spatial scales $(u,v)=(0,0)$ are captured by the light curve,

$$L(t) = \iint I(x,y,t)dxdy \tag{3.8}$$

with a similar definition of L_T containing the total flux of I_T. When $L(t)$ is known, it can be used to remove the variability on the largest spatial scale by "light-curve-normalizing" the image,

$$\mathcal{I}(x,y,t) \equiv \frac{I(x,y,t)}{L(t)}. \tag{3.9}$$

The corresponding light-curve-normalized version of $\langle P_T\rangle$ is given by,

$$\langle \hat{P}_T(u,v,t)\rangle = \left\langle \left|\mathcal{F}_{xy}\left[\mathcal{I}(x,y,t) - \bar{\mathcal{I}}_T(x,y,t)\right]\right|^2\right\rangle, \tag{3.10}$$

with $\bar{\mathcal{I}}_T(x,y,t)$ defined by replacing $I\to\mathcal{I}$ in Equation 3.7. We refer to Equation 3.10 as the "residual" PSD, with an analogous image-domain counterpart $\mathcal{I}(x,y,t)-\bar{\mathcal{I}}_T(x,y,t)$

that we refer to as the residual image. In practice, removing the large-scale variability in this way also removes some fraction of the variability contained on all spatial scales that are correlated with $(u, v) = (0, 0)$. The impact of this light curve normalization procedure is explored in Subsection 3.4.4.

Practically, the angled brackets in Equation 3.6 are an average over all of the frames in a simulation, that is

$$\langle f(t) \rangle \approx \frac{\sum_i f(t_i) \Delta t}{\sum_i \Delta t}, \tag{3.11}$$

where Δt is the frame spacing[2]. The integral in the definition of the time-smoothed image in Equation 3.7 is similarly replaced with a sum,

$$\bar{I}_T(x, y, t_j) \approx \frac{\sum_i I(x, y, t_i) \exp\left[-(t_i - t_j)^2/(2T^2)\right] \Delta t}{\sum_i \exp\left[-t_i^2/(2T^2)\right] \Delta t}, \tag{3.12}$$

where t_i is the frame time. Therefore, outside of the time region of the simulation, the Gaussian weighting kernel is set to zero and when T approaches the length of the simulation, the time-smoothed image approaches the uniformly-weighted mean.

To compute the Fourier transforms in Equation 3.3 and Equation 3.4, we use the fast Fourier transform algorithm implemented in `NumPy` (Harris et al., 2020). When we later use Φ for a qualitative comparison, we average the temporal Fourier Transform over multiple time windows of a simulation to reduce the statistical error. The spatial Fourier transforms in the calculation of $\langle P_T \rangle$ and $\langle \hat{P}_T \rangle$ of a compact source are done on images padded with a large zero array. No significant change occurs when increasing the pixel size or reducing the field of view of the original images. The time-averaged power spectra constructed as described in Equation 3.5 and Equation 3.10 are averaged over many images, and therefore many independent realizations of the underlying random processes of interest, reducing the statistical uncertainty in the variability PSDs.

3.3.2 PSD Behavior on Short Baselines

While the detailed structure of the PSD depends on the specifics of the source variability, we can understand its broader structural characteristics in light of a few relatively simple

---

[2] The `KORAL` runs used for the exploration in Subsection B.1.3 have non-uniform frame spacing. For these, each frame's Δt is found as half the time difference between the previous and successive frames. For the first and last frames, the simulation-averaged frame spacing is used. No temporal Fourier transforms are taken for these simulations.

considerations. The PSD structure on the largest spatial scales can be described by an expansion of the PSD about $(u, v) = (0, 0)$. For any function $f(x, y)$, its spatial Fourier transform expanded around this point is given by

$$\mathcal{F}_{xy}(f) \approx \iint \left[1 - 2\pi i(ux + vy) + \dots\right] f(x, y) dxdy. \tag{3.13}$$

Because $\langle P_T \rangle$ is an even function in (u, v), it must flatten out on large enough spatial scales and will thus be dominated by the first term in the Fourier expansion

$$\langle P_T \rangle \approx \langle P_T \rangle(0, 0) = \left\langle \left[L(t) - \bar{L}_T(t)\right]^2 \right\rangle \tag{3.14}$$

for small enough (u, v), typically those measuring structures larger than the source size. As $T \to \infty$, the value of $\langle P_T \rangle$ on the largest scales approaches the variance in the light curve. For $T \to 0$, the zero-baseline behavior is determined by the temporal power spectral density of the light curve; see Subsection 3.3.3 and Subsection 3.4.3.

However, the first term in the expansion of $\langle \hat{P}_T \rangle$ is zero by construction. Therefore, the leading term in the expansion is

$$\begin{aligned} \langle \hat{P}_T \rangle &\approx 4\pi^2 \left\langle \left[\iint (ux + vy) \left(\mathcal{I} - \bar{\mathcal{I}}_T\right) dxdy\right]^2 \right\rangle & (3.15) \\ &= 4\pi^2 \left\langle \left[u \left(\mathcal{X} - \bar{\mathcal{X}}_T\right) + v \left(\mathcal{Y} - \bar{\mathcal{Y}}_T\right)\right]^2 \right\rangle & (3.16) \end{aligned}$$

where we have rewritten the equation in terms of the centroids of emission as

$$\mathcal{X}(t) = \frac{1}{L(t)} \iint xI(x, y, t) dxdy \tag{3.17}$$

and a similar definition for $\mathcal{Y}(t)$. It follows that on the largest spatial scales, $\langle \hat{P}_T \rangle$ along any direction in (u, v) will behave like a power law in baseline length with an index of 2. It is also interesting to note that for a symmetric (i.e., even) function, all the odd terms of the expansion in Equation 3.13 are zero. Therefore, a large-scale power law index of 2 is an indicator of asymmetric variability in the residual image. More generally, the coefficient in Equation 3.16 is a measure of the amount of variability in the centroid of emission.

The identification of the short-baseline variability described here with an astrometric quantity (the position of the image centroid) raises a practical question. Where an absolute phase reference is available, such a measurement is well defined (see, e.g., Reid & Honma, 2014; Broderick et al., 2011). However, for the current EHT specifically and VLBI studies typically, atmospheric phase delays preclude a full reconstruction of the visibility phases on timescales longer than the atmospheric coherence time; at 230 GHz for the EHT, this is on the timescale of seconds (M87* Paper II). These unknown station-based phase errors are degenerate with an arbitrary shift in the image centroid. We discuss the implications for the short-baseline structure of $\langle \hat{P}_T \rangle$ further in Section B.3. Because our focus in this paper is on the theoretical predictions of variability and the signatures of astrophysical processes contained therein, in the remainder of this paper we presume that absolute phase information is available.

3.3.3 Temporal PSD

The dependence of $\langle P_T \rangle$ on averaging timescale T can be represented by an integral of the temporal PSD. Here, we show how $\langle P_T \rangle$ and $\langle \hat{P}_T \rangle$ contain information about variability on different timescales. Applying Parseval's theorem to Equation 3.5 as $\tau \to \infty$, the integral transforms into the Fourier domain,

$$\lim_{\tau \to \infty} \langle P_T \rangle = \lim_{\tau \to \infty} \frac{1}{\tau} \int d\omega \Phi(u,v,\omega) \left(1 - e^{-2\pi^2 \omega^2 T^2}\right)^2 . \tag{3.18}$$

Subtracting a moving average on different timescales is thus equivalent to applying a high-pass filter to the temporal variability at each spatial scale. Where Φ is a red-noise process[3], this integral is dominated by frequencies near $\omega \sim 1/2\pi T$, so

$$\langle P_T \rangle \approx \Phi\left(u, v, 1/2\pi T\right). \tag{3.19}$$

This equivalence is demonstrated for a GRMHD simulation in Subsection 3.4.2. As $T \to \infty$, $\langle P_T \rangle$ is formally affected by the inability to measure infinite times, but will theoretically converge to the variance map of $\mathcal{F}_{xy}(I)$. On finite timescales, it can be shown that

$$\frac{d \ln \langle P_T \rangle}{d \ln T} = 1 - \frac{1}{\langle P_T \rangle \tau} \int d\omega \frac{d\Phi}{d \ln \omega} \left(1 - e^{-2\pi^2 \omega^2 T^2}\right)^2 . \tag{3.20}$$

For a red-noise process, we can thus directly relate the temporal behavior of $\langle P_T \rangle$ to the spatiotemporal PSD,

$$\left.\frac{d \ln \langle P_T \rangle}{d \ln T}\right|_T \approx 1 - \left.\frac{d \ln \Phi}{d \ln \omega}\right|_{\omega = 1/2\pi T} . \tag{3.21}$$

---

[3] The Gaussian smoothing procedure we employ during averaging effectively softens the high-pass filter, to the point where there exists a maximum power-law index (equal to 4) beyond which the spectrum remains red even after filtering. The GRMHD simulations used in this work typically have power-law spectral indices less than this value in the spatial and temporal scales of interest.

More specifically, if $\Phi \sim \omega^{-\eta}$ near the dominating frequencies at $\omega \sim 1/2\pi T$, then $\langle P_T \rangle \sim T^{\eta-1}$. An equivalent characterization exists for the light-curve-normalized version, with the replacement $P_T \rightarrow \hat{P}_T$ and $I \rightarrow \mathcal{I}$.

3.4 Application to GRMHD

We now apply the formalism presented in Section 3.3 to the GRMHD library, and we characterize the resulting variability PSDs. Specifically, we explore how the variability PSDs depend on the averaging timescale, the light curve normalization, and how the spatial component depends on the orientation. For the representative examples provided in this section, we use a SANE simulation from model set A, with $a_* = 0.5$, $R_{\mathrm{high}} = 40$, and an inclination of 90 degrees[4]. It includes five realizations each of length $5,000\ GM/c^3$ from $5,000\ GM/c^3$-$30,000\ GM/c^3$ (only the last three are listed in Table 3.2)[5]. Where otherwise stated, we instead use simulation set B, which samples frames at a higher cadence.

In Subsection 3.4.1 and Subsection 3.4.2, we discuss the red-noise of the spatiotemporal PSD and describe the equivalence between the mean-subtraction procedure and applying a filter to the spatiotemporal PSD. In Subsection 3.4.3, we show the dependence on the averaging timescale. In Subsection 3.4.4, we demonstrate the effect of normalizing by the light curve. In Subsection 3.4.5, we discuss the short and long baseline limits, and we explore azimuthal dependence in Subsection 3.4.6.

3.4.1 Red-Noise Power Spectrum

Figure 3.2 shows the three-dimensional spatiotemporal PSD of a representative GRMHD simulation, and Figure 3.3 shows two-dimensional and one-dimensional slices of the same PSD. A single point (u, v, ω) in $\Phi(u, v, \omega)$ is a measure of the variability at a temporal frequency ω and at a spatial frequency (u, v). The GRMHD simulations produce characteristically red-noise variability in all dimensions, with the spatiotemporal PSD generically increasing towards small (u, v, ω) and exhibiting its maximum power at $(u, v, \omega) = (0, 0, 0)$.

---

[4] Georgiev et al. (2022) erroneously labeled this simulation as $a_* = 0.94$.

[5] For reference, $5,000\ GM/c^3$ is about 30 hours for Sgr A* and 5 years for M87*.

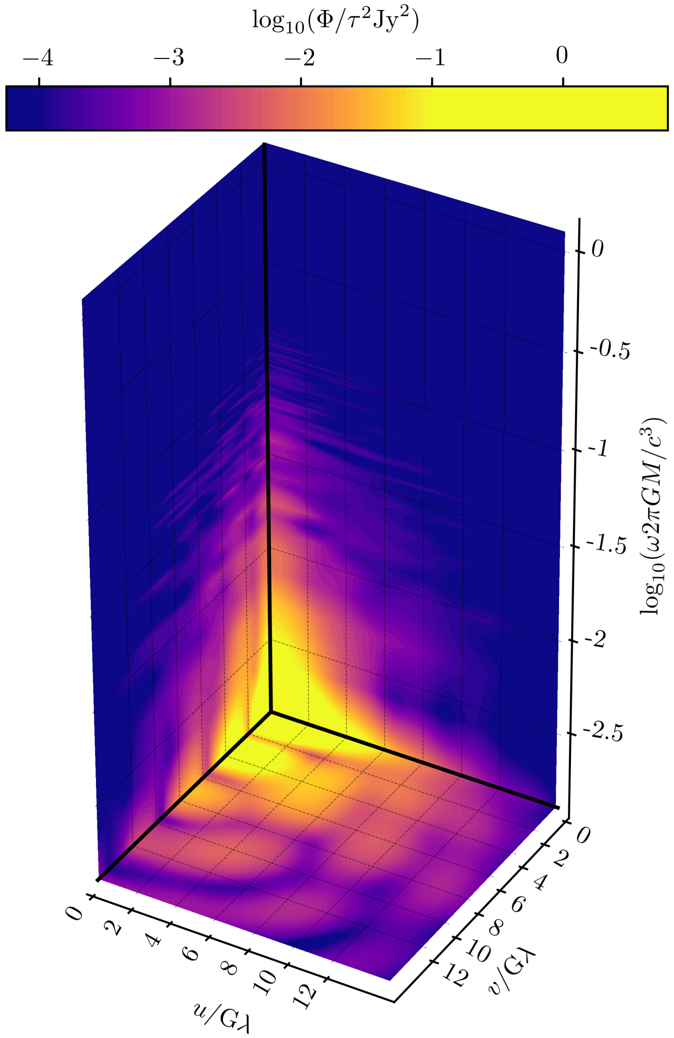

Figure 3.2: A graphical representation of the three-dimensional PSD, $\Phi(u, v, \omega)$, with the different cube faces showing two-dimensional slices along $u = 0$, $v = 0$, and $\omega = 0$. The PSD is a red-noise process in all dimensions, with the power being maximized on both the largest spatial scales and the longest temporal scales. For this example, we use a SANE simulation from model set A with $a_* = 0.5$, $R_{\mathrm{high}} = 40$, an inclination of 90 degrees, and its power spectrum has been averaged over its five time windows.

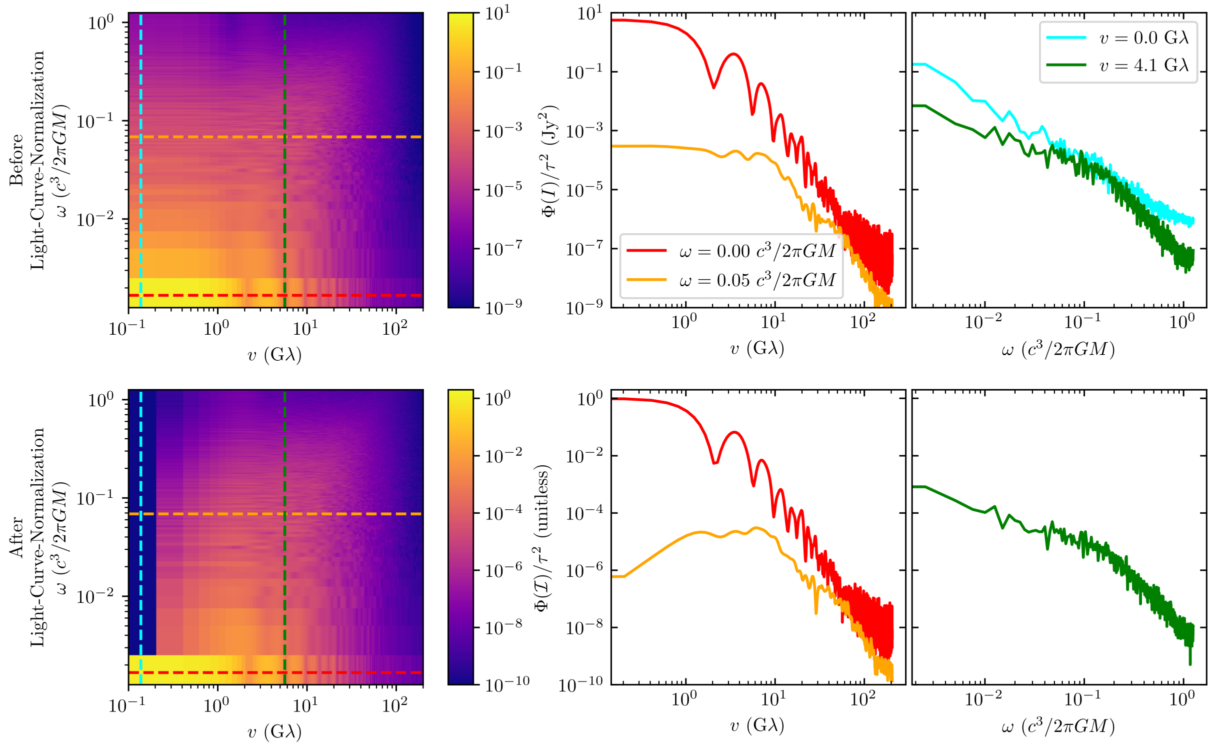

Figure 3.3: *Top:* Two-dimensional and one-dimensional slices of an example spatiotemporal PSD for the same simulation as shown in Figure 3.2 (i.e., a SANE simulation from model set A with $a_* = 0.5$, $R_{\mathrm{high}} = 40$, and $i = 90°$). The left panel shows $\Phi(0, v, \omega)$, corresponding to the left face of the cube in Figure 3.2. Due to the red-noise nature of GRMHD simulations, the majority of the power lives on the largest spatial scales and on the longest temporal scales. The various dashed lines show the locations of one-dimensional slices along the spatial and temporal dimensions, whose corresponding power spectra are shown in the middle and right panels, respectively. The cyan and green curves in the right panel correspond to $\Phi(0, 0, \omega)$ and $\Phi(0, 4.1\ \mathrm{G}\lambda, \omega)$, respectively, and the red and orange curves in the middle panel correspond to $\Phi(0, v, 0)$ and $\Phi(0, v, 0.05\ c^3/2\pi GM)$, respectively. $\Phi(0, v, 0)$ corresponds to the visibilities of the average image, which contains a prominent ring of emission that produces the ringing signature evident in the red curve. *Bottom:* Same as the top row, but showing the spatiotemporal PSD for the light-curve-normalized images. The light-curve-normalization procedure removes all power from $(u, v) = (0, 0)$, and some non-unit fraction of the power at non-zero spatial frequencies (i.e., those small enough to be strongly correlated with the light curve).

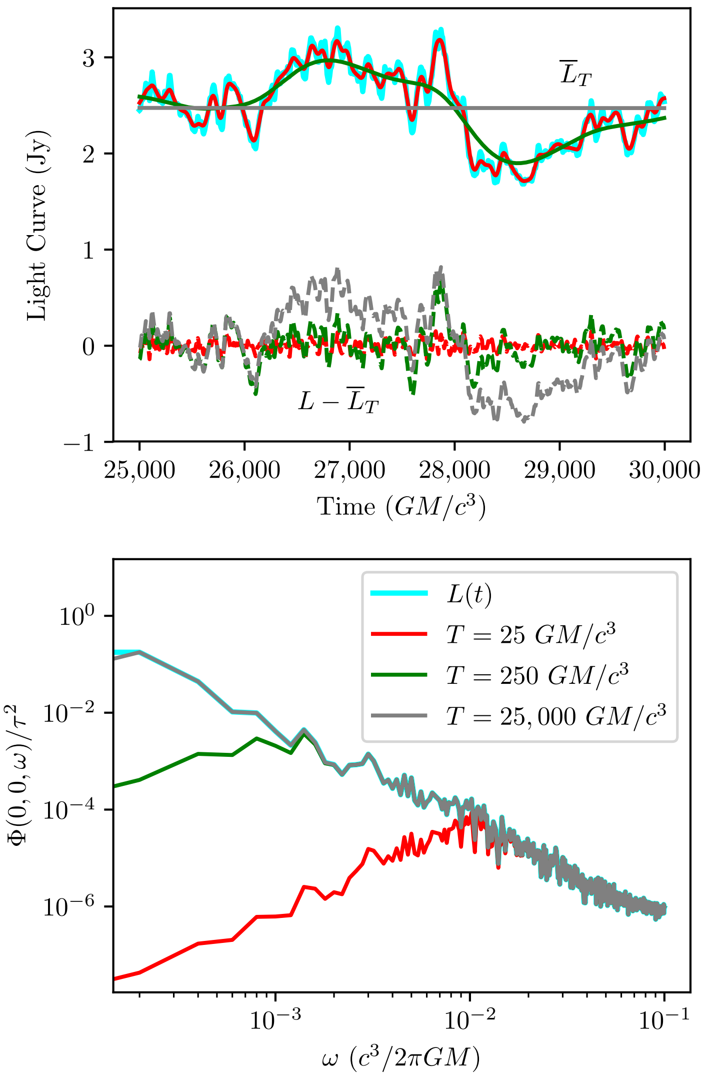

Figure 3.4: Illustration of the Gaussian mean-subtraction procedure for the same simulation as in Figure 3.2 (model set A, SANE, $a_* = 0.5$, $R_{\mathrm{high}} = 40$, and $i = 90°$). *Top:* Solid lines show the light curve and its mean on different timescales T. The dashed lines show the mean-subtracted light curve. *Bottom:* Solid lines show the power spectral density of the mean-subtracted light curves and the original light curve. The effect of mean-subtraction is equivalent to multiplying the power spectral density of $L(t)$ by the filter in Equation 3.18. The cyan curve in the bottom panel is the same as that in the top right panel of Figure 3.3.

3.4.2 Mean Subtraction

The spatiotemporal PSD evaluated at $(u, v) = (0, 0)$ represents the complex square of the Fourier transform of the light curve. The top panel of Figure 3.4 shows an example light curve $L(t)$, time-smoothed versions of that light curve $\bar{L}_T(t)$ for several choices of the averaging timescale T (solid lines), and the residual light curve after subtracting out these averaged versions (dashed lines). The bottom panel of Figure 3.4 shows the temporal PSD of these residual light curves, equivalent to the $(u, v) = (0, 0)$ slice of the spatiotemporal PSD of the mean-subtracted image.

The act of constructing an average mean-subtracted PSD on an averaging timescale T is equivalent[6] to multiplying Φ by the filter described in Equation 3.18. Owing to the red-noise nature of the system, the resulting (i.e., filtered) spatiotemporal PSD peaks near $\omega = (2\pi T)^{-1}$; this effect is visually apparent in the bottom panel of Figure 3.4. One consequence of this effect is that when performing an integral of the mean-subtracted PSD with respect to ω (such as in Equation 3.18), the result will be dominated by a single timescale; we thus have $\langle P_T \rangle(0, 0) \approx \Phi(0, 0, 1/2\pi T)$.

This equivalence between the mean subtraction procedure and a simple evaluation of the spatiotemporal PSD extends beyond just the light curve (i.e., it holds even when u or v are nonzero). For all spatial scales on which the temporal variability is a red-noise process, $\langle P_T \rangle(u, v)$ will approximately measure $\Phi(u, v, 1/2\pi T)$.

3.4.3 Spatial Properties and Averaging Timescale

The spatiotemporal PSD evaluated at $\omega = 0$ represents the complex square of the Fourier transform of the average image (see, e.g., the red curves in the middle column of Figure 3.3). Because these GRMHD simulations designed for Sgr A* typically exhibit a ring of emission in the image domain (see Figure 3.1), they often show a Bessel function-like behavior in the Fourier domain with a first minimum near $\sim$2–3 Gλ.

The spatio-temporal PSD evaluated at $\omega > 0$ measures the variability power at a particular temporal frequency, as a function of the spatial size of the contributing fluctuations in the image. Figure 3.5 shows a one-dimensional slice of the average mean-subtracted PSD, $\langle P_T \rangle(0, v)$, for a MAD simulation with $a_* = 0.94$, $R_{\mathrm{high}} = 40$ and an inclination

---

[6] This equivalence only strictly holds in the limit where the total duration of the light curve is much longer than the averaging timescale T. Effects when this property does not hold are explored in Subsection B.1.3.

of 60 degrees (the slowlight simulation from model set B_s) on several different averaging timescales T. From the arguments in Subsection 3.4.2, $\langle P_T \rangle(0, v) \approx \Phi(0, v, 1/2\pi T)$, and thus $\langle P_T \rangle(0, v)$ contains information about the variability on approximately a single timescale. We can see that for large T, the amount of variability as measured by this average mean-subtracted PSD is larger. As $T \to \infty$, $\langle P_T \rangle(0, v)$ shows a broken power-law behavior, exhibiting a flat spectrum in v up to some threshold spatial frequency, and then a falling power law at higher spatial frequencies. On spatial scales larger than about $200\,\mu$as (corresponding to spatial frequencies smaller than $1\,\mathrm{G}\lambda$) the simulations contain no significant structure, and thus no significant variability, so the average mean-subtracted PSD on these scales must be flat. On smaller spatial scales (i.e., higher spatial frequencies), the observed power-law behavior reflects a combination of the average image structure as well as the dissipative MHD turbulence cascade; we explore this small-scale behavior in more detail in Subsection 3.4.5. For shorter averaging timescales, the total amount of variability power decreases and the location of the break moves to higher spatial frequencies.

The zero-baseline limit of $\langle P_T \rangle(0, v)$ only measures properties of the light curve. Specifically, as in Equation 3.14, it measures the variance of the light curve on the timescale T. Figure 3.5 shows that $\langle P_T \rangle(0, 0)$ for this simulation roughly scales as T^2 for small T, so the temporal power spectrum of the light curve should fall as ω^{-3} (see Subsection 3.3.3). On large finite spatial scales, for all T where $\langle P_T \rangle(0, v)$ is constant in v, the temporal power spectrum of the variability should scale with ω identically to the light curve. We note that other simulations (i.e., with different accretion parameters) can in general have different temporal spectral power law indices, but this falling power-law behavior at large ω is common to all.

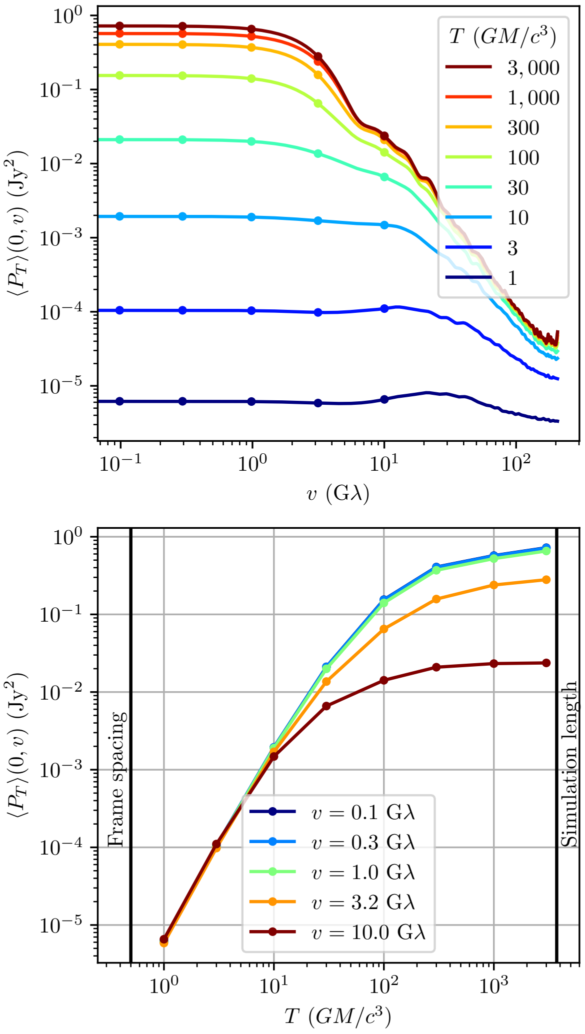

Figure 3.5: Effect of the averaging timescale on the average mean-subtracted PSD. The top panel shows $\langle P_T \rangle(0, v)$ for different T. All show a flat spectrum up until some breaking spatial frequency, after which the spectrum falls as some power law. As the averaging time decreases, the amount of power in the variability similarly decreases and the break location moves to larger values of v. The bottom panel shows the same information as a function on T. For small T, every spatial scale follows the same rising power law with a flattening turnover at different time scales for different v. The simulation used is the single slowlight simulation in model set B$_s$, which is MAD, with $a_* = 0.94$, $R_{\text{high}} = 40$, and an inclination of 60 degrees.

3.4.4 Effects of Normalization by the Light Curve

On short baselines, we expect the PSD to be dominated by variability that is heavily correlated with the light curve variability. By light-curve-normalizing as described in Subsection 3.3.1, we thus expect to remove a large fraction of the short-baseline variability. The two rows in Figure 3.3 show the same simulation before and after normalizing by the light curve. By definition, all variability is removed from $(u, v) = (0, 0)$, but we further see a decrease in the PSD at all spatial frequencies below $\sim 1\,\mathrm{G}\lambda$. We interpret this decrease in power on small spatial frequencies to mean that the majority of variability on large spatial scales is indeed strongly correlated with the variability in the light curve. At higher spatial frequencies, the PSD remains unchanged, indicating that the variability on small spatial scales is independent of the large-scale variability.

Figure 3.6 shows the dependence of $\langle \hat{P}_T \rangle$ on the averaging timescale T. For spatial frequencies larger than the break location, $\langle \hat{P}_T \rangle$ shows the same behavior as $\langle P_T \rangle$, i.e. a power-law with the same index. The break locations of the two PSDs are not equivalent, with $\langle \hat{P}_T \rangle$ breaking at smaller spatial scales than $\langle P_T \rangle$. On larger spatial scales, the flat part of $\langle P_T \rangle$ has been reduced to an increasing power-law in v with a power-law index that is independent of T.

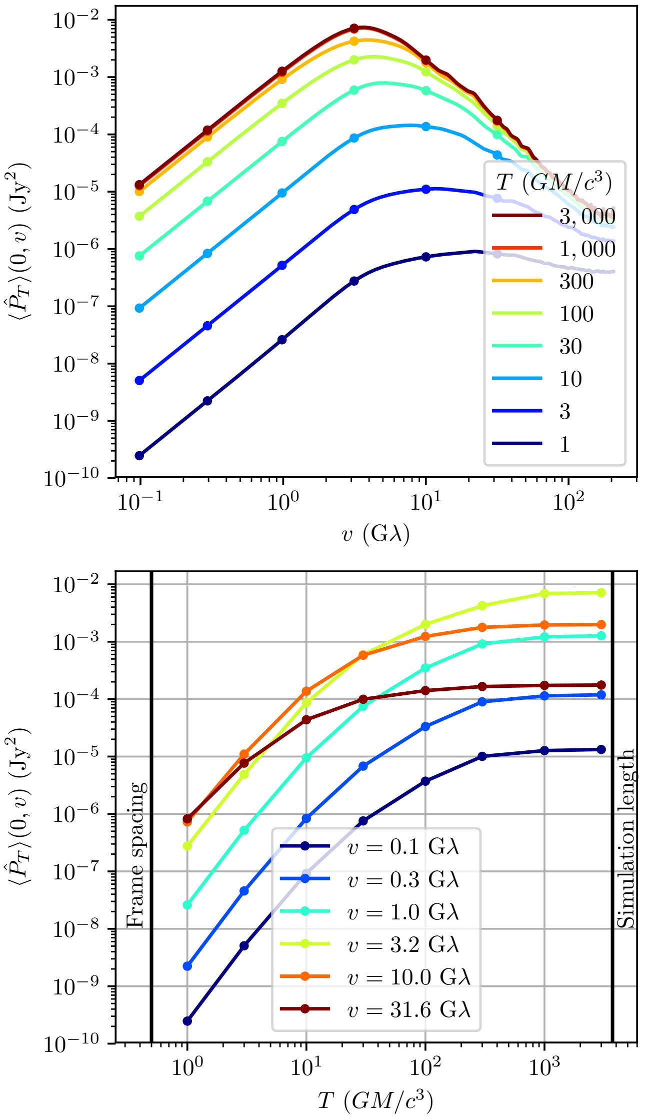

Figure 3.6: Same as Figure 3.5, but for the average residual PSD. $\langle \hat{P}_T \rangle$ shows an increasing power law at short baselines up to some break, after which the spectrum falls as some different power law. Similar to $\langle \hat{P}_T \rangle$, as the averaging time decreases, the amount of power in the variability decreases, and the break moves to larger values. The simulation used is the single slowlight simulation in model set B_s, which is MAD, with $a_* = 0.94$, $R_{\mathrm{high}} = 40$, and an inclination of 60 degrees.

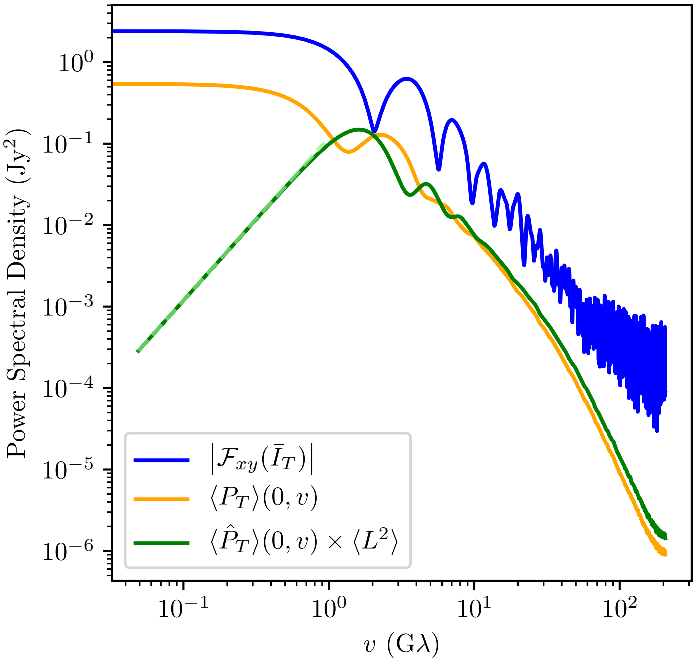

Figure 3.7: Profiles of the Fourier amplitudes of the average image, average mean-subtracted PSD, and the average residual PSD. The timescale used is $T = \infty$, and the average residual PSD has been rescaled by $\langle L^2\rangle$ such that it retains units of Jy^2. On long baselines, $\langle P_T\rangle$ and $\langle \hat{P}_T\rangle$ are the same, as the variability there is almost independent from the light curve. On short baselines, $\langle P_T\rangle \sim v^0$ and $\langle \hat{P}_T\rangle \sim v^2$ (green dashed line) as expected from Subsection 3.3.2. The simulation used from model set A is the same as in Figure 3.2 and is in the SANE state with $a_* = 0.5$, $R_{\mathrm{high}} = 40$, and $i = 90°$. This figure differs from figure 7 of Georgiev et al. (2022), which erroneously labeled the blue curve as the square Fourier amplitudes.

3.4.5 Short and Long Baseline Power Law

Figure 3.7 shows $\langle P_T \rangle(0,v)$ and $\langle \hat{P}_T \rangle(0,v)$ for the same simulation as used in Subsection 3.4.1, as well as the PSD of the average image. On the smallest spatial frequencies, we can see the $\langle \hat{P}_T \rangle(0,v) \propto v^2$ behavior expected from Subsection 3.3.2, and this behavior continues up to a spectral break near $\sim$1 Gλ. The higher-order terms from Equation 3.13 must therefore remain subdominant to the v^2 term below the break. From this empirical behavior, we can deduce that the location of the break lies approximately on the v^2 power law defined in Equation 3.16. The break therefore contains information about the variability of the centroid of emission in the total image.

For any image structure $I(x,y,t)$, we can generically decouple the time-variable component from the average image as

$$\begin{aligned} I(x,y,t) &= \bar{I}(x,y) + \delta I(x,y,t) \\ &= \bar{I}(x,y)\left[1+\delta(x,y,t)\right], \end{aligned} \tag{3.22}$$

where $\bar{I}(x,y)$ is understood to be the time-averaged image as $T \to \infty$ and we have defined $\delta \equiv \delta I/\bar{I}$. Upon applying our mean-subtracting procedure (see Equation 3.5), we would measure

$$\begin{aligned} P_T(u,v,t) &= \left|\mathcal{F}_{xy}\left[I(x,y,t) - \bar{I}(x,y)\right]\right|^2 \\ &= \left|\mathcal{F}_{xy}\left[\bar{I}(x,y)\delta(x,y,t)\right]\right|^2 \\ &= \left|\mathcal{F}_{xy}\left[\bar{I}(x,y)\right] * \mathcal{F}_{xy}\left[\delta(x,y,t)\right]\right|^2, \end{aligned} \tag{3.23}$$

where the $*$ symbol represents convolution. If the Fourier transforms, $\mathcal{F}_{xy}\left[\bar{I}\right]$ and $\mathcal{F}_{xy}\left[\delta\right]$, have uncorrelated phases, and if both have amplitudes that fall as a power-law at large (u,v), then $P_T(u,v,t)$ will be dominated by the shallower power-law component at long baselines.

Figure 3.7 shows that $\langle P_T \rangle$ follows a similar power law at large spatial frequencies as the Fourier transform of the average image, suggesting that the variability in GRMHD simulations is perhaps better described by an uncorrelated multiplicative random field δ, rather than an uncorrelated additive field δI. Under this assumption, we would expect the long-baseline power law of the variability in GRMHD simulations to match that of their average images and would conclude that the spectrum of the "true" variability δ is steeper than that of the average image. Alternatively, the average image power law could be steeper, in which case the long baseline behavior of the variability PSDs would be dominated by δ.

At the longest baselines, there is a slight upturn in the squared amplitudes of the average image, and in both variability PSDs. We attribute this to the limited numerical accuracy with which image data from the GRMHD image library is saved (keeping only six nonzero digits). The associated truncation error introduces a white noise floor that while negligible for the baselines probed by the EHT, can dominate above 100 Gλ.

3.4.6 Azimuthal Dependence

Figure 3.8 shows $\langle \hat{P}_T \rangle(u, v)$ as $T \to \infty$ for the same simulation as in Figure 3.2 and Figure 3.3. Near $(u, v) = (0, 0)$, the average residual PSD has little power due to the light-curve-normalization. On short baselines, from Equation 3.16, a constant $\langle \hat{P}_T \rangle(u, v)$ occurs on an elliptical contour, with the minor and major axis in the direction of the largest and smallest variation in the centroid of emission, respectively. This elliptical behavior remains as long as the u^2 dependence, up until the break location. The major and minor axes are often either aligned with or perpendicular to the direction of the black hole spin axis (see Section B.4). On spatial frequencies larger than the break location, the power-law falls with an index determined by the average image in every direction (see Subsection 3.4.5).

3.5 Variability for Different Accretion Flows

Given our understanding of the morphology of the GRMHD variability PSDs from Section 3.4, we now explore the trends seen with different simulation types.

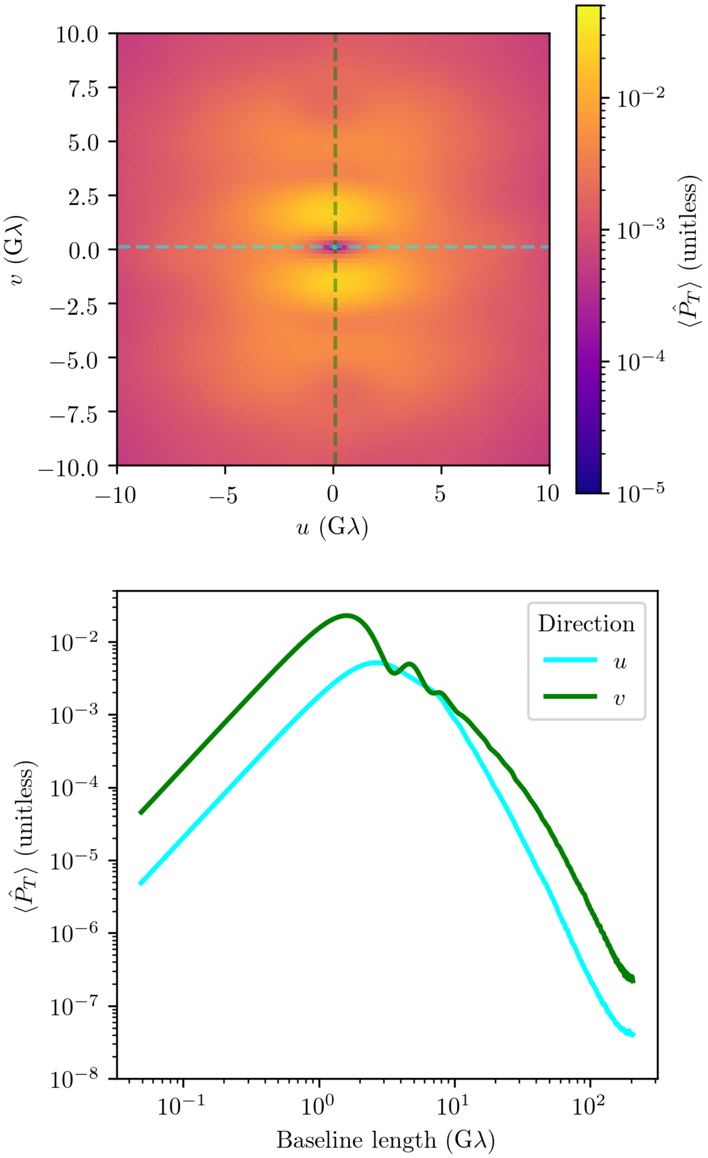

Figure 3.8: Azimuthal dependence of $\langle \hat{P}_T \rangle$ for the same simulation as in Figure 3.2 (i.e., a SANE simulation from model set A with $a_* = 0.5$, $R_{\text{high}} = 40$, and $i = 90°$). The top panel shows $\langle \hat{P}_T \rangle$ in the infinite T limit, with zero power at the center corresponding to the light-curve-normalization. The bottom panel shows a horizontal and vertical slice.

3.5.1 Universal Power Law

Figure 3.9 shows the visibility amplitudes of the average image, $\langle P_T \rangle$, and $\langle \hat{P}_T \rangle$ for all the simulations in model set A for infinite averaging time T. We explore these PSDs only along and orthogonal to the direction of the black hole spin axis, as these typically encompass the extrema of the variability (see Subsection 3.4.6 and Section B.4). The variability for each simulation follows the behavior described in Section 3.4, and each variability PSD can be approximated as a broken power law. Moreover, the different simulations are remarkably consistent over different magnetic field configurations, black hole spins, temperature prescriptions, and viewing inclinations, with similar broken power law parameters.

The short baseline behavior of flat $\langle P_T \rangle$ and a $\langle \hat{P}_T \rangle$ that rises as $|u|^2$ match the expectations from a finite size geometry. The amount of power at short baselines and the location of the break are unique to each simulation and likely contain physical information about the underlying variability. Between about 2 Gλ and 30 Gλ, the long-baseline behavior is similar to that of the underlying average image. This suggests a common origin for both, e.g., the variability could be described as an average image multiplied by an uncorrelated random field (see Subsection 3.4.5).

This behavior can occur if the local variability in the fluid is random and caused by processes that scale with the density (e.g., rotational velocity, shearing). The global variability would then scale with the average density, and where the emitted intensity scales with the density, the variability in the observed image would appear multiplicative to the average image. However, the departure between the variability PSDs and that of the background average image above 30 Gλ implies that the relationship between the small-scale fluctuations and the time-averaged image structure is more complicated. More precise identification of the physical source of the universality of the variability PSDs is complicated by the interactions of the global accretion state, the emissivity, and the radiative transfer along photon trajectories, and we leave its study to future work.

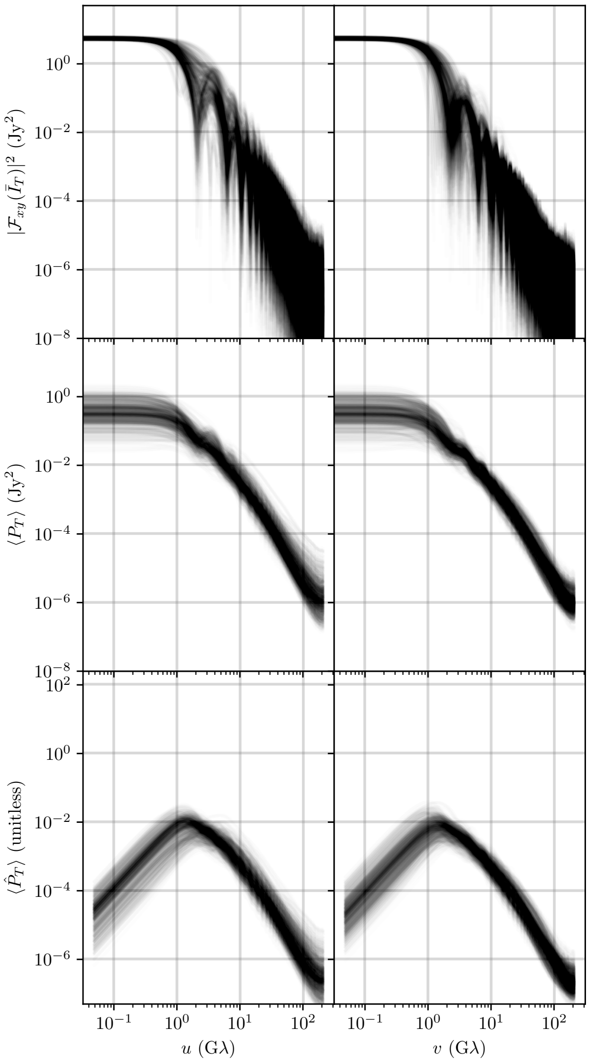

Figure 3.9: The average image visibilities (top), average mean-subtracted PSD (middle), and average residual PSD (bottom) along the u (left) and v (right) directions for the 360 simulations in model set A (KHARMA). All PSDs use infinite averaging time. For $\langle P_T \rangle$, all simulations show a flat spectrum for small baselines and a power-law dropoff at large baselines. For $\langle \hat{P}_T \rangle$, all simulations show a broken power-law spectrum, rising as $|u|^2$ for small baselines, and a power-law dropoff at large baselines. The long-baseline power laws are remarkably consistent among different simulation types.

3.5.2 PSD Dependence on Flow Parameters

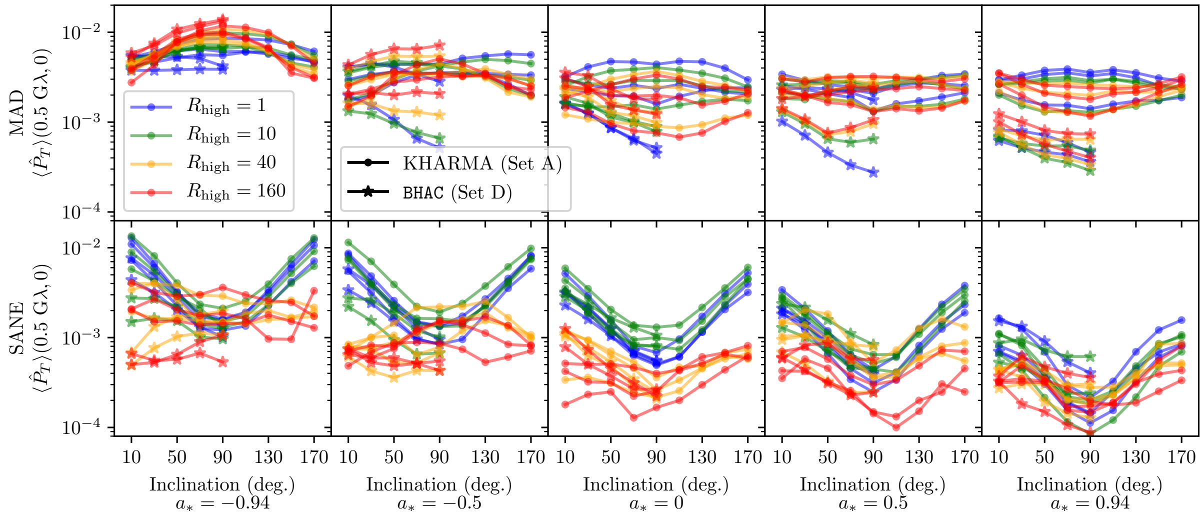

Figure 3.10: Effect of simulation parameters on the average residual PSD at (u, v)=(0.5 Gλ, 0) and infinite averaging time T. The simulations used are from model sets A and D, and each of their windows is shown. The measurement is proportional to the variance in the centroid of emission orthogonal to the black hole spin axis.

We now turn our attention to understanding the dependence of the variability PSDs on GRMHD simulation parameters. The parameter space and possible PSD measurements are vast, and trends across simulations are typically limited to small pieces of the parameter space. Therefore, in this section, we limit ourselves to explaining the trends across models for two measurements from the variability PSDs and briefly comment on other extensions.

Figure 3.10 shows the measurement of $\langle \hat{P}_T \rangle (0.5\ G\lambda, 0)$ for infinite averaging time T for each of the 360 models in model set A and 200 from model set D. The primary source of uncertainty is that each window is short enough to have a separate realization of turbulence. We show all three or two windows separately to provide a visual representation of the spread in this measurement, which matches the 25% estimate from the analysis in Subsection B.1.3. This measurement at short baselines is dominated by the $|u|^2$ dependence and thus only probes variability caused by the variance in the centroid of emission orthogonal to the spin axis (i.e., along the disk). For a visual representation to accompany the following arguments, we refer the reader to Section B.4, which shows this variance relative to its location in the average images.

MAD simulations show little change across spin, inclination or R_{high}. This is not necessarily surprising: in MAD flows, the emission is typically produced close to the event horizon and is centered around the midplane, in the region where filamentary accretion strands are separated by the strong magnetic pressures and plunge from intermediate radius down to the horizon. When the majority of the emission comes from this violent, disorganized region, the precise details of the accretion model and thermodynamics may not strongly influence the detailed morphology of the emission.

For SANEs, however, there is a clear decrease in variability for larger corotating spin, larger R_{high}, and more edge-on viewer inclination. In contrast with the MADs, the morphology of the SANE emission region can be much more strongly affected by the details of the model, and the split between $R_{\mathrm{high}} = 10$ and $R_{\mathrm{high}} = 40$ coincides with a change from emission dominated by the disk or the jet sheath, respectively. For disk-dominated $R_{\mathrm{high}} \lesssim 10$, the variability is much more pronounced for face-on inclinations, where any part of the disk may light up, than for edge-on inclinations, where the emission, and therefore variability, is limited to the Doppler-boosted gas moving toward the observer on only one side of the black hole. For jet-dominated $R_{\mathrm{high}} \gtrsim 40$, the emission comes from further off the midplane and is not as beamed, thus changing inclination has little effect on the emission region size orthogonal to the spin axis. Note that for these jet-dominated models the component along the spin axis does depend on inclination.

Figure 3.11 is analogous to Figure 3.10, but showing $\langle \hat{P_T} \rangle(0, 0.5\ \mathrm{G}\lambda)/\langle \hat{P_T} \rangle(0.5\ \mathrm{G}\lambda, 0)$ for infinite average time T. This is in the short-baseline regime, and therefore we can interpret trends through the axial ratio of the covariance ellipse of the distribution of centroids of emission. MAD and SANE simulations are split, with MADs seeing more variability along the disk where the average bulk of the emission is concentrated, and SANEs seeing more along the jet, which is especially lit up in models with large R_{high}. SANE simulations show a split between disk- and jet-dominated emission as in Figure 3.10. The simulations dominated by jet emission have small ratios (i.e., < 1) since the variability comes from different parts of the jet and counter-jet lighting up. This effect becomes more pronounced at edge-on inclinations where the average image resembles two disjoint emitting regions along the direction of the spin axis.

For long baselines, the variability PSDs are related to the power spectrum of the average image. Therefore, any trends seen in the variability would likely be similar to trends in the average image, which can more easily be obtained from data. On short timescales (small T), simulations show a clear power law dependence on T with a temporal break, which, like the two PSD measurements discussed in this section, can discriminate accretion flow types. However, on shorter timescales, it becomes increasingly difficult to provide a physical origin for the PSD measurement trends, and we leave its study to future work.

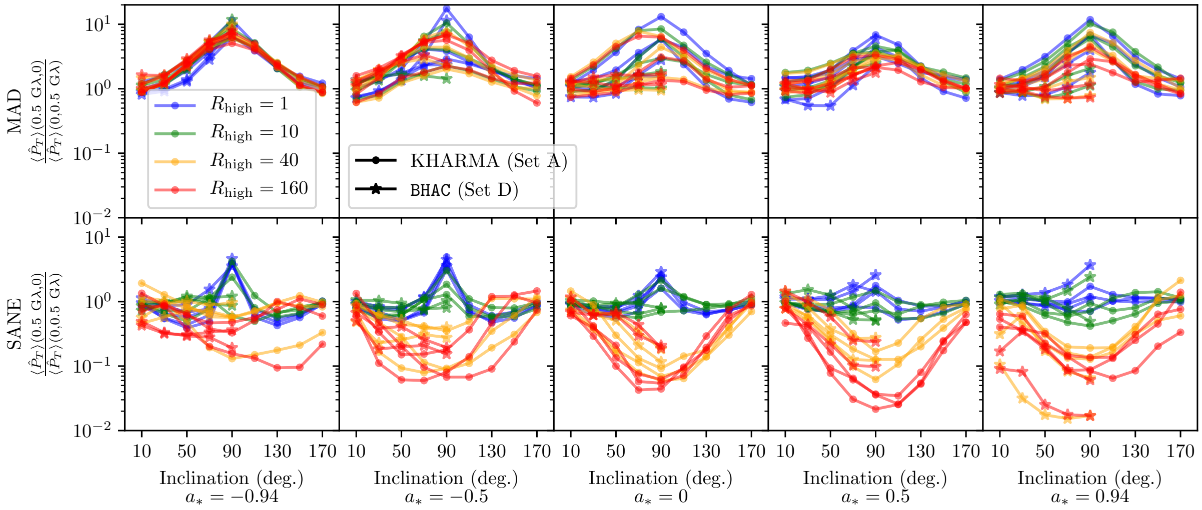

Figure 3.11: Effect of simulation parameters on the ratio of the average residual PSDs at (u, v)=(0.5 Gλ,0) and (u, v)=(0,0.5 Gλ) for infinite averaging time T. The simulations used are the same as in Figure 3.10. The measurement is equivalent to the ratio of the variance of the centroid of emission orthogonal to the black hole spin axis to that along the black hole spin axis.

3.5.3 Application to Sgr A*

Current EHT observations of Sgr A* provide only sparse temporal and spatial coverage of the (u, v)-domain. A measurement of the variability PSDs from EHT data is further obscured by instrumental and calibration effects. However, tests on synthetic data from GRMHD simulations in Broderick et al. (2022b) show that $\langle \hat{P}_T \rangle$ can be well constrained with current observations for 2 G$\lambda \lesssim |u| \lesssim 6$ Gλ.

Motivated by these findings, we show measurements of $\langle \hat{P}_T \rangle$ from GRMHD simulations that have been generated to mimic the EHT observations of Sgr A*. For each simulation, we diffractively scatter the image with the kernel from Johnson et al. (2018) to include the effects of interstellar scattering. We then average its variability PSDs over all possible position angles of the scattering screen, and then azimuthally average. Finally, we fit each resulting scattered $\langle \hat{P}_T \rangle_\phi$ with a power law between 2 G$\lambda < |u| < 6$ Gλ and take our measurement to be the value of the fit at $|u| = 4$ Gλ. This procedure is shown for the simulation used in Subsection 3.4.1 in Figure 3.12.

The equivalent of the full scattered $\langle \hat{P}_T \rangle_\phi$ is measured from EHT data in Paper IV,

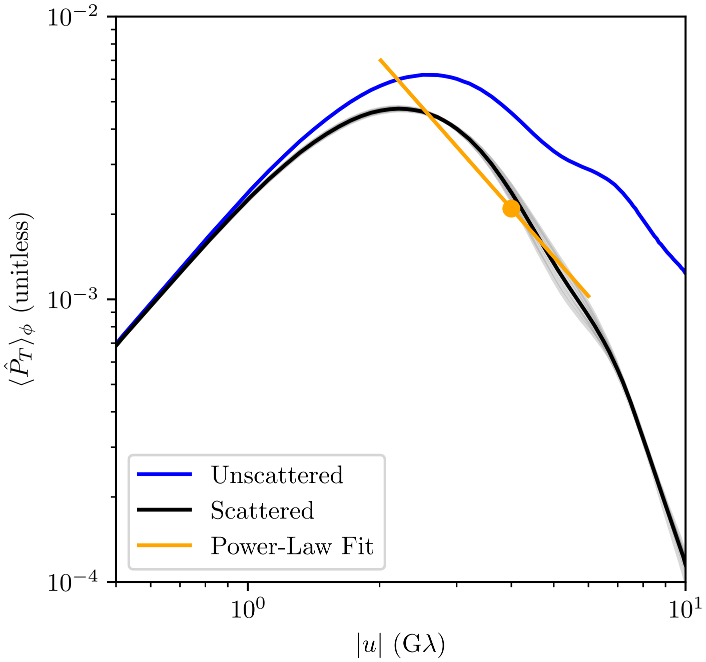

Figure 3.12: Example of the diffractive scattering and fitting procedure. The blue line shows the azimuthally-averaged $\langle \hat{P}_T \rangle$ for the same simulation as in Subsection 3.4.1 for infinite averaging time T. The six overlapping gray lines correspond to $\langle \hat{P}_T \rangle_\phi$ after multiplication by a diffractive scattering kernel with an orientation of 0°–180° relative to the black hole spin axis. The average over these is shown as a black line, to which we fit a power-law between 2 G$\lambda < |u| <$ 6 Gλ (orange line), and take as a measurement the value of this fit at $|u| =$ 4 Gλ (orange point).

and the amplitude and slope of this power-law measurement are compared to data and interpreted in Paper V[7].

Figure 3.13 shows this measurement constructed for infinite averaging time T. Unlike at short baselines, $\langle \hat{P}_T \rangle$ near 4 Gλ is similar for many of the considered accretion flow parameters, predicting a universal amount of GRMHD variability to within an order of magnitude.

The main trend is with inclination, with face-on models showing less variability than edge-on models. This is likely mirroring average visibilities, where face-on models have a

---

[7] The convention in Paper V and Section B.2 differs by a factor of 2, due to the use of visibility amplitudes instead of complex visibilities. The rest of the procedure is identical.

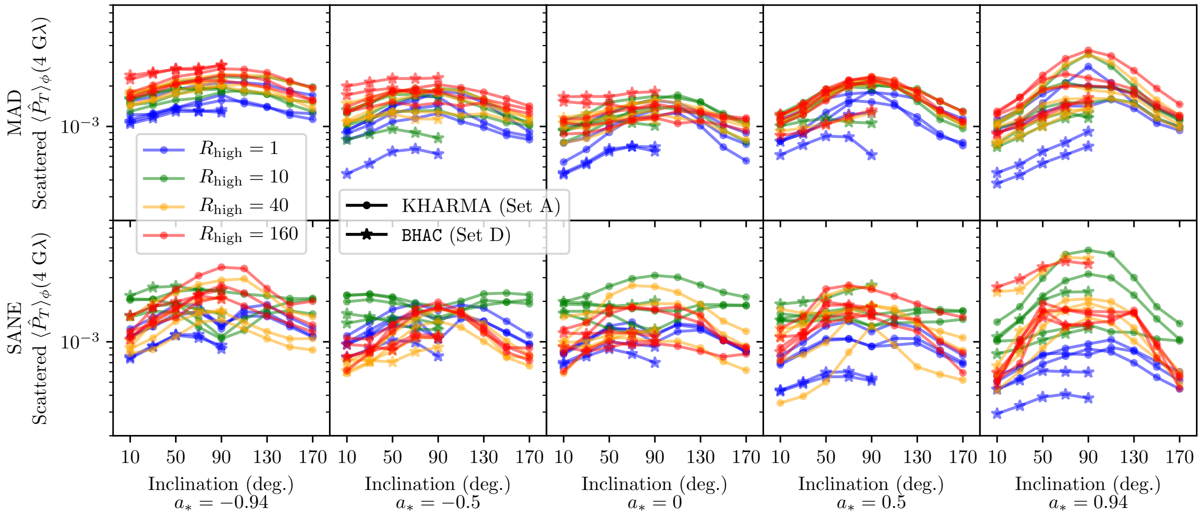

Figure 3.13: Azimuthally averaged $\langle\hat{P}_T\rangle$ evaluated at observationally tractable baseline lengths for infinite averaging time T. The average residual PSD is diffractively scattered, azimuthally averaged, and averaged over orientations of the scattering screen. The value of the scattered $\langle\hat{P}_T\rangle_\phi$ at 4 Gλ comes from a power-law fit at long EHT baseline lengths. Though it can be well constrained by data, this quantity is from a region of (u, v)-space where the average residual PSD is common among many accretion flow types.

larger area of emission, thus larger variability. Further discriminatory power is limited by the noisiness of the simulations, with separate windows and codes providing the leading uncertainty in the measurement of $\langle\hat{P}_T\rangle$.

3.6 Summary and Conclusions

Motivated by the EHT observations of Sgr A*, in this paper, we have developed a framework for characterizing the spatiotemporal power spectrum of variability observed towards an accreting supermassive black hole system. We have applied this framework to the library of GRMHD simulations and associated GRRT images produced by Paper V, which span a range of physical properties (magnetic flux configurations, black hole spin, temperature prescription), numerical properties (code, resolutions, initialization, simulation length, time from initialization, and other approximations), and observer locations. While these models were produced specifically to match the expected emission from Sgr A* during the 2017

EHT observation period, one can reasonably expect our findings to be applicable to a range of systems in the optically-thin regime, following a proper scaling of a black hole mass and total flux. We find that the variability power spectrum is generically a red-noise process in both the temporal and spatial dimensions; i.e., the simulations exhibit the highest variability power on the longest timescales and on the largest spatial scales. This find will become important when separating coherent variability in Chapter 4. To focus on variability in excess of the spatiotemporal mean, we have considered primarily the PSD behavior after subtracting out the time-averaged source structure and normalizing it by the light curve.

All of the GRMHD simulations show remarkably similar structures in their variability PSDs. The spatial PSDs along any direction can be described by a broken power law, with a small range of long-baseline indices ($\sim 2-3$) and break locations ($\sim 1-3$ Gλ). The long-baseline indices correspond to a combination of the turbulent power spectrum and the average image. The break location informs the largest spatial scale on which variability is present, which is roughly the size of the emission region. On long baselines, the power-law index of the variability PSD tracks that of the Fourier transform of the average image, which is itself similar across simulations. The observed universality in the long-baseline PSD behavior could arise if the variability came in the form of an average image whose structure is modulated by an uncorrelated, multiplicative random field. In this picture, it is the Fourier transforms of the average images which are similar among simulations, and the random field could be caused by any process (e.g., MHD turbulence) that has a long-baseline power-law index steeper than that of the average image.

On any given spatial scale, the power in the variability about the average image increases with averaging time up to some breaking timescale, beyond which it remains constant. That is, for every spatial scale there exists a maximum necessary averaging time, above which the variability power on that spatial scale remains unchanged with increased averaging. On short baselines (below ~ 1 Gλ), this breaking timescale is greater than $10^3\ GM/c^3$, and it is shorter on longer baselines. The implication for EHT observations of M87* is that the variability "noise" seen across a single night of observing should be approximately constant with baseline length, and the expectation from GRMHD simulations is that the average magnitude of this variability should be less than 1% of the total flux. Across a week of observations, the expected variability power increases to several percent, though it remains flat with baseline length on the range of baselines probed by the EHT. These expectations qualitatively match the observed behavior of M87*, which the EHT found to exhibit little or no intra-day variability, and a modest amount of inter-day variability (M87* Paper IV). Significant structural variation can be expected on timescales of months and years, which has been observationally confirmed by Wielgus et al. (2020). For Sgr A*, we expect that the variability should manifest as a broken power law "noise," with large contributions on short baselines and decreasing contributions on longer baselines. Normalizing the visibility amplitudes by the light curve removes the largest component of this variability noise.

The azimuthal structure of the variability PSDs tracks the average images, that is, the variability typically occurs where the emission does. Especially on longer baselines where the variability PSDs and average images are similar, we expect the variability to discriminate between accretion flows less than the average images could, and the latter are observationally easier to measure. The variability can better distinguish different accretion flows from measurements on short baselines ($|u| \lesssim 1$ Gλ), where the variability is larger. Here, the variability power is proportional to the variance in the centroid of emission, the most variable spatial mode after the light curve. The entire temporal PSD of a short baseline measurement of the variability (i.e. for a time series of the centroid of emission) for timescales greater than $\sim 30\ GM/c^3$ contains unique discriminating power. However, with the current EHT instrumental effects, such a measurement would contain degeneracies with the average image.

We find that a measurement of the average residual PSD on short baselines – and thus a measurement of the covariance ellipse of the centroid of emission – wields some power to discriminate between the different accretion flows in our GRMHD library. In particular, the MAD and SANE accretion states exhibit distinct emission centroid behaviors relative to the direction of the black hole spin axis. A centroid ellipse that is strongly elongated along the direction of the spin axis suggests a variable jet, which is often seen in SANE simulations that support hot funnels. In such a case, the magnitude of variability orthogonal to the spin axis can constrain the black hole spin, with more positive spins producing less variability in this direction. Centroid ellipses that are strongly elongated along the direction of the disk (i.e., assumed perpendicular to the spin axis) are instead more typical of edge-on MAD states, which have more equatorial emission structures. In this case, spin constraints are more difficult because the variability in MAD simulations does not significantly change with spin or temperature prescription. Circular centroid ellipses suggest a face-on accretion flow, or perhaps an emission region that is further from the black hole. In such cases, the amount of variability in the direction of the disk could distinguish whether the accretion flow is in the SANE state with emission dominated by hot electrons in the disk. For each of the cases described, the temporal behavior also holds significant discriminating power, though it remains a more uncertain measurement from the simulations used here, and is left for future study.

Lastly, we note that the work and conclusions rely entirely on the assumption that the accretion flows around black holes and Sgr A* in particular are described by ideal GRMHD and that the assumptions within those simulations are valid. We have attempted to address the latter by individually studying small excursions in resolution, the fastlight approximation, accretion flow evolution, random initialization, and further assumptions made with different simulation codes. Primarily, we find that the variability evolves on long timescales and that even simulations run for as long as $30,000\ GM/c^3$ may be insufficient for the characterization of variability statistics.

Chapter 4

A Set of Coherent Variability Metrics to Distinguish Accretion Flow Properties from Synthetic Images of M87*

4.1 Introduction

The EHT can provide independent images for the black hole M87* on each day of observation. The dynamical timescale of M87* is $GM/c^3 \approx 9$ hrs, and thus subsequent days are expected to have strongly correlated images. Combined, these images are a slowly varying movie, one in which we can track features and measure coherent variability. Presumably, moving features in images originate as similarly moving features in the accretion flow; thus variability can inform about the underlying emission mechanisms, dynamics, or the black hole spacetime.

This idea has been applied to M87* in previous work by tracking blobs of emission moving along the jet on scales of 10-1,000 GM/c^2 (Hada et al., 2016; Walker et al., 2016). Closer to the black hole, magnetic reconnection events could create a local spot of accelerated electrons, which either fall into the black hole or are expelled in the jet (Broderick & Loeb, 2006a; Tiede et al., 2020; Jeter et al., 2020). GRMHD simulations of black holes consistently create coherent features, though they aren't necessarily transported with the flow (see Section 1.3 and references therein; Conroy et al. 2023). Thus, it is of interest to connect the coherent variability in EHT images to a physical process in a measurable way to learn about the physics of the accretion flow.

Label	Code	Flux	a_*	$\Gamma_{\rm ad}$	$r_{\rm in}$	$r_{\rm max}$	Resolution	Cadence
A	KHARMA	MAD	0, ±0.5, ±0.94	4/3	20	41	288×128×128	5
A	KHARMA	SANE	0, ±0.5, ±0.94	4/3	10	20	288×128×128	5
B	`iharm3d`	MAD	0.94	13/9	20	41	384×192×192	0.5

Table 4.1: GRMHD fluid simulation parameters. $r_{\rm in}$ and $r_{\rm max}$ are in units of GM/c^2, and the cadence is in units of GM/c^3.

In this chapter, we measure three quantities from a library of GRMHD movies, which are in principle measurable from EHT data, and relate them to details of the accretion flow. The first two are a correlation timescale and an image-averaged rotation rate. Besides relating physical processes to an image, these also help validate the assumptions of independent snapshots in existing analyses of EHT data (M87* Paper V; M87* Paper VI). We find results inconsistent with the idea that coherent variability is caused by features moving with the fluid velocity. Motivated by the first EHT results of M87* and by results in this chapter, we introduce the third measurement as a position angle of a dipole moment, mimicking a rotating spot. A similar analysis of the position angle of GRMHD simulations and EHT data has been carried out in Wielgus et al. (2020). We compare to EHT data and show consistency with the simulations, and present predictions for variability in future EHT observations.

We introduce the library of GRMHD simulations and images in Section 4.2 and describe the methodology of correlation measurements in Section 4.3. We present the results applied to one simulation in Section 4.4 and to the whole library in Section 4.5, which also explores potential origins for an unexpectedly low rotation rate. Finally, in Section 4.6, we measure and compare the dipole moment of the simulation library and EHT data.

4.2 GRMHD Simulations and Images of M87*

We utilize a library of GRMHD fluid simulations which are then imaged to produce movies of M87*. For a complete description of the details of the simulations, see Section 1.3, and for a description of the codes, see Chapter 3. From observations of the jet in M87*, we have a strong constraint on the inclination angle, $i = 17°$ or $i = 163°$ depending on the orientation of the black hole spin and fluid rotation (Walker et al., 2016). Half of these inclinations can be further excluded by the relative position angles of the brightest spot in the EHT images and the large-scale jet (M87* Paper V). The other half of the images correspond to an image flip. Therefore, we take an inclination of 163° for $a_* \geq 0$ and

Label	Code	$R_{\rm high}$	i (deg)	Flux (Jy)	FOV	Res.	Start	End	Cadence
A	`ipole`	1, 10, 40, 160	17,163	0.5-0.56	200	0.5	25,000	30,000	5
A_e	`ipole`	40	17	0.45	160	0.5	5,000	9,000	5
B_s	`ipole`	40	163	0.42	160	1	6.506.5	10,375.5	0.5,3
B	`ipole`	40	163	0.42	160	1	5,488	9,357	0.5,3

Table 4.2: GRMHD image parameters. The simulation A_e is a MAD, $a_* = -0.5$, and is split into elevation slices. The simulation B_s is imaged using slowlight, and all others use the fastlight approximation. The field of view and resolution are in μas. Inclinations are 17° for $a_* < 0$ and 163° for $a_* \geq 0$.

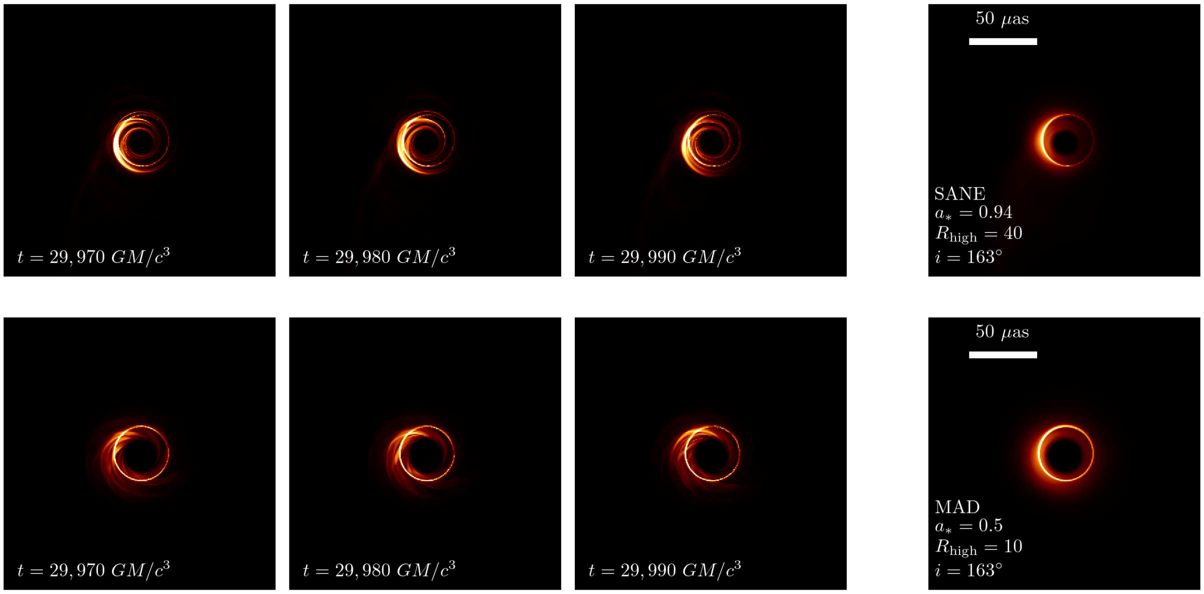

Figure 4.1: Example frames from the GRMHD simulation library. Shown here are a representative MAD (top) and SANE (bottom) simulation with three sequential frames (left) and the average image (right). The inclination used is nearly face-on, with frames typically showing rotating spiral features.

17° for $a_* < 0$. The latter of these has the black hole spin axis pointing away from the observer, and the angular momentum of the disk pointing at the observer. By convention, we choose the y-direction to correspond to these spin axis.

The spatial and density scales of the images are irrelevant in this work as long as the flux is low enough for the gas to be optically thin; following M87* Paper V, the simulations have $M = 6.5 \times 10^9\, M_\odot$ and $D = 16.8$ Mpc. We do not know details of the magnetic flux or the electron temperature, so we survey over the typical MAD/SANE models and $R_{\rm high} = 1, 10, 40, 160$. One of these simulations is split into elevation regions and only emission from that region is considered, resulting in a movie for each elevation slice. A full set of fluid parameters are given in Table 4.1 and corresponding image parameters in Table 4.2. We also use a separate simulation to test the effects of the fastlight approximation and to test the effects of cadence.

Figure 4.1 shows representative images of two example simulations. GRMHD images of M87* show an average image of a bright ring and diffuse emission, dominated by gravitational lensing effects in the optically-thin regime. On top of this, we see spiral features that rotate in the same direction as the accretion flow and are roughly centered at the black hole. The total brightness, asymmetry, and other large-scale image modes also vary in the simulations. We make an inherent assumption in this chapter that the large-scale effects are caused by boundary conditions, i.e., by the structure of the infalling material at large radii. The small-scale spiraling features are likely emerging from the saturated state of the turbulence. These fluctuations depend on details of the fluid, so we now measure their correlation timescales and rotation rates to classify them.

4.3 Image-Averaged Correlations

We wish to measure the correlations of the fluctuations present in the GRMHD movies. Finding them is challenging since the simulations have a red-noise power spectrum, with the largest fluctuations occurring on the longest timescales (see Chapter 3), and the finite simulations do not necessarily probe sufficiently long timescales. We attempt to mitigate this in two ways: by lightcurve-normalizing and by mean-subtracting on a fixed timescale. We also explore here the effects of smoothing the images as an approximation to the resolution of the EHT.

We define an image-integrated correlation of two images $A(x, y, t)$ and $B(x, y, t)$ in a standard way and incorporate a way to focus on single timescales. That is,

$$\mathcal{C} = \frac{\iint (A - \langle A \rangle)(B - \langle B \rangle) dx dy}{\sqrt{\left[\iint (A - \langle A \rangle)^2 dx dy\right] \left[\iint (B - \langle B \rangle)^2 dx dy\right]}}, \tag{4.1}$$

where the integral over x, y is approximated with a sum over pixels, and the angled brackets represent some temporal average. We can get a measure of the simulation correlation by averaging the correlation over all sets of images (A, B) either a time Δt or angle $\Delta\phi$ apart. We denote this $\langle\mathcal{C}\rangle(\Delta t, \Delta\phi)$, and by definition, $\langle\mathcal{C}\rangle(0,0) = 1$.

Note that this quantity is not easily measurable from independent EHT images as the lack of an absolute complex phase in the Fourier data makes each observation independently shifted, scrambling the concept of an average image. This problem goes away in the presence of a stationary feature, e.g., a bright ring of emission. Instrument improvements, such as frequency phase transfer, or phase referencing can ameliorate the inability of the EHT to measure absolute position. We leave the exact resolution of this issue and application to/validation with EHT data to future work.

4.3.1 Lightcurve-Normalization

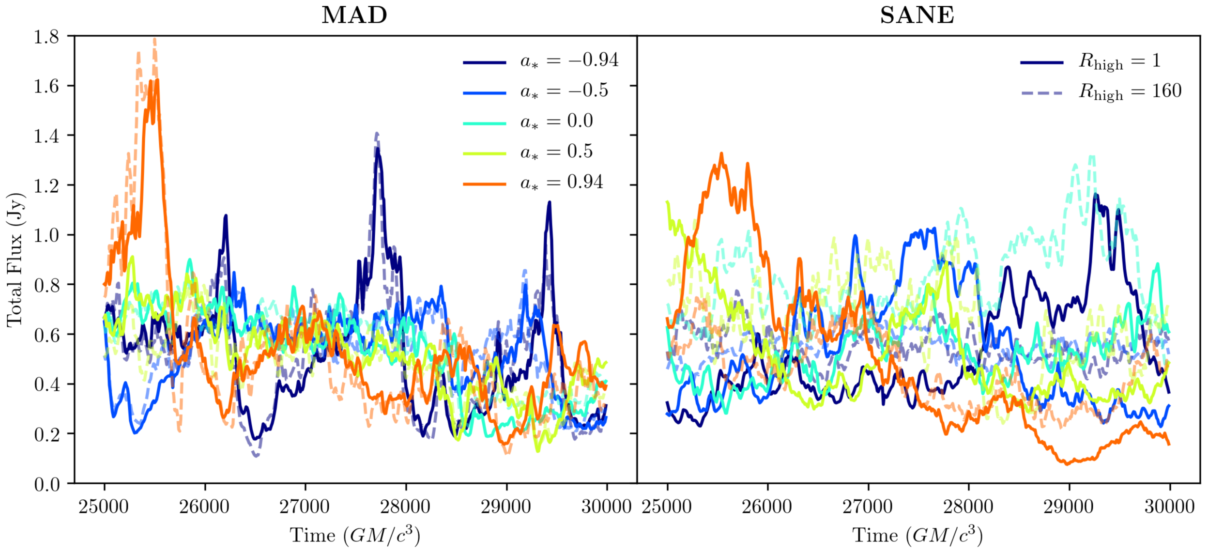

Figure 4.2: Total flux lightcurves of a sample of the simulations used. The left panel contains MAD models for all spin values, and the right panel contains the same for SANE simulations, with only $R_{\rm high} = 1$ (solid) and $R_{\rm high} = 160$ (dashed) shown. The short timescale variability is roughly stochastic, but there are long-timescale trends, such as a decaying flux in the SANE, $a_* = 0.94$ model.

One mitigation strategy is to assume that the long-timescale effects are dominated by a fluctuation in the total flux lightcurve. This is, again, motivated by the findings in Chapter 3, where we found that the largest variability is on the largest spatial scales. Figure 4.2

shows the total flux lightcurves for a representative subset of the models used. The total flux of frames can vary by up to an order of magnitude during the $5,000\,GM/c^3$ time window shown. The existence of these long-timescale trends indicates that the simulations have not been run long enough to reach a steady state. The mass within the simulation, initially organized as a hydrodynamically stable torus, also drains, causing the mass accretion rate (related to the total flux) to drop on timescales of about $15,000\,GM/c^3$. This is particularly evident in the SANE, $R_{\rm high} = 1$ simulation. To counteract this effect, we normalize each frame to have a unit total flux, that is

$$\mathcal{I}(x, y, t) = \frac{I(x, y, t)}{L(t)} = \frac{I(x, y, t)}{\iint I(x, y, t)dxdy}. \tag{4.2}$$

Where the optical depth is large, this is equivalent to changing the mass scale of the simulations. However, near the photon ring, the longer photon path decreases the optical depth, and this will lead to artifacts in the correlations.

We find the lightcurve by summing over pixels and use a sum over frames to find the average of $\mathcal{I}$.

4.3.2 Mean-Subtraction

The second mitigation strategy is to impose a temporal filter on the calculation of the mean. Given an image, $I(x, y, t)$, we define a Gaussian-weighed mean, so that images separated by long times correspond to a correlation of 0, and we are not affected by the finite simulation length. Similarly to Chapter 3, we define a moving mean as

$$\langle I \rangle = \bar{I}_T(x, y, t) = \int I(x, y, t') \frac{1}{\sqrt{2\pi}T} \exp\left[-\frac{(t-t')^2}{2T^2}\right] dt. \tag{4.3}$$

Numerically, this integral becomes a sum over frames, and the normalization is over all included frames. Practically, we need only choose a T longer than any correlation timescale of the small-scale spiral fluctuations, but small enough to get many independent samples in a simulation length.

4.3.3 Smoothing

In order to approximate the effects of the EHT beam size, we convolve the images with a circular gaussian with a standard deviation of $10\,\mu$as, or a FWHM of $\sim 23.6\,\mu$as. Practically, this is done with the multidimensional gaussian filter in `SciPy` (Jones et al., 2001–).

4.3.4 Centering

Lastly, to calculate correlations between images at different angles, we need to choose a center, and there is no unique choice. Most GRMHD images show a bright ring corresponding to lensed emission, which is expected to be near the shadow edge. For a face-on black hole, this curve is a circle whose center[1] changes with spin,

$$x_{\text{center}} = 2a_* \sin 17^\circ \frac{GM}{Dc^2} \approx 2.22 a_* \mu\text{as}. \tag{4.4}$$

We instead find this center from the average image using the following procedure applicable to EHT images. First, we guess the center location and draw 36 rays emanating from this position at 10-degree intervals. For each, we find the radius to the maximum brightness location and find the standard deviation of these radii among the different rays. We choose the center of the image as the one that minimizes this standard deviation. This center is used for all frames of the simulation and is not found independently on the smoothed images or the lightcurve-normalized images. Note that applying this procedure to EHT data may not key in on the same feature, as our procedure requires a higher resolution to find a brightness peak. This is partly explored in Section 4.6, but will ultimately need to be validated on synthetic data.

Figure 4.3 shows the images before and after shifting to the center of the bright ring. The simulation used is a MAD, with $a_* = 0.94$, $R_{\text{high}} = 1$, and an inclination of 163°, so it is maximally offset by gravitational effects. Before centering, the ring is offset, and the peaks along the cardinal directions do not align with a flip of the image. After measuring a center and shifting the images, the brightness peaks in all directions are better aligned. For the majority of simulations, this center agrees with the center of the shadow (red dot). This procedure will result in measuring rotations about the center of the bright ring, which does not necessarily match every rotating feature in the system.

4.4 Application to One Simulation

We now apply the procedures of the previous section and interpret the generic features simulations produce. We also test the effects of smoothing, lightcurve-normalization, and averaging timescale. In particular, we are interested in the coherent fluctuations which manifest as spiral arms instead of the more prominent total flux and brightness asymmetry variations that are expected from changes in the mass accretion rate far away from the black hole.

---

[1] The GRMHD coordinates are centered at the black hole with the spin axis pointing vertically.

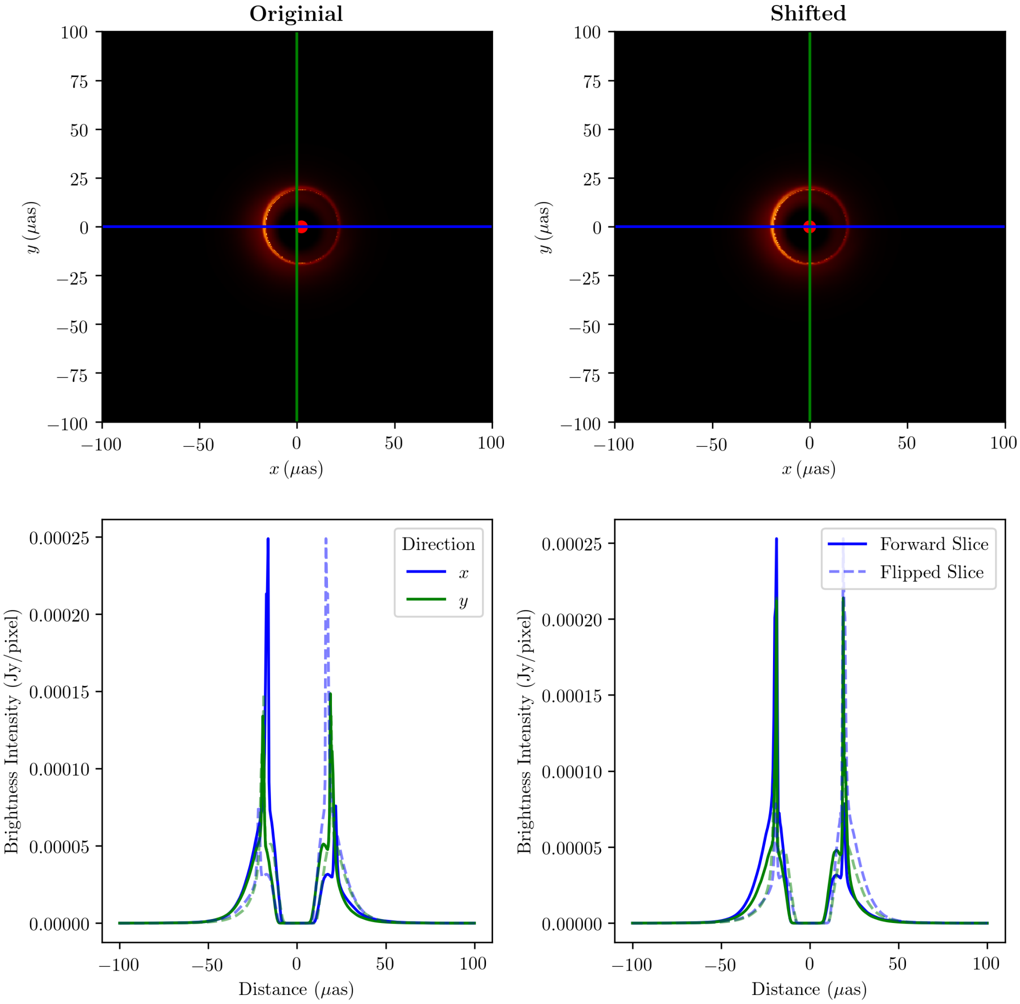

Figure 4.3: Demonstration of the centering procedure. The left column shows the raw GRMHD frame, while the right shows the frame after shifting. The top row contains the images, with the red dot representing the theoretical center of the photon ring given by Equation 4.4. The bottom row shows horizontal and vertical slices along the center of the images as solid lines, with dashed lines flipped along the respective axis for comparison purposes. Before centering, there is a clear offset in the bright ring structure, and after centering, the peaks all lie on top of each other.

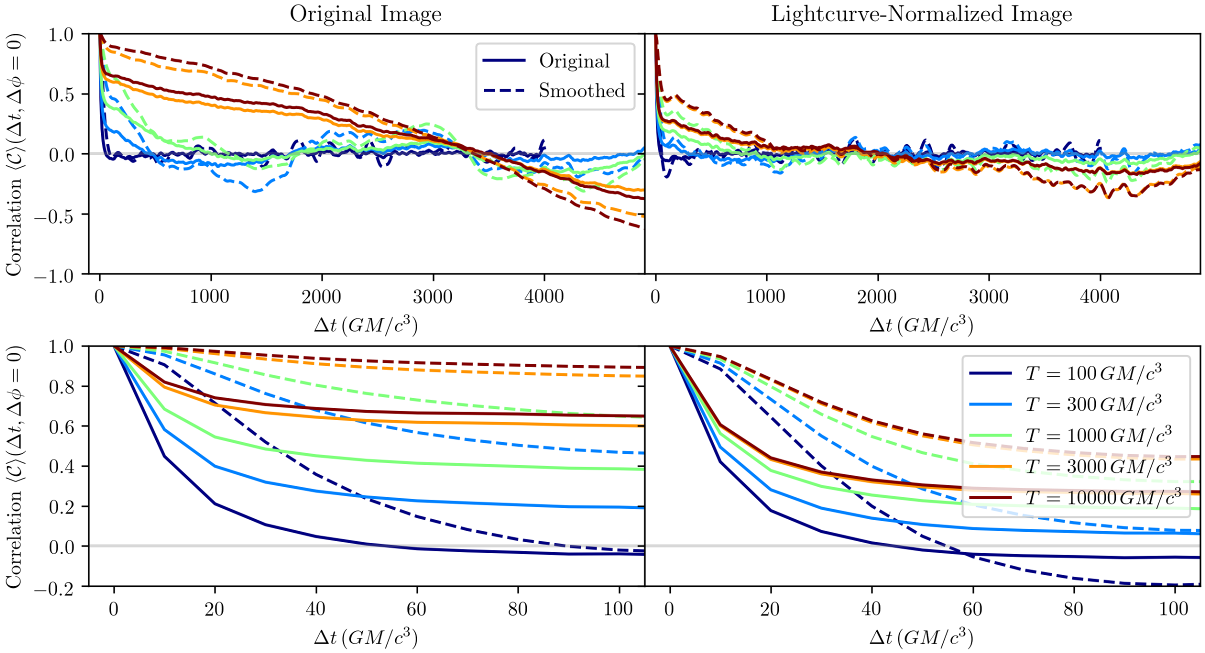

Figure 4.4: Dependence of the correlation on mean-subtraction, lightcurve-normalization, and smoothing for a simulation with the strongest long-timescale trends. Different colors show mean-subtraction on different timescales. Solid lines show the unsmoothed image, and dashed lines show the image smoothed with a gaussian with a standard deviation of $10\,\mu$as. The left panels are the original images, while the right panels use the images after lightcurve normalization. The bottom panels contain the same lines as the top but zoomed in to highlight the short-timescale behaviors.

4.4.1 Long-Timescale Variability

Figure 4.4 shows the effects of smoothing the images, lightcurve-normalization, and the mean image subtraction procedure. The correlation shown involves no rotation and is independent of a center choice. We specifically choose a simulation with one of the strongest long-timescale trends to explore effective mitigation strategies (SANE, with $a_* = 0.94$ and $R_{\rm high} = 1$). With no mitigation, there is a hint of two separate timescales, one on the order of $30\,GM/c^3$ and another of more than $1,000\,GM/c^3$. It is unclear whether this long timescale is physically important or a numerical artifact of either finite simulation length or a draining of the torus.

Applying lightcurve-normalization has a massive effect. Since Equation 4.1 is independent under a scaling factor, the lightcurve-normalization procedure only affects the weighting of each image in the average. If the long-timescale effects are limited to a scaling of the mass accretion rate and the gas were optically thin everywhere, this procedure would remove the effects relating to a draining torus. The correlation function is certainly reduced; however, there is still a residual effect. For large T, there still appear to be two correlation timescales. We attribute this to the photon ring being optically thick, and thus the rescaling has not equally reduced the long timescale correlation in the bright ring structure.

The mean subtraction procedure is equivalent to a filter applied to a power spectrum, and as we take mean images on shorter timescales, all correlations at long Δt are deliberately suppressed. For this simulation, as we suppress the long-timescale behavior, the measured correlation timescale is entirely determined by the choice of T. If applied to EHT data, this timescale would be set by the time over which observations were taken, noting that we can always subdivide data to measure averages on smaller timescales.

For the rest of this chapter, we use a combination of lightcurve-normalization and an averaging timescale to focus on the short-timescale effects. We take $T = 100\,GM/c^3$ and do not consider timescales longer than this. EHT reconstructions will inherently smooth flux on scales smaller than a beam, and smoothing images has a net effect of lengthening correlation timescales. In this work, we do not include smoothing and leave its study to future work when it is applied to EHT data.

4.4.2 Rotation

Figure 4.5 shows a simulation's correlation map as a function of both time and rotation. The left panel shows the correlation when only lightcurve-normalization is used to mitigate the long-timescale effects. For $|\Delta\phi| \lesssim 15°$, the correlation timescale extends longer than any other structure and is distinct from the structure at small Δt. In order to separate this, we perform the mean subtraction procedure, thus entirely removing the long-timescale correlations. The shape of the short-timescale correlation is preserved; however the correlation value is suppressed everywhere, so we measure the timescale at $\langle \mathcal{C} \rangle = 0.2$. The timescale of the mean-subtracted correlation roughly matches an e-folding time of the correlation without mean-subtraction.

Immediately apparent in Figure 4.5 is that the correlation has a preferred direction, about 1 degree per GM/c^3. This rotation is expected, as the GRMHD movies look like rotating spirals, and is precisely the effect we wish to measure. To better measure these rotations, we first calculate the contour of $\langle \mathcal{C} \rangle = 0.2$ from an interpolated two-dimensional mea-

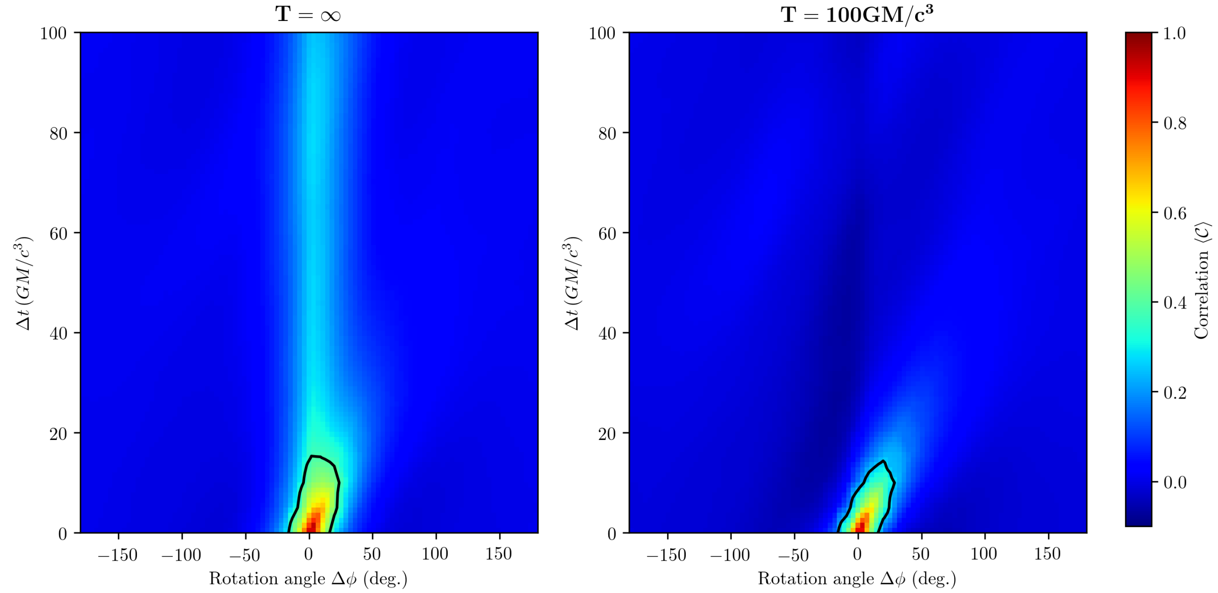

Figure 4.5: Correlation as a function of time and angle. The left panel shows an infinite averaging time, while the right has $T = 100\,GM/c^3$, thus suppressing long-timescale correlations. As a result, while the left contour is $\langle \mathcal{C} \rangle = 1/e$, the matching right contour is about $\langle \mathcal{C} \rangle = 0.2$. The correlation maps have been interpolated from the original $5\,GM/c^3$ temporal spacing for visual clarity and for smoother contours.

surement. Then, we define the maximum correlation timescale as the correlation timescale at the angle where it reaches its maximum. Lastly, we define the rotation rate as $\Delta\phi/\Delta t$ to this location. We also define an uncertainty in the rotation rate as the collection of rotation rates that have a correlation timescale of 95% of the maximum timescale. In practice, this uncertainty will be set by the ability to reconstruct images from EHT data.

4.5 Application to Many Simulations

Figure 4.6 shows the contours of $\langle \mathcal{C} \rangle = 0.2$ for all the simulations in model set A covering the two accretion states (MAD/SANE), 5 spin values, and 4 R_{high} values. Most simulations appear to have a preferred rotation rate, and a correlation timescale of about $30\,GM/c^2$. The flip in rotation rate is caused by the flip in inclinations for $a_* < 0$, such that the observed rotation is always in the direction of the large-scale disk. From these contours, we measure the angle of the maximum timescale and measure an error as all the rotation

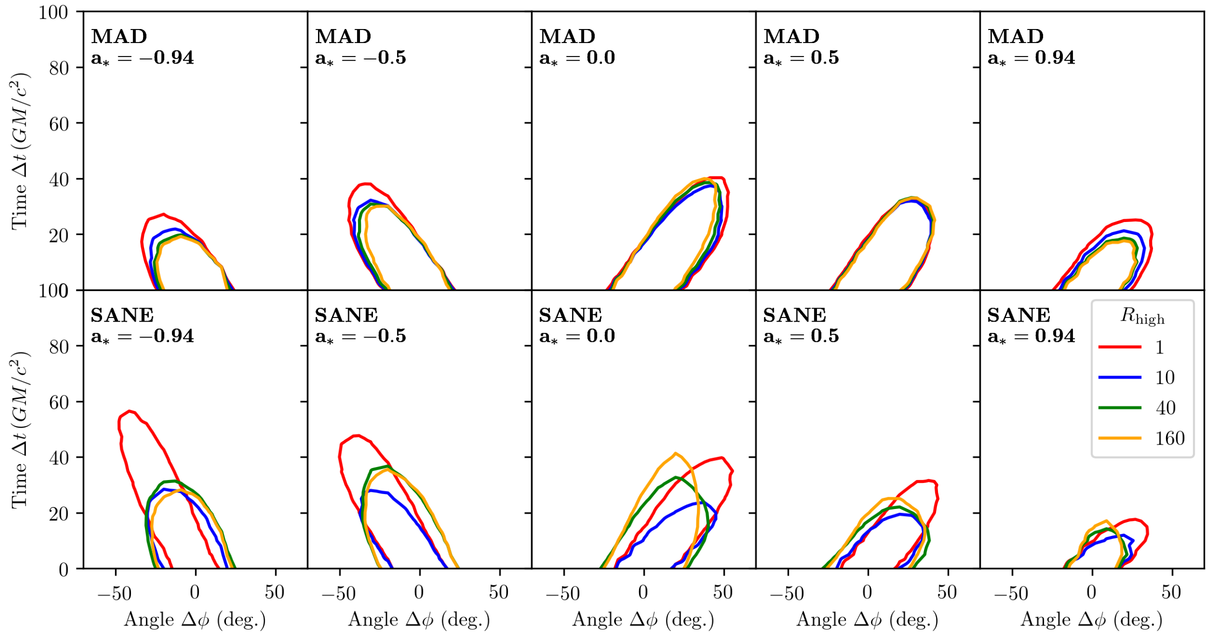

Figure 4.6: Contours of constant correlation $\langle \mathcal{C} \rangle = 0.2$ for different accretion flow types. The specifics are identical to those used for the simulation in Figure 4.5. The rotation of features is aligned with the direction of the accretion disk, even for retrograde models.

rates which give a timescale 95-100% of the maximum timescale. The primary source of uncertainty in this procedure is the low cadence of the frames. These measurements are summarized in Figure 4.7.

MAD and SANE accretion flows behave differently. All MAD models give the nearly constant rotation rate of about ±0.8 degrees per GM/c^3, independent of $R_{\rm high}$. Since the MAD models are dominated by streams of irregular accretion near the midplane, the emission tends not to depend on the specifics of the electron temperatures. There is a weak dependence of the rotation rate on the spin, possibly due to a slower azimuthal velocity at retrograde spin. The correlation timescale inversely depends on the magnitude of the spin. Since the black hole event horizon is smaller for larger spin magnitudes, this could originate from an emission region at smaller radii where timescales are shorter.

SANE models are split into three groups. For $R_{\rm high} = 1$, the emission comes from the midplane and can extend significantly further out than the photon ring, particularly for retrograde spin. This is seen in the correlation timescales, which seem to be related to the typical radius of emission and decrease with a_*. For $R_{\rm high} = 10$, the emission is still

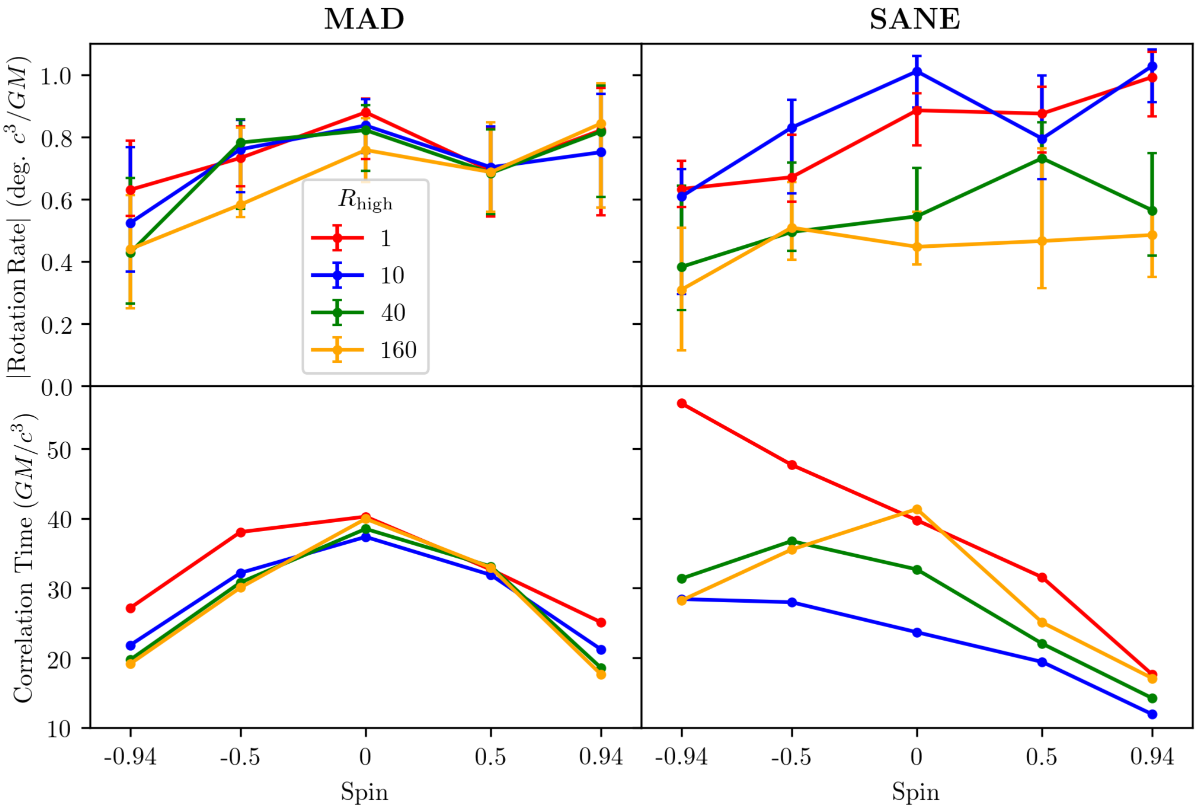

Figure 4.7: Synthesis of the rotation rate and maximum correlation timescale for the simulations in Figure 4.6. The rotation rates for $a_* < 0$ have been flipped to better show the trend with spin.

dominated by the midplane, but now a significant portion of the emission originates closer to the black hole. This gas is evolving faster, thus the correlation timescale decreases. The gas at higher values of $R_{\rm high}$ originates from regions closer to the jet. These results indicate that the gas features in the corona/jet region that are responsible for the main image emission rotate slower than features in the midplane.

4.5.1 Origin of the Rotation Rate

Most GRMHD simulations predict a rotation rate of about 0.8° per GM/c^3. A natural assumption is that the spiral features rotating in the image are caused by spiral features

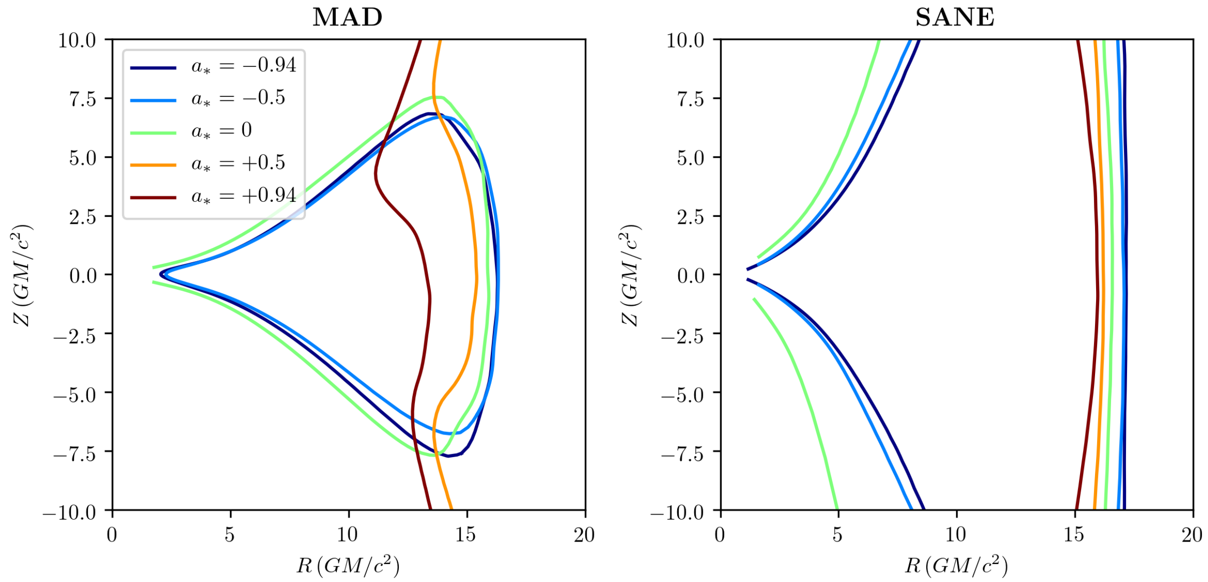

Figure 4.8: Contours of $u^\phi/u^t = 0.8^\circ/(GM/c^3)$ for the simulations used. In the midplane, this rotation rate occurs at a radius of about $15\,GM/c^3$, far away from the emission region. Retrograde and zero spin models can reach this rate close to the black hole but above the midplane.

in the accretion flow rotating at the same angular velocity, $\Omega = u^\phi/u^t \approx 0.8^\circ\, c^3/GM$. In the midplane, the GRMHD simulations are roughly Keplerian, with

$$\Omega \approx \Omega_K = \frac{GM}{r^{3/2} - a}. \tag{4.5}$$

Equating these, the rotation rate in the images would need to come from $r \approx 17\,GM/c^2$, well outside of the emission region. Taking a more rigorous approach, Figure 4.8 shows the contours of this rotation rate, that is, a map of where emission would need to come from so that the rotation rate in the images matches the rotational fluid velocity. For $a_* = 0.5$ and $a_* = 0.94$, this only happens far away from the black hole. For zero and retrograde spin, however, it is possible for the emission to originate from slightly above the midplane and thus match the image rotation rate with Ω.

In this section, we argue that the rotational rates observed in GRMHD movies correspond to features moving in the accretion flow at a different speed than the fluid velocity, akin to a backward traveling spiral wave. First, we rule out a lensing effect, by testing the effects of the fastlight approximation and by a rotating hot spot. Then, we study the effects of including emission only from a particular latitude to show that each latitude contributes to the total correlation map with the same rotation rate.

Slowlight vs Fastlight

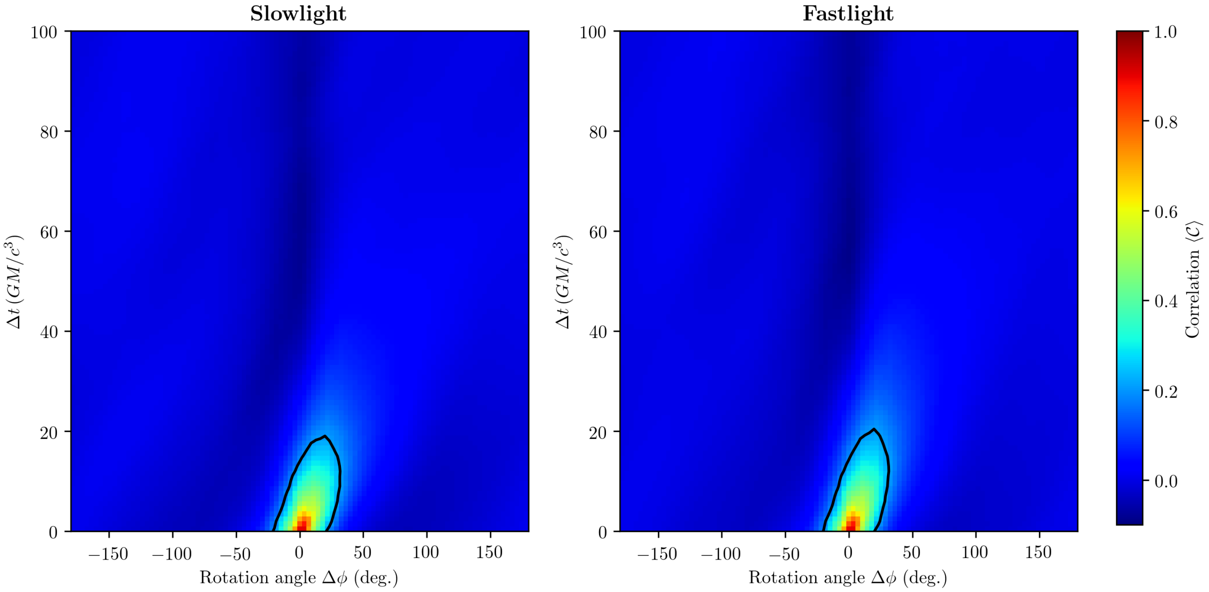

Figure 4.9: Correlation as a function of time and angle for the same simulation imaged with (right) and without (left) the fastlight approximation. The contours correspond to $\langle\mathcal{C}\rangle = 0.2$ after interpolating in the same way as in Figure 4.5. The two simulations produce similar results.

All of the images so far have used the fastlight approximation, in which the accretion flow is assumed to be constant during the imaging process. However, the velocity of the gas is comparable to the speed of light near the black hole, and it is possible this affects an observed rotation rate. To test the effects of this fastlight approximation, we use model sets B and B_s, composed of the same fluid simulation, but imaged with and without the fastlight approximation, respectively. Similarly as done in Chapter 3, we align the two simulations' time windows by maximizing the cross-correlation of their total flux lightcurves with respect to a time offset. After this, the lightcurves are very similar, and both simulations have the same time cadence and window length. Imaging is done with dumps spaced $0.5\,GM/c^3$, but only every sixth frame (i.e., a cadence of $3\,GM/c^3$) is used for the correlation analysis.

Figure 4.9 shows the time and azimuthal correlation of the two simulations. The correlation is a measure averaged over many images, and although individual snapshots of the two simulations are visually distinct, $\langle\mathcal{C}\rangle$ appears similar. Quantitatively, a typical correlation measurement of the two simulations differs by $5-10\%$, slightly above the statistical noise from using ~ 1000 images. Certainly, there is not enough of a difference to explain the image features rotating slower than the gas rotation.

Hotspot

An image of a black hole accretion flow mixes the emission from many locations of gas, so it is possible that a feature moving in the accretion flow may not appear to move the same way after being lensed. For this test, we use the RIAF+hotspot model in Tiede et al. (2020).

The constant RIAF background is from Broderick et al. (2016) with the parameters slightly changed to match the total flux of M87*, and thus the approximate amount of optical depth. On top of this diffuse background, there is an overdensity of electrons placed in the midplane at $r = 8.5\,GM/c^2$. The center of this spot has

$$u^r = u^r_K + \alpha(u^r_{\text{ff}} - u^r_K) \tag{4.6}$$

and

$$\Omega = \Omega_K + (1-\kappa)(\Omega_{\text{ff}} - \Omega_K) \tag{4.7}$$

where the K subscript refers to Keplerian motion, and ff refers to freefall. For this spot, we set $\kappa = 1$ and $\alpha = 0.005$, so that the spot has a fixed rotation rate as it falls into the black hole. Since the spot has a finite size, it shears and diffuses and does not keep its original shape. The model is imaged at an inclination of 18.2°.

Figure 4.10 shows images of this hotspot model and the correlation map. Also shown is the rotation rate of the spot at its initial location of $r = 8.5\,GM/c^2$. There is a prominent rotation rate that matches the Keplerian rate, and thus the approximate rotation of the spot in space. Thus, we conclude that for a face-on black hole such as M87*, an orbiting feature in the gas creates a similarly orbiting feature in the images with the same rotation rate.

Elevation

The contours of u^ϕ/u^t in Figure 4.8 suggest that, for zero or retrograde spin, a slow rotation rate could be present in the gas close to the black hole, but above the midplane.

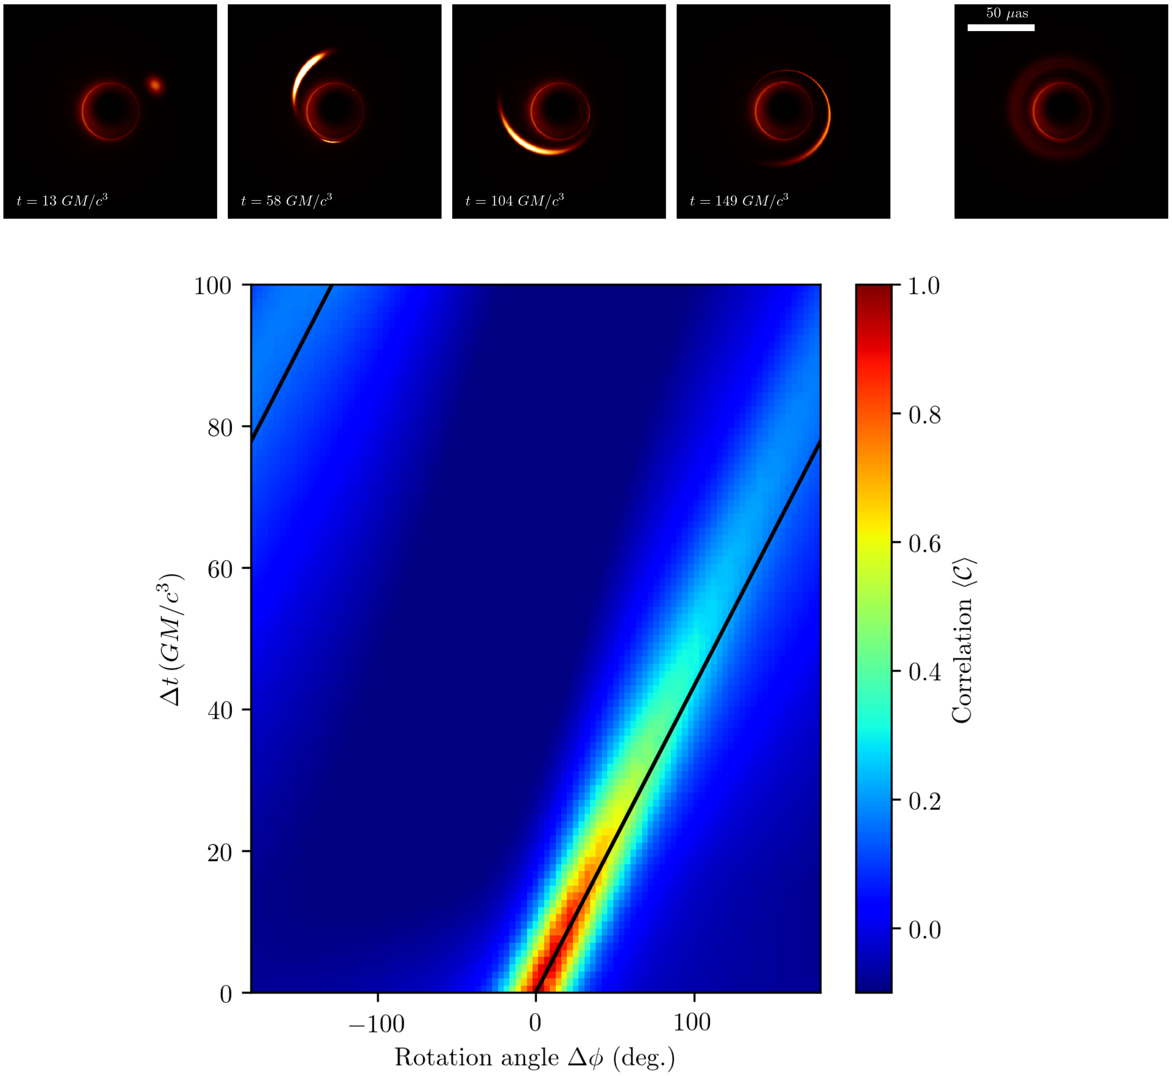

Figure 4.10: Example frames and correlation map of an orbiting hotspot model. The top panels show four sequential frames (left) and the average image (right). The bottom panel shows the same quantity as in Figure 4.5. The solid black line is the fluid rotational rate of the Keplerian spot at a radius of $8.5\,GM/c^2$, and matches the observed rotation in the images.

To test this, we decompose images of a simulation into elevation slices labeled as model set A_e. Each image of each slice corresponds to emission only taken in some range of z.

From the image sets of each slice, we independently measure the correlations, although they are all centered on the same position. The simulation used is MAD, with $R_{\mathrm{high}} = 40$ and $a_* = -0.5$ and imaged at an inclination of $17°$. The spacings in z are every $0.5\,GM/c^3$ from $-9.5\,GM/c^3$ to $9.5\,GM/c^3$, and we take 1000 images from the time window $25,000$-$30,000\,GM/c^2$. With this inclination, positive z corresponds to emission coming from nearer to the observer than negative z.

Figure 4.11 shows example images and the correlation contours in the same style as Section 4.5 for the slices near the midplane. The bulk of the emission comes from near the midplane as that is where the gas density is highest. The images constructed as a sum of elevation slices are similar to the images created using all the gas, implying that optical depth effects are minor. The correlation contours similarly match between the two simulations (solid black and grey dashed lines).

The different elevation slices have roughly the same rotation rate, at least for the dominating emission near the midplane. The correlation timescale increases with the elevation away from the observer. However, the relative contribution to the emission decreases, so the correlation of the full simulation is some combination of bright midplane emission, and prominently correlated emission away from the midplane. Each elevation slice contributes with the same rotation rate.

Figure 4.11 also shows the average fluid rotational velocity at $r = 4\,GM/c^3$ as rays emanating from $(\Delta t, \Delta\phi) = (0, 0)$. These do not equal the rotation seen in the images and are sometimes in the opposite direction. To make them match, the emitting radius would need to be further out for higher elevations. For $z = 2\,GM/c^3$ above the midplane, the gas rotates in the same direction as the black hole (opposite the rotation seen in images) all the way to $r \sim 7\,GM/c^2$. Therefore, it is unlikely that the coherent fluctuations are so far away from the black hole and rotate with the average fluid velocity.

It is noteworthy that the fluid velocity often varies by 100%, and so it is possible that the rotating features only appear when the value of Ω briefly lowers in magnitude. It is also possible that the features are not moving gas, but rather a spiral wave backward-propagating with about half the fluid velocity and vertical structure such that there is not much change in correlation rotation rate with elevation. We leave an exact connection for future work and note that the vertical structure of the fluid velocity is insufficient to explain the discrepancy between the average fluid velocity and the observed correlation rotation rate.

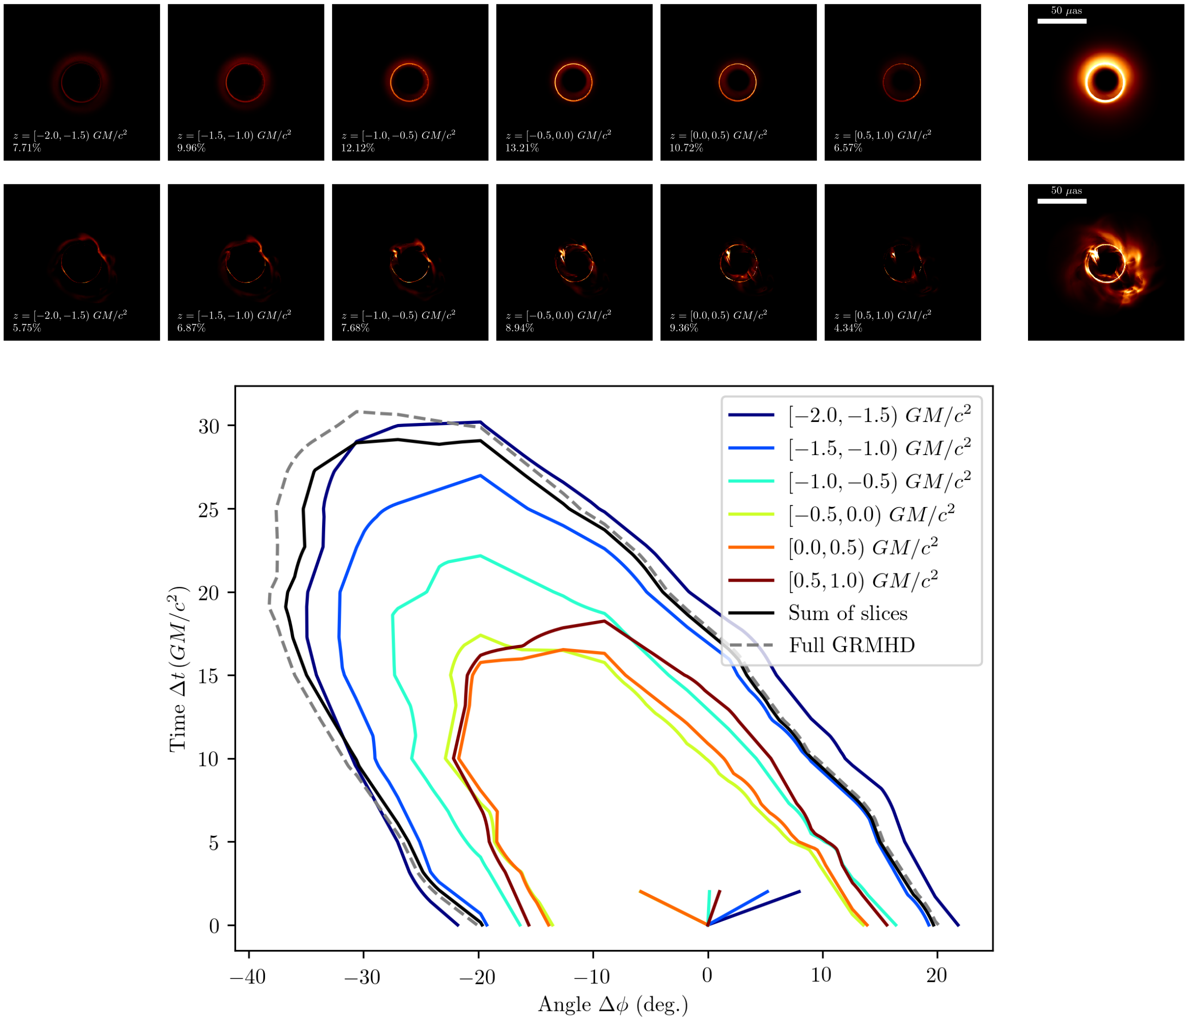

Figure 4.11: Example frames and correlation contours of a simulation sliced along elevation. The top row of panels shows the time-averaged images for each slice (left) and the sum over slices (right). The second row of panels shows a single image at $5,000\,GM/c^2$ from each slice (left) and the combined image (right). The percentages measure the contributions of a slice to the total image flux. The bottom plot shows the contours of correlation for each slice, the sum over the slices, and the simulation including all gas in the imaging process. The lines near (0, 0) show the average Ω of the gas in the corresponding slice at $r = 4\,GM/c^2$.

4.6 Image Dipole Moments

The results of the correlation analysis clearly indicate that the dominant variability is expected to be rotating features. In fact, EHT observations of M87* see a general morphology of a bright ring with some asymmetry that changes on a timescale of a week (M87* Paper I; see Figure 1.2). In this section, we turn this into a possible constraint on accretion flows by comparing the dipole moments of simulations and data.

4.6.1 Methodology

For both the simulations and the data, we define the dipole moment angle as

$$\mathcal{D}(t) = \arctan\left[\frac{\iint dxdyI(x,y,t)\sin\phi}{\iint dxdyI(x,y,t)}, \frac{\iint dxdyI(x,y,t)\cos\phi}{\iint dxdyI(x,y,t)}\right], \tag{4.8}$$

where

$$\sin\phi = \frac{x}{\sqrt{x^2+y^2}} \quad \text{and} \quad \cos\phi = \frac{y}{\sqrt{x^2+y^2}}. \tag{4.9}$$

The integrals are evaluated as a sum over pixels, and the GRMHD images are centered before this calculation as described in Subsection 4.3.4.

For EHT data, we take the results of Broderick et al. (2022a). These fit 2017 EHT observations of M87* with a Bayesian image model, and thus the resulting image samples represent a posterior distribution. For each image on each day, we center as before and calculate a dipole moment. There is no guarantee that the angle measured on reconstructed images corresponds to the same angle on a "truth" image; however, we note that this Bayesian imaging model can recover brightness asymmetries reasonably well (Tiede, Paul, 2021). In particular, the dipole moment measurement is sensitive to a convolution with a beam size, since the center of the images changes when the standard deviation of the convolution kernel becomes larger than $5\,\mu$as. We leave a more rigorous validation and calibration of the measured position angle to future work and include EHT data merely as indicative.

As we are primarily concerned with variability, we will define as our observable the change in position angle of the dipole moment over a period of time. Published EHT data of M87* has observations separated by 1 day ($3\,GM/c^3$) and by 5 days ($13\,GM/c^3$). In addition, there are ongoing observations spaced by approximately 1 year ($1,000\,GM/c^3$).

4.6.2 Sampling Rate and Slowlight

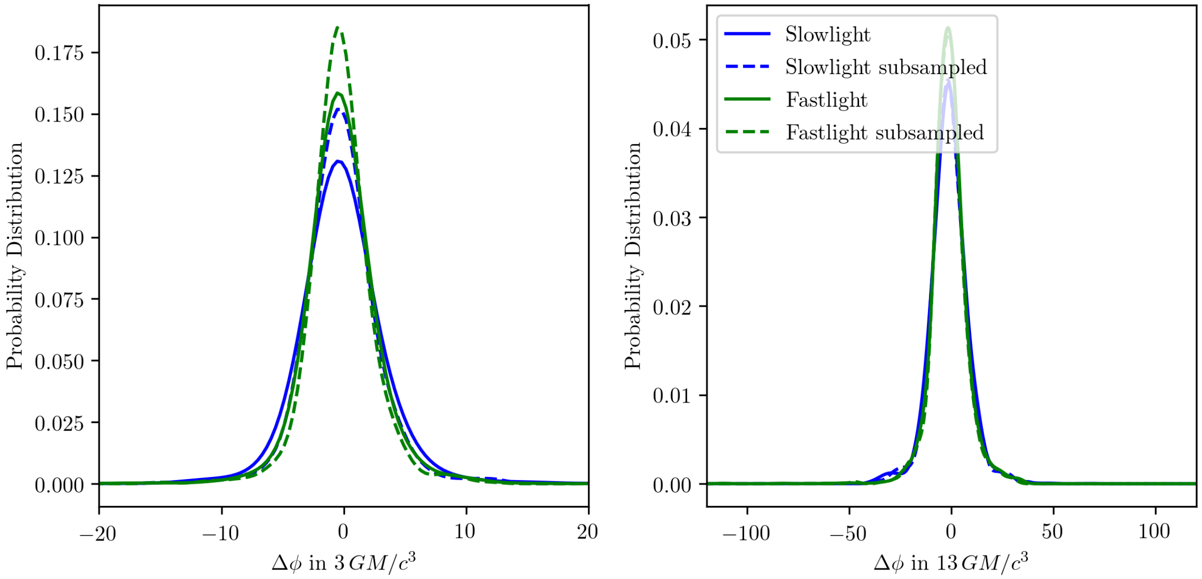

Figure 4.12: Change in the angle of the dipole moment in $3\,GM/c^3$ (left) and $13\,GM/c^3$ (right) for the same simulation imaged with slowlight (blue) and fastlight (green). Also shown is the result when each simulation only has each 10th frame, or $5\,GM/c^3$ temporal spacing (dashed).

As before, there is a worry that the fastlight approximation may impact the variability. Here we perform a test whereby we calculate the dipole moment for both the slowlight and fastlight simulations. Additionally, the simulation cadence does not match the observation cadence, so we linearly interpolate the dipole angle to match. Since the slowlight simulation has a small cadence of $0.5\,GM/c^3$, we can test whether the interpolation changes the recovered amount of variability. Figure 4.12 shows the predicted change in the dipole moment for the simulations in model sets B and B_s. Both the fastlight approximation and the linear interpolation have negligible effects on the measurement.

4.6.3 Dipole Moments of the GRMHD Simulation Library

Figure 4.13 and Figure 4.14 show the dipole angle change in 1 day and 5 days, respectively. Over 1 day, the simulations all predict that the dipole moment shouldn't change more than

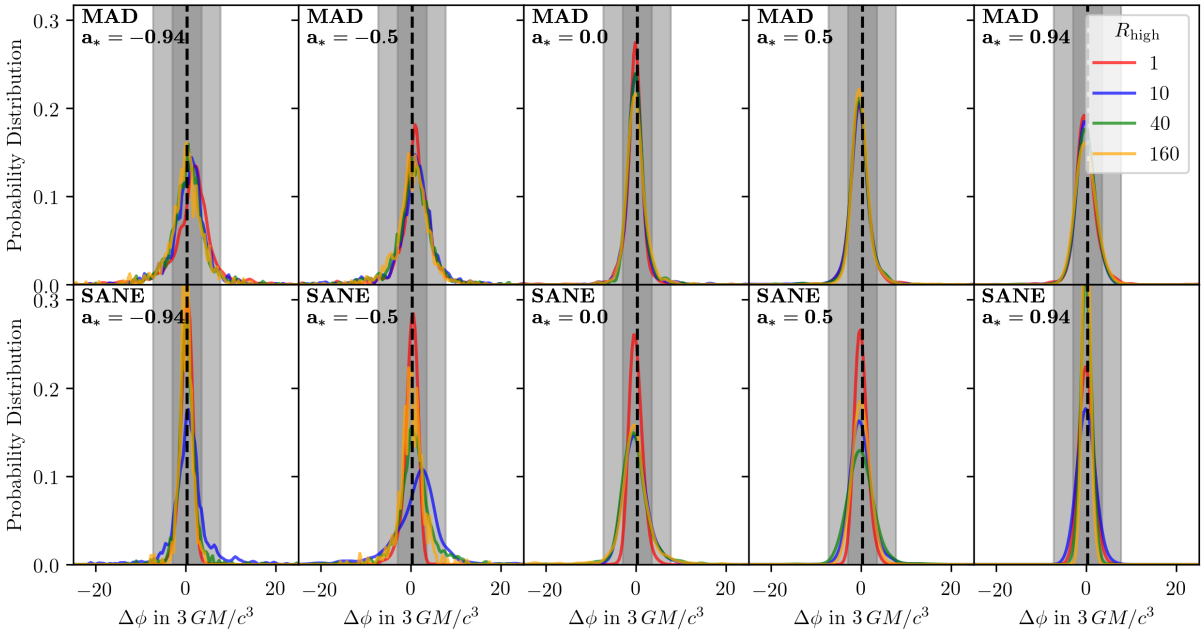

Figure 4.13: Change in the dipole moment over 1 day for all simulations in model set A. The vertical dashed line and surrounding bands are the EHT measurements between April 5 and April 6.

about 5°, regardless of the type of accretion flow. This is expected, as even the fastest moving features couldn't travel more than about $10\,\mu$as in a day. Five days, however, allows for more complex motion, while still being shorter than the maximum correlation timescales found in Section 4.5. Figure 4.14 shows that the retrograde simulations exhibit larger coherent variability, especially for MAD models. The EHT data prefers a clockwise change in the position angle from the 2017 observations. The amount is consistent with the prediction from the simulations.

Figure 4.15 shows the predicted change in the dipole moment over a year. Retrograde models allow a broader range of possibilities, typically ±90°, as these simulations have messier emission structures. The SANE, $a_* = 0.94$, $R_{\text{high}} > 40$ models are dominated by emission in the jet and corona region, and evidently, this region does not produce much variability. Thus, a one-year measurement of the change in position angle can inform on the regions where the emission originates.

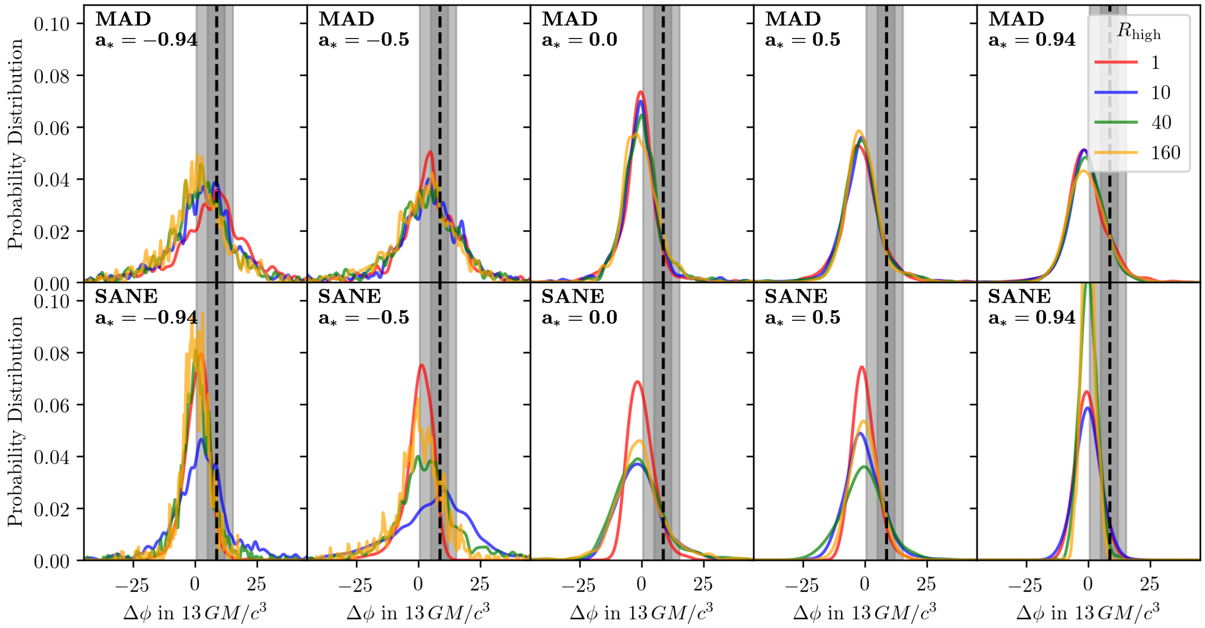

Figure 4.14: Same as Figure 4.13, but for 5 day spacing. EHT data is taken between April 6 and April 11.

4.7 Summary and Conclusions

In this chapter, we have presented a set of metrics to measure coherent variability and applied them to a library of GRMHD simulations of black hole accretion flows suited for M87*. These simulations covered the two accretion flow types (MAD/SANE), 5 spin values, and 4 values of the $R_{\rm high}$ electron temperature prescription. For each simulation, we construct a map of the average correlation for images spaced apart in time and rotation angle.

The first metric we present is the maximum correlation timescale, which measures how long fluctuations in simulations last until they diffuse out. This timescale is between 20-50 GM/c^3, providing a baseline for analyses that use one simulation to generate many independent snapshots (see, e.g., M87* Paper V). We find that MAD models have a longer correlation timescale for non-spinning black holes, and it increases with spin magnitude. SANE models have a more complex behavior which roughly depends on the location of the emitting region. Generally, the further away the emission region is from the black hole, the longer the correlation timescale.

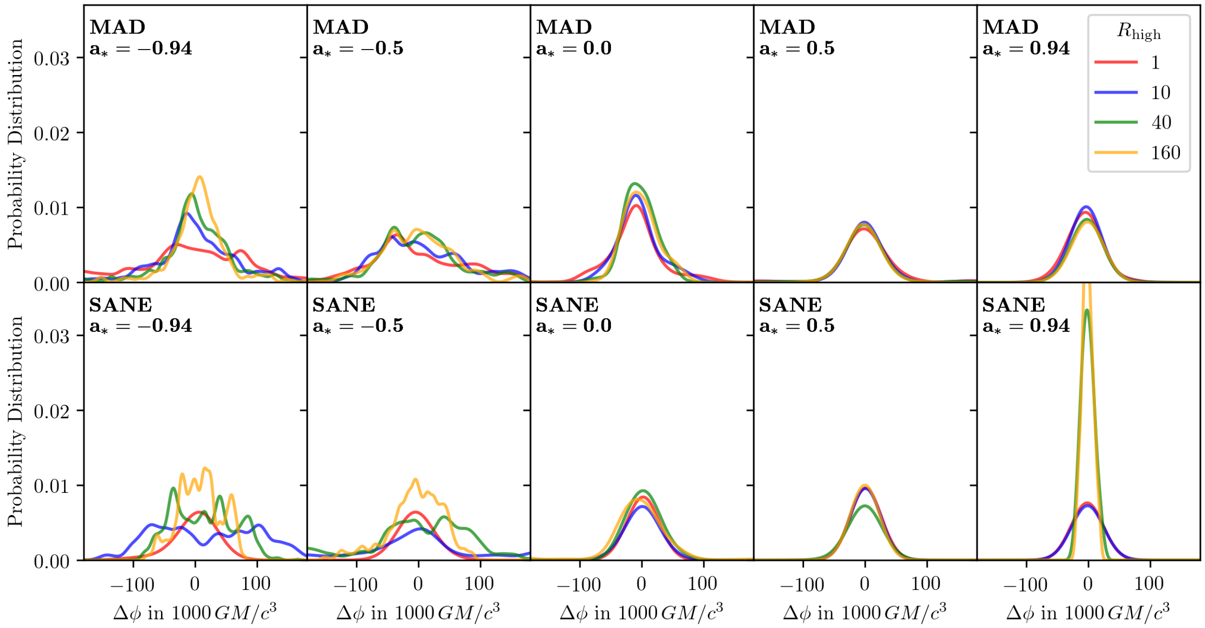

Figure 4.15: Same as Figure 4.15, but for 1 year spacing.

The second measurement was a rotation rate. The GRMHD library gives a consistent rotation rate of about $0.8^\circ\,c^3/GM$, aligned with the angular momentum of the large-scale disk. There is a weak dependence on spin and a larger dependence on the emission region through the $R_{\rm high}$ parameter. Prograde models have slightly higher rotation rates and jet/corona-dominated models have moderately higher rotation rates. The measured rotation rates do not match the average rotation of the fluid velocity in the emission region. To provide further rigor to this statement, we performed a series of tests confirming that neither lensing effects nor the vertical locations of the emitting gas are responsible for this discrepancy.

The third measurement is the position angle of the dipole moment, mimicking a bright region on a ring of emission. We measure this quantity on both simulation and EHT data and provide an interpretational comparison. Essentially all of the simulations are consistent with the observed variability in M87* over both one and five days. We also present a prediction for the maximum change of the dipole moment for the 2018 EHT observation of ±90°.

These three measurements provide a means to connect observed coherent variability to physical processes. They are not unique, and any measure of short-timescale variability which is well-constrained by data can be plugged into the procedures set up here and calculated from the GRMHD image library. It is, however, necessary to validate these methods on synthetic data to ensure that the process of reconstructing images from EHT data does not bias the measurement of interest or create variability where none exists. We leave the application to further real and synthetic EHT and ngEHT data for future work.

Chapter 5

Summary and Conclusions

With EHT observations of Sgr A* and M87* reaching precisions that challenge current models, we need methods that collect all of the physical and systematic effects in order to robustly constrain theories of plasma physics and gravity. Along with these, we need an understanding of what features in the data drive constraints, both to build confidence in results and to design better instruments. In this thesis, we contribute to these tasks in two ways. First, we study the variability present in dynamical simulations and attempt to separate the variability from the average accretion flows in order to build a more consistent variability prescription for semi-analytical models. The resulting model we build can then be directly fit to data and used to extract parameters of interest. Second, we analyze a vast library of simulated black hole images and extract measures of variability in the short- and long-timescale regimes relevant for M87* and Sgr A*, respectively. In the case of Sgr A*, this disentangling of variability is a necessity for the reconstruction of average images. For both sources, these variability measures provide a means by which we can use variability in EHT data to discriminate between accretion flow models.

5.1 On a Semi-Analytical Model of Black Hole Accretion Flows

In Chapter 2, we further the understanding of magnetohydrodynamical turbulence in black hole accretion flows by building a static midplane model and comparing it against variable GRMHD simulations. We use 10 simulations, spanning five spin values and the two accretion flow types, MAD and SANE. All the simulations converge to a quasi-steady state, indicating that there is a stable configuration within which the variability impacts the background flow.

We first measure how well the average GRMHD fluid variables satisfy their constituent equations, essentially quantifying the effects of the nonlinear fluid interactions. We find that due to numerical shortcomings, we can only robustly measure averages and fluctuations in the disk and corona regions where the density is sufficiently large. We then find that the effects of the variability are less important in the midplane, with different behaviors in MAD and SANE models. SANE models are quieter, and most of their variability is driven by the variance in the velocity and the vertical momentum flux. MAD models are much more variable, and they are driven by variability in almost every constituent equation.

Motivated by these findings, and by the simplicity of the GRMHD equations in the midplane, we construct a static axisymmetric GRMHD midplane model assuming vertical symmetry. It takes the form of 5 auxiliary variables that are constant in radius. These physically correspond to a specific entropy, energy, angular momentum, and a temporal and azimuthal magnetic flux, and, along with a given density profile, fully describe the flow. Importantly, a solution that exists at one radius may not exist everywhere. We then describe the behavior of the static accretion flow for several cases relevant to GRMHD flows. When the angular momentum is too high, the gas from far away cannot reach the black hole. When the entropy or energy is too large, the gas could not have originated from far away.

Upon incorporating the ability of one auxiliary variable to vary in time, we find that solutions converge to a piecewise solution. If a solution constructed from constant auxiliary variables would not exist at some radius, then instead the radial velocity is zero. Depending on the boundary conditions of the accretion flow, the average auxiliary variables become some combination of radial constants and their critical values that lead to a zero radial velocity.

From this behavior, we construct an approximation for variability in accretion flows as an average of two piecewise solutions. The first solution has a boundary condition at infinite radius, with approximately no radial momentum, no pressure, and no magnetic field, but some finite angular momentum. In this solution, as the gas falls in, it eventually stops and enters a region with zero radial velocity. Near the black hole, the solution reverts to a constant angular momentum, equal to the maximum that a particle can have at infinity and still reach the black hole. The second solution has a boundary condition at the event horizon and has some finite energy, entropy, angular momentum, and magnetic field. As information travels outward, this solution also reaches a zero radial velocity and stops, and this applies to all radii past a certain point. Finally, we average the two solutions to complete the approximation.

The approximation works well for SANE models while being many orders of magnitude faster to compute. Due to their relatively low variability, the radial profiles of SANE models can be well described by their median values, which contain spin-dependent breaks and radial dependencies. Our approximation takes as input the "true" density profiles and matches the auxiliary variable profiles and their spin dependencies. MAD models are more variable, and so their description requires more than a median midplane value, so our approximation fails. However, the MAD model profiles and their variability strongly indicate these models are also driven by the critical solutions with zero radial velocity.

In order to extract physical information from EHT data, three main future extensions are required. The first is a two-dimensional static model upon which we could construct our inner and outer critical solutions. It is likely that the static model would then require numerical integration and a more complex boundary condition. The second is the integration of this GRMHD approximation into existing parameter estimation frameworks such that EHT data directly informs on the metric and boundary parameters. A key necessity will be the validation that the images created from the static model correspond to averages of GRMHD movies of black holes. Lastly, the third main extension is the understanding of how the distributions of the auxiliary variables evolve in space. This is a requirement for MAD models, where variability is large, but also for interpreting the inner saturated critical solutions. We leave these extensions to future work.

5.2 On Using Variability in EHT Data to Constrain Accretion Flows

In Chapter 3 and Chapter 4, we construct two libraries of simulated images relevant for M87* and Sgr A*. These cover the two accretion states, MAD and SANE, multiple spin values, different properties of electron thermodynamics (i.e., $R_{\rm high}$ values), and observer inclination. M87* varies on timescales of an observing night, and Sgr A* varies on timescales of a minute. Combined with the observational necessities of the EHT, these sources either give a picture of moving emission structures or a map of the statistical properties of the variability. Thus, we develop two separate methods to measure this variability.

5.2.1 M87*

Measurements of the jet in M87* indicate that the black hole inclination is nearly face-on. Thus, images are composed of a bright ring and spiral arms that rotate with the flow. We focus on these spiral features and create three measurements of their coherent variability. The first is a correlation timescale measuring how long features remain before diffusing out. The second is a rotation rate of the spiral features, and the third is the position angle of the dipole moment. For the first two measurements, we first construct a correlation map, i.e., a time-average of the image-averaged correlations for frames spaced in time and position angle. We also separate out the variability associated with the mass accretion rate or with long-timescale trends not sufficiently captured by the simulations by mean-subtracting and lightcurve-normalizing.

The simulations show a clear rotation rate in their correlation maps in the same direction as the large-scale accretion flow, as well as a unique correlation timescale. MAD models have nearly the same rotation rate of $\sim 0.8^\circ c^3/GM$ regardless of spin or $R_{\rm high}$, and their correlation timescale peaks for zero spin. SANE models have more variance in the rotation rate depending on the location of the flow parameters, though still averaging to the same rotation rate as the MAD models. The correlation timescales are 10-50 GM/c^3 and their dependence on model type can be explained by a change in the emission region. Importantly, this short correlation timescale validates assumptions about the independence of snapshots in EHT analyses that have used one simulation to generate a population of turbulent realizations.

The typical rotation rate we measure in the GRMHD images is about half the amount expected from the fluid velocity. We argue that the spiral features in the images correspond to similar features in the accretion flow, yet move at a different speed than the fluid velocity (e.g., a backward-traveling wave). To build confidence in this assertion, we use a toy model composed of a static accretion flow and a shearing orbiting hot spot. The resultant correlation maps clearly show a rotation rate that matches the velocity of the orbiting spot. We also test the effects of the fastlight approximation, confirming that there are no spurious lensing effects altering rotation rate measurements. When looking only at emission from a small latitude range, we find that emission coming from every part of the simulation has the same rotation rate, inconsistent with the average fluid velocity.

EHT images of M87* support the idea that the event horizon emission can be composed of some ring-like structure and moving spots on that ring. To quantify these and connect them to the GRMHD simulations, we measure the dipole moment of all the images in the simulation library. We then choose a measurement of the change in this dipole moment over a day ($3\,GM/c^3$), a week ($13\,GM/c^3$), and a year ($1,000\,GM/c^3$), as these are the timescales we can probe with EHT data. We find that most simulations predict a similar narrow range of the change in the dipole moment, which varies more on longer timescales. Current EHT observations are consistent with the predictions of the simulations. The prediction for the 2018 EHT observations is that the dipole moment will have changed by fewer than 90 degrees.

To fully use the results of the correlation analysis, there needs to be a similar way to measure it from the data. Few observations are available, so any average image or correlation map will be subjected to a high amount of sampling noise. A key necessary test is to determine whether the correlation measure in the GRMHD images is altered after it has gone through the EHT image reconstruction process. The correlation we have looked at is image-averaged, and more information would exist in, e.g., an optical flow map that determines where in the image and in what direction fluctuations move. Furthermore, it may be possible to directly correlate along the tracks in the Fourier domain for which the EHT has data, though this increases the difficulty of interpreting motions. The dipole moment analysis can also be extended by measuring variability in parameters of more sophisticated models in both simulations and EHT observations. We leave these possible extensions for future work.

5.2.2 Sgr A*

Sgr A* is variable on timescales much shorter than current EHT observations, so a statistical measure of variability is needed. We quantify it with a lightcurve-normalized mean-subtracted PSD, which can be calculated from both GRMHD images and EHT data. We apply it to a library of simulations constructed for Sgr A*, and find that variability is greatest on the longest timescales and largest spatial scales. We also find a universal broken power-law behavior. On large scales, the power-law index is 2, and the amplitude directly measures the variance in the centroid of emission. The break location and small-scale power-law index contain information about a mix of the average image and the MHD turbulence. On shorter timescales, as more applicable to M87* or a subset of Sgr A* data, the power in the variability lessens and the break location moves to longer baselines.

In particular, we find that the variability PSDs have the most discriminatory power on short baselines, where the variability is driven by the covariance ellipse of the centroid of emission. Much of its dependencies on accretion flow type are driven by the shape of the average emission, which in turn is driven by the fluid emission regions. The variability is most sensitive to the inclination of the source, whether the flow is a MAD or a SANE, and in the case of SANE models, whether the emission comes from the disk or corona.

Motivated by the ability of EHT data to measure the variability PSDs, we compare the amplitude of the azimuthally-averaged scattered residual PSD at 4 Gλ. The simulation library predicts a narrow range near 10^{-3}.

These results are consistent with EHT observations. M87* has very little variability over a one-night observation, and a small, but detectable, amount when considering a whole week. The variability in Sgr A* is only robustly quantifiable on long baselines and averaged over azimuth, and it follows a power law. The power-law index matches the predictions from simulations, and the magnitude is on the lower half of the range at 4 Gλ. This agreement implies the GRMHD simulations are producing variability similar to the EHT data, or that our universal measurement or variability is dominated by some other quantity, such as the average emission.

Crucially, these results are used as the basis of a broken power-law variability noise model, which is required to produce average images of Sgr A*. As EHT data improves, the noise model will need to expand to cover a two-dimensional structure and timescale dependencies. We provide descriptions of both in this work and as the data more precisely measures the level of variability, it can provide discriminatory power for libraries of simulations. Since the variability is strongest on large scales, in order to best use Sgr A* to study variable accretion flows, our work points to improving the EHT instrument to have better short-baseline coverage and gain calibration.

5.3 Thesis Conclusions in an Astrophysical Context

Much of the work in this thesis directly addresses difficulties in extracting information about physical processes from EHT observations of black holes at event horizon scales. In particular, the EHT can constrain deviations from GR and make precise measurements of metric parameters when the astrophysics are known. The chief systematic uncertainty is the state of the accreting gas and its turbulent variability. We've mitigated this by constructing models and measurements which both remove variability as an uncertainty and use variability to extract information about the saturated MHD processes.

We have found that the structural variability predicted by GRMHD simulations is consistent with data for both short and long timescales. This did not have to be the case; this initial success of GRMHD in describing black hole accretion flows adds confidence to the applicability of GRMHD in other astrophysical systems.

Our work in measuring the power spectra of GRMHD images has broad implications for VLBI. The expectations of the shapes of these power spectra have been used by the EHT

Collaboration to model and reconstruct the stochastic variability and average images for a source that breaks traditional assumptions about Earth-aperture synthesis. The success in modeling Sgr A* variability opens the door to observing other rapidly varying sources with high angular resolution using VLBI.

We've also identified a manner in which the turbulence in accretion flows saturates. Importantly, it can straightforwardly be extended to non-GR spacetimes and more complex vertical structures. For SMBHs, the type of models we created can rapidly test a wide variety of GRMHD-like accretion flows. They can also describe how the relativistic outflows behave, and thus how SMBHs affect their cosmic environment's evolution. For other accreting systems, from protoplanetary disks to quasars, our work has implications about the type of turbulent angular momentum and energy transport. Even more broadly, it can be applied to many astrophysical (and terrestrial) systems where fluid turbulence takes a leading role.

References

Abbott, B. P., Abbott, R., Abbott, T. D., et al. 2016a, Phys. Rev. Lett., 116, 221101, doi: 10.1103/PhysRevLett.116.221101

—. 2016b, Phys. Rev. Lett., 116, 061102, doi: 10.1103/PhysRevLett.116.061102

Akiyama, K., Ikeda, S., Pleau, M., et al. 2017a, AJ, 153, 159, doi: 10.3847/1538-3881/aa6302

Akiyama, K., Kuramochi, K., Ikeda, S., et al. 2017b, ApJ, 838, 1, doi: 10.3847/1538-4357/aa6305

Anantua, R., Emami, R., Loeb, A., & Chael, A. 2020, ApJ, 896, 30, doi: 10.3847/1538-4357/ab9103

Arras, P., Frank, P., Haim, P., et al. 2022, Nature Astronomy, 6, 259, doi: 10.1038/s41550-021-01548-0

Balbus, S. A., & Hawley, J. F. 1991, ApJ, 376, 214, doi: 10.1086/170270

Bisnovatyi-Kogan, G. S., & Ruzmaikin, A. A. 1974, Astrophysics and Space Science, 28, 45, doi: 10.1007/BF00642237

Blandford, R. D., & Payne, D. G. 1982, MNRAS, 199, 883, doi: 10.1093/mnras/199.4.883

Blandford, R. D., & Znajek, R. L. 1977, MNRAS, 179, 433, doi: 10.1093/mnras/179.3.433

Boccardi, B., Krichbaum, T. P., Ros, E., & Zensus, J. A. 2017, A&A Rev., 25, 4, doi: 10.1007/s00159-017-0105-6

Bolton, C. T. 1972, Nature, 235, 271, doi: 10.1038/235271b0

Bondi, H. 1952, MNRAS, 112, 195, doi: 10.1093/mnras/112.2.195

Bower, G. C., Goss, W. M., Falcke, H., Backer, D. C., & Lithwick, Y. 2006, ApJ, 648, L127, doi: 10.1086/508019

Bower, G. C., Markoff, S., Brunthaler, A., et al. 2014, ApJ, 790, 1, doi: 10.1088/0004-637X/790/1/1

Bransgrove, A., Ripperda, B., & Philippov, A. 2021, Phys. Rev. Lett., 127, 055101, doi: 10.1103/PhysRevLett.127.055101

Broderick, A. E., Fish, V. L., Doeleman, S. S., & Loeb, A. 2009, ApJ, 697, 45, doi: 10.1088/0004-637X/697/1/45

Broderick, A. E., & Loeb, A. 2006a, MNRAS, 367, 905, doi: 10.1111/j.1365-2966.2006.10152.x

—. 2006b, ApJ, 636, L109, doi: 10.1086/500008

—. 2009, ApJ, 697, 1164, doi: 10.1088/0004-637X/697/2/1164

Broderick, A. E., Loeb, A., & Reid, M. J. 2011, ApJ, 735, 57, doi: 10.1088/0004-637X/735/1/57

Broderick, A. E., & Narayan, R. 2006, ApJ, 638, L21, doi: 10.1086/500930

Broderick, A. E., Pesce, D. W., Tiede, P., Pu, H.-Y., & Gold, R. 2020a, ApJ, 898, 9, doi: 10.3847/1538-4357/ab9c1f

Broderick, A. E., Fish, V. L., Johnson, M. D., et al. 2016, ApJ, 820, 137, doi: 10.3847/0004-637X/820/2/137

Broderick, A. E., Gold, R., Karami, M., et al. 2020b, ApJ, 897, 139, doi: 10.3847/1538-4357/ab91a4

Broderick, A. E., Pesce, D. W., Gold, R., et al. 2022a, ApJ, 935, 61, doi: 10.3847/1538-4357/ac7c1d

Broderick, A. E., Gold, R., Georgiev, B., et al. 2022b, ApJ, 930, L21, doi: 10.3847/2041-8213/ac6584

Burrows, A., & Vartanyan, D. 2021, Nature, 589, 29, doi: 10.1038/s41586-020-03059-w

Calderón, D., Cuadra, J., Schartmann, M., Burkert, A., & Russell, C. M. P. 2020, ApJ, 888, L2, doi: 10.3847/2041-8213/ab5e81

Cantiello, M., Blakeslee, J. P., Ferrarese, L., et al. 2018, ApJ, 856, 126, doi: 10.3847/1538-4357/aab043

Carilli, C. L., & Thyagarajan, N. 2022, ApJ, 924, 125, doi: 10.3847/1538-4357/ac3cba

Carr, B., & Kühnel, F. 2020, Annual Review of Nuclear and Particle Science, 70, 355, doi: 10.1146/annurev-nucl-050520-125911

Carter, B. 1971, Phys. Rev. Lett., 26, 331, doi: 10.1103/PhysRevLett.26.331

Castelló-Mor, N., Kaspi, S., Netzer, H., et al. 2017, MNRAS, 467, 1209, doi: 10.1093/mnras/stx153

Chael, A., Bouman, K., Johnson, M., et al. 2019, eht-imaging: v1.1.0: Imaging interferometric data with regularized maximum likelihood, v1.1.0, Zenodo, Zenodo, doi: 10.5281/zenodo.2614016

Chael, A. A., Johnson, M. D., Narayan, R., et al. 2016, ApJ, 829, 11, doi: 10.3847/0004-637X/829/1/11

Chandra, M., Foucart, F., & Gammie, C. F. 2017, ApJ, 837, 92, doi: 10.3847/1538-4357/aa5f55

Chatterjee, K., Younsi, Z., Liska, M., et al. 2020, MNRAS, 499, 362, doi: 10.1093/mnras/staa2718

Combi, L., & Siegel, D. M. 2023, ApJ, 944, 28, doi: 10.3847/1538-4357/acac29

Comisso, L., & Sironi, L. 2018, Phys. Rev. Lett., 121, 255101, doi: 10.1103/PhysRevLett.121.255101

Conroy, N. S., Baubock, M., Dhruv, V., et al. 2023, arXiv e-prints, arXiv:2304.03826, doi: 10.48550/arXiv.2304.03826

Cuadra, J., Nayakshin, S., & Martins, F. 2008, MNRAS, 383, 458, doi: 10.1111/j.1365-2966.2007.12573.x

Curtis, H. D. 1918, Publications of Lick Observatory, 13, 9

Davis, S. W., Stone, J. M., & Pessah, M. E. 2010, ApJ, 713, 52, doi: 10.1088/0004-637X/713/1/52

Dexter, J., Agol, E., & Fragile, P. C. 2009, ApJ, 703, L142, doi: 10.1088/0004-637X/703/2/L142

Dexter, J., & Fragile, P. C. 2013, MNRAS, 432, 2252, doi: 10.1093/mnras/stt583

Dexter, J., McKinney, J. C., & Agol, E. 2012, MNRAS, 421, 1517, doi: 10.1111/j.1365-2966.2012.20409.x

Dhruv, V. 2023, in prep

Di Matteo, T., Allen, S. W., Fabian, A. C., Wilson, A. S., & Young, A. J. 2003, ApJ, 582, 133, doi: 10.1086/344504

Do, T., Hees, A., Ghez, A., et al. 2019, Science, 365, 664, doi: 10.1126/science.aav8137

Doeleman, S., Akiyama, K., Blackburn, L., et al. 2019, Bulletin of the AAS, 51. https://baas.aas.org/pub/2020n3i537

Doeleman, S. S., Weintroub, J., Rogers, A. E. E., et al. 2008, Nature, 455, 78, doi: 10.1038/nature07245

EHT MWL Science Working Group, Algaba, J. C., Anczarski, J., et al. 2021, ApJ, 911, L11, doi: 10.3847/2041-8213/abef71

Einstein, A. 1915, Sitzungsberichte der Königlich Preussischen Akademie der Wissenschaften, 844

Event Horizon Telescope Collaboration, Akiyama, K., Alberdi, A., et al. 2019a, ApJ, 875, L1, doi: 10.3847/2041-8213/ab0ec7

—. 2019b, ApJ, 875, L2, doi: 10.3847/2041-8213/ab0c96

—. 2019c, ApJ, 875, L4, doi: 10.3847/2041-8213/ab0e85

—. 2019d, ApJ, 875, L3, doi: 10.3847/2041-8213/ab0c57

—. 2019e, ApJ, 875, L5, doi: 10.3847/2041-8213/ab0f43

—. 2019f, ApJ, 875, L6, doi: 10.3847/2041-8213/ab1141

—. 2022a, ApJ, 930, L13, doi: 10.3847/2041-8213/ac6675

—. 2022b, ApJ, 930, L12, doi: 10.3847/2041-8213/ac6674

—. 2022c, ApJ, 930, L14, doi: 10.3847/2041-8213/ac6429

—. 2022d, ApJ, 930, L15, doi: 10.3847/2041-8213/ac6736

—. 2022e, ApJ, 930, L16, doi: 10.3847/2041-8213/ac6672

—. 2022f, ApJ, 930, L17, doi: 10.3847/2041-8213/ac6756

Falcke, H., & Markoff, S. 2000, A&A, 362, 113, doi: 10.48550/arXiv.astro-ph/0102186

Falcke, H., Melia, F., & Agol, E. 2000, ApJ, 528, L13, doi: 10.1086/312423

Fishbone, L. G., & Moncrief, V. 1976, ApJ, 207, 962, doi: 10.1086/154565

Foucart, F., Chandra, M., Gammie, C. F., & Quataert, E. 2016, MNRAS, 456, 1332, doi: 10.1093/mnras/stv2687

Foucart, F., Chandra, M., Gammie, C. F., Quataert, E., & Tchekhovskoy, A. 2017, MNRAS, 470, 2240, doi: 10.1093/mnras/stx1368

Foucault, L. 1851, Annalen der Physik, 158, 458, doi: 10.1002/andp.18511580315

Frank, J., King, A., & Raine, D. J. 2002, Accretion Power in Astrophysics: Third Edition

Gammie, C. F., McKinney, J. C., & Tóth, G. 2003, ApJ, 589, 444, doi: 10.1086/374594

Genzel, R., Eckart, A., Ott, T., & Eisenhauer, F. 1997, MNRAS, 291, 219, doi: 10.1093/mnras/291.1.219

Genzel, R., Eisenhauer, F., & Gillessen, S. 2010, Reviews of Modern Physics, 82, 3121, doi: 10.1103/RevModPhys.82.3121

Georgiev, B., Pesce, D. W., Broderick, A. E., et al. 2022, ApJ, 930, L20, doi: 10.3847/2041-8213/ac65eb

Ghez, A. M., Klein, B. L., Morris, M., & Becklin, E. E. 1998, ApJ, 509, 678, doi: 10.1086/306528

Goddi, C., Martí-Vidal, I., Messias, H., et al. 2019, PASP, 131, 075003, doi: 10.1088/1538-3873/ab136a

Gold, R., Broderick, A. E., Younsi, Z., et al. 2020, ApJ, 897, 148, doi: 10.3847/1538-4357/ab96c6

Gómez, J. L., Lobanov, A. P., Bruni, G., et al. 2016, ApJ, 817, 96, doi: 10.3847/0004-637X/817/2/96

GRAVITY Collaboration, Abuter, R., Amorim, A., et al. 2018, A&A, 618, L10, doi: 10.1051/0004-6361/201834294

Gravity Collaboration, Abuter, R., Amorim, A., et al. 2019, A&A, 625, L10, doi: 10.1051/0004-6361/201935656

—. 2020, A&A, 638, A2, doi: 10.1051/0004-6361/202037717

Greisen, E. W. 2003, in Astrophysics and Space Science Library, Vol. 285, Information Handling in Astronomy - Historical Vistas, ed. A. Heck, 109, doi: 10.1007/0-306-48080-8_7

Guo, F., Li, H., Daughton, W., & Liu, Y.-H. 2014, Phys. Rev. Lett., 113, 155005, doi: 10.1103/PhysRevLett.113.155005

Hada, K., Kino, M., Doi, A., et al. 2016, ApJ, 817, 131, doi: 10.3847/0004-637X/817/2/131

Haggard, D., Nynka, M., Mon, B., et al. 2019, ApJ, 886, 96, doi: 10.3847/1538-4357/ab4a7f

Harris, C. R., Millman, K. J., van der Walt, S. J., et al. 2020, Nature, 585, 357, doi: 10.1038/s41586-020-2649-2

Hawley, J. F., Gammie, C. F., & Balbus, S. A. 1995, ApJ, 440, 742, doi: 10.1086/175311

Igumenshchev, I. V., Narayan, R., & Abramowicz, M. A. 2003, The Astrophysical Journal, 592, 1042, doi: 10.1086/375769

Israel, W. 1968, Communications in Mathematical Physics, 8, 245, doi: 10.1007/BF01645859

Issaoun, S., Johnson, M. D., Blackburn, L., et al. 2019, ApJ, 871, 30, doi: 10.3847/1538-4357/aaf732

Issaoun, S., Wielgus, M., Jorstad, S., et al. 2022, ApJ, 934, 145, doi: 10.3847/1538-4357/ac7a40

Janssen, M., Falcke, H., Kadler, M., et al. 2021, Nature Astronomy, 5, 1017, doi: 10.1038/s41550-021-01417-w

Jeter, B., & Broderick, A. E. 2021, ApJ, 908, 139, doi: 10.3847/1538-4357/abda3d

Jeter, B., Broderick, A. E., & Gold, R. 2020, MNRAS, 493, 5606, doi: 10.1093/mnras/staa679

Johnson, M., Haworth, K., Pesce, D. W., et al. 2019, in Bulletin of the American Astronomical Society, Vol. 51, 235. https://arxiv.org/abs/1909.01405

Johnson, M. D., Kovalev, Y. Y., Lisakov, M. M., et al. 2021, ApJ, 922, L28, doi: 10.3847/2041-8213/ac3917

Johnson, M. D., Narayan, R., Psaltis, D., et al. 2018, ApJ, 865, 104, doi: 10.3847/1538-4357/aadcff

Jones, E., Oliphant, T., Peterson, P., et al. 2001–, SciPy: Open source scientific tools for Python. http://www.scipy.org/

Jorstad, S., Wielgus, M., Lico, R., et al. 2023, ApJ, 943, 170, doi: 10.3847/1538-4357/acaea8

Kardashev, N. S., Alakoz, A. V., Andrianov, A. S., et al. 2017, Solar System Research, 51, 535, doi: 10.1134/S0038094617070085

Kerr, R. P. 1963, Phys. Rev. Lett., 11, 237, doi: 10.1103/PhysRevLett.11.237

Kim, J.-Y., Krichbaum, T. P., Broderick, A. E., et al. 2020, A&A, 640, A69, doi: 10.1051/0004-6361/202037493

Kokubo, M. 2018, PASJ, 70, 97, doi: 10.1093/pasj/psy096

Kormendy, J., & Ho, L. C. 2013, ARA&A, 51, 511, doi: 10.1146/annurev-astro-082708-101811

Kunz, M. W., Schekochihin, A. A., & Stone, J. M. 2014, Phys. Rev. Lett., 112, 205003, doi: 10.1103/PhysRevLett.112.205003

Lee, D., & Gammie, C. F. 2021, ApJ, 906, 39, doi: 10.3847/1538-4357/abc8f3

Levis, A., Srinivasan, P. P., Chael, A. A., Ng, R., & Bouman, K. L. 2022, arXiv e-prints, arXiv:2204.03715, doi: 10.48550/arXiv.2204.03715

Li, L.-X., Zimmerman, E. R., Narayan, R., & McClintock, J. E. 2005, ApJS, 157, 335, doi: 10.1086/428089

Liepold, E. R., Ma, C.-P., & Walsh, J. L. 2023, ApJ, 945, L35, doi: 10.3847/2041-8213/acbbcf

Lockhart, W., & Gralla, S. E. 2022, MNRAS, 509, 3643, doi: 10.1093/mnras/stab3204

Lu, R.-S., Asada, K., Krichbaum, T. P., et al. 2023, Nature, 616, 686, doi: 10.1038/s41586-023-05843-w

Luminet, J. P. 1979, A&A, 75, 228

Lynden-Bell, D. 1969, Nature, 223, 690, doi: 10.1038/223690a0

Mahadevan, R., & Quataert, E. 1997, ApJ, 490, 605, doi: 10.1086/304908

McKinney, J. C., Tchekhovskoy, A., & Blandford, R. D. 2012, MNRAS, 423, 3083, doi: 10.1111/j.1365-2966.2012.21074.x

McNamara, B. R., & Nulsen, P. E. J. 2012, New Journal of Physics, 14, 055023, doi: 10.1088/1367-2630/14/5/055023

Medeiros, L., Psaltis, D., Lauer, T. R., & Özel, F. 2023, ApJ, 947, L7, doi: 10.3847/2041-8213/acc32d

Merritt, D., Berczik, P., & Laun, F. 2007, AJ, 133, 553, doi: 10.1086/510294

Montero-Castaño, M., Herrnstein, R. M., & Ho, P. T. P. 2009, ApJ, 695, 1477, doi: 10.1088/0004-637X/695/2/1477

Moran, J. M. 2008, in Astronomical Society of the Pacific Conference Series, Vol. 395, Frontiers of Astrophysics: A Celebration of NRAO's 50th Anniversary, ed. A. H. Bridle, J. J. Condon, & G. C. Hunt, 87, doi: 10.48550/arXiv.0804.1063

Moriyama, K., Mineshige, S., Honma, M., & Akiyama, K. 2019, ApJ, 887, 227, doi: 10.3847/1538-4357/ab505b

Mościbrodzka, M., Falcke, H., & Shiokawa, H. 2016, A&A, 586, A38, doi: 10.1051/0004-6361/201526630

Mościbrodzka, M., Falcke, H., & Shiokawa, H. 2016, Astronomy and Astrophysics, 586, A38, doi: 10.1051/0004-6361/201526630

Mościbrodzka, M., Falcke, H., Shiokawa, H., & Gammie, C. F. 2014, A&A, 570, A7, doi: 10.1051/0004-6361/201424358

Mościbrodzka, M., & Gammie, C. F. 2018, Monthly Notices of the Royal Astronomical Society, 475, 43, doi: 10.1093/mnras/stx3162

Mościbrodzka, M., Gammie, C. F., Dolence, J. C., Shiokawa, H., & Leung, P. K. 2009, ApJ, 706, 497, doi: 10.1088/0004-637X/706/1/497

Most, E. R., & Noronha, J. 2021, Phys. Rev. D, 104, 103028, doi: 10.1103/PhysRevD.104.103028

Most, E. R., Noronha, J., & Philippov, A. A. 2021, arXiv e-prints, arXiv:2111.05752. https://arxiv.org/abs/2111.05752

Mösta, P., Richers, S., Ott, C. D., et al. 2014, ApJ, 785, L29, doi: 10.1088/2041-8205/785/2/L29

Mościbrodzka, M., & Gammie, C. F. ????, ipole: Polarized covariant radiative transfer in C, https://github.com/AFD-Illinois/ipole, GitHub

Müller, H., & Lobanov, A. P. 2022, A&A, 666, A137, doi: 10.1051/0004-6361/202243244

Narayan, R., Chael, A., Chatterjee, K., Ricarte, A., & Curd, B. 2021, arXiv e-prints, arXiv:2108.12380. https://arxiv.org/abs/2108.12380

Narayan, R., Igumenshchev, I. V., & Abramowicz, M. A. 2003, Publications of the Astronomical Society of Japan, 55, L69, doi: 10.1093/pasj/55.6.L69

Narayan, R., Johnson, M. D., & Gammie, C. F. 2019, ApJ, 885, L33, doi: 10.3847/2041-8213/ab518c

Narayan, R., Sädowski, A., Penna, R. F., & Kulkarni, A. K. 2012, Monthly Notices of the Royal Astronomical Society, 426, 3241, doi: 10.1111/j.1365-2966.2012.22002.x

Narayan, R., & Yi, I. 1995, ApJ, 452, 710, doi: 10.1086/176343

Nättilä, J., & Beloborodov, A. M. 2021, ApJ, 921, 87, doi: 10.3847/1538-4357/ac1c76

Netzer, H. 2015, ARA&A, 53, 365, doi: 10.1146/annurev-astro-082214-122302

Novikov, I. D., & Thorne, K. S. 1973, in Black Holes (Les Astres Occlus), 343–450

Olivares, H., Porth, O., Davelaar, J., et al. 2019, A&A, 629, A61, doi: 10.1051/0004-6361/201935559

Osorno, J., Nagar, N., Richtler, T., et al. 2023, arXiv e-prints, arXiv:2304.11264, doi: 10.48550/arXiv.2304.11264

Palenzuela, C., Lehner, L., Liebling, S. L., et al. 2013, Phys. Rev. D, 88, 043011, doi: 10.1103/PhysRevD.88.043011

Palumbo, D. C. M., Gelles, Z., Tiede, P., et al. 2022, ApJ, 939, 107, doi: 10.3847/1538-4357/ac9ab7

Paschalidis, V., Ruiz, M., & Shapiro, S. L. 2015, ApJ, 806, L14, doi: 10.1088/2041-8205/806/1/L14

Penrose, R., & Floyd, R. M. 1971, Nature Physical Science, 229, 177, doi: 10.1038/physci229177a0

Pesce, D. W. 2021, AJ, 161, 178, doi: 10.3847/1538-3881/abe3f8

Pesce, D. W., Seth, A. C., Greene, J. E., et al. 2021a, ApJ, 909, 141, doi: 10.3847/1538-4357/abde3d

Pesce, D. W., Palumbo, D. C. M., Narayan, R., et al. 2021b, ApJ, 923, 260, doi: 10.3847/1538-4357/ac2eb5

Porth, O., Komissarov, S. S., & Keppens, R. 2014, MNRAS, 438, 278, doi: 10.1093/mnras/stt2176

Porth, O., Olivares, H., Mizuno, Y., et al. 2017, Computational Astrophysics and Cosmology, 4, 1, doi: 10.1186/s40668-017-0020-2

Porth, O., Chatterjee, K., Narayan, R., et al. 2019, ApJS, 243, 26, doi: 10.3847/1538-4365/ab29fd

Prather, B. 2023, in prep

Prather, B. S., Wong, G. N., Dhruv, V., et al. 2021, Journal of Open Source Software, 6, 3336, doi: 10.21105/joss.03336

Pretorius, F. 2005, Phys. Rev. Lett., 95, 121101, doi: 10.1103/PhysRevLett.95.121101

Pu, H.-Y., & Broderick, A. E. 2018, ApJ, 863, 148, doi: 10.3847/1538-4357/aad086

Pu, H.-Y., & Takahashi, M. 2020, ApJ, 892, 37, doi: 10.3847/1538-4357/ab77ab

Quataert, E., Narayan, R., & Reid, M. J. 1999, ApJ, 517, L101, doi: 10.1086/312035

Ramakrishnan, V., Nagar, N., Arratia, V., et al. 2023, Galaxies, 11, 15, doi: 10.3390/galaxies11010015

Raymond, A. W., Palumbo, D., Paine, S. N., et al. 2021, ApJS, 253, 5, doi: 10.3847/1538-3881/abc3c3

Reid, M. J., & Honma, M. 2014, ARA&A, 52, 339, doi: 10.1146/annurev-astro-081913-040006

Reid, M. J., Readhead, A. C. S., Vermeulen, R. C., & Treuhaft, R. N. 1999, ApJ, 524, 816, doi: 10.1086/307855

Remillard, R. A., & McClintock, J. E. 2006, ARA&A, 44, 49, doi: 10.1146/annurev.astro.44.051905.092532

Ressler, S. M., White, C. J., Quataert, E., & Stone, J. M. 2020, ApJ, 896, L6, doi: 10.3847/2041-8213/ab9532

Ripperda, B., Bacchini, F., & Philippov, A. A. 2020, ApJ, 900, 100, doi: 10.3847/1538-4357/ababab

Ripperda, B., Liska, M., Chatterjee, K., et al. 2021, arXiv e-prints, arXiv:2109.15115. https://arxiv.org/abs/2109.15115

Sądowski, A., Narayan, R., Penna, R., & Zhu, Y. 2013, Monthly Notices of the Royal Astronomical Society, 436, 3856, doi: 10.1093/mnras/stt1881

Saurabh, & Nampalliwar, S. 2023, ApJ, 947, 43, doi: 10.3847/1538-4357/acc6d3

Schwarzschild, K. 1916, Sitzungsberichte der Königlich Preussischen Akademie der Wissenschaften, 189

Shakura, N. I., & Sunyaev, R. A. 1973, A&A, 500, 33

Shepherd, M. C. 1997, in Astronomical Society of the Pacific Conference Series, Vol. 125, Astronomical Data Analysis Software and Systems VI, ed. G. Hunt & H. Payne, 77

Shiokawa, H., Krolik, J. H., Cheng, R. M., Piran, T., & Noble, S. C. 2015, ApJ, 804, 85, doi: 10.1088/0004-637X/804/2/85

Simon, D. A., Cappellari, M., & Hartke, J. 2023, arXiv e-prints, arXiv:2303.18229, doi: 10.48550/arXiv.2303.18229

Sironi, L., & Spitkovsky, A. 2014, ApJ, 783, L21, doi: 10.1088/2041-8205/783/1/L21

Sądowski, A., Narayan, R., McKinney, J. C., & Tchekhovskoy, A. 2014, MNRAS, 439, 503, doi: 10.1093/mnras/stt2479

Sądowski, A., Narayan, R., Tchekhovskoy, A., & Zhu, Y. 2013, MNRAS, 429, 3533, doi: 10.1093/mnras/sts632

Soltan, A. 1982, MNRAS, 200, 115, doi: 10.1093/mnras/200.1.115

Stone, J. M., & Pringle, J. E. 2001, MNRAS, 322, 461, doi: 10.1046/j.1365-8711.2001.04138.x

Sun, H., Bouman, K. L., Tiede, P., et al. 2022, ApJ, 932, 99, doi: 10.3847/1538-4357/ac6be9

Tchekhovskoy, A., McKinney, J. C., & Narayan, R. 2008, MNRAS, 388, 551, doi: 10.1111/j.1365-2966.2008.13425.x

Tchekhovskoy, A., Narayan, R., & McKinney, J. C. 2011, Monthly Notices of the Royal Astronomical Society, 418, L79, doi: 10.1111/j.1745-3933.2011.01147.x

Thompson, A. R., Moran, J. M., & Swenson, Jr., G. W. 2017, Interferometry and Synthesis in Radio Astronomy, 3rd Edition, doi: 10.1007/978-3-319-44431-4

Tiede, P. 2022, Journal of Open Source Software, 7, 4457, doi: 10.21105/joss.04457

Tiede, P., Pu, H.-Y., Broderick, A. E., et al. 2020, ApJ, 892, 132, doi: 10.3847/1538-4357/ab744c

Tiede, Paul. 2021, PhD thesis. http://hdl.handle.net/10012/17199

Toro, E. F. 2009, Riemann Solvers and Numerical Methods for Fluid Dynamics (Springer Berlin Heidelberg), doi: 10.1007/b79761

Urry, C. M., & Padovani, P. 1995, PASP, 107, 803, doi: 10.1086/133630

Wagner, S. J., & Witzel, A. 1995, ARA&A, 33, 163, doi: 10.1146/annurev.aa.33.090195.001115

Wald, R. M. 1984, General Relativity (Chicago, USA: Chicago Univ. Pr.), doi: 10.7208/chicago/9780226870373.001.0001

Walker, R. C., Hardee, P. E., Davies, F., et al. 2016, Galaxies, 4, 46, doi: 10.3390/galaxies4040046

Webster, B. L., & Murdin, P. 1972, Nature, 235, 37, doi: 10.1038/235037a0

Werner, G. R., Uzdensky, D. A., Cerutti, B., Nalewajko, K., & Begelman, M. C. 2016, ApJ, 816, L8, doi: 10.3847/2041-8205/816/1/L8

White, C. J., Quataert, E., & Gammie, C. F. 2020, ApJ, 891, 63, doi: 10.3847/1538-4357/ab718e

Wielgus, M., Akiyama, K., Blackburn, L., et al. 2020, ApJ, 901, 67, doi: 10.3847/1538-4357/abac0d

Wielgus, M., Marchili, N., Martí-Vidal, I., et al. 2022, ApJ, 930, L19, doi: 10.3847/2041-8213/ac6428

Witzel, G., Eckart, A., Bremer, M., et al. 2012, ApJS, 203, 18, doi: 10.1088/0067-0049/203/2/18

Witzel, G., Martinez, G., Willner, S. P., et al. 2021, ApJ, 917, 73, doi: 10.3847/1538-4357/ac0891

Wong, G. N., Prather, B. S., Dhruv, V., et al. 2022, arXiv e-prints, arXiv:2202.11721. https://arxiv.org/abs/2202.11721

Xie, F.-G., & Yuan, F. 2012, MNRAS, 427, 1580, doi: 10.1111/j.1365-2966.2012.22030.x

Yang, H. Y. K., Ruszkowski, M., & Zweibel, E. G. 2022, Nature Astronomy, 6, 584, doi: 10.1038/s41550-022-01618-x

Younsi, Z., Porth, O., Mizuno, Y., Fromm, C. M., & Olivares, H. 2020, in Perseus in Sicily: From Black Hole to Cluster Outskirts, ed. K. Asada, E. de Gouveia Dal Pino, M. Giroletti, H. Nagai, & R. Nemmen, Vol. 342, 9–12, doi: 10.1017/S1743921318007263

Younsi, Z., Wu, K., & Fuerst, S. V. 2012, A&A, 545, A13, doi: 10.1051/0004-6361/201219599

Yuan, F., & Narayan, R. 2014, ARA&A, 52, 529, doi: 10.1146/annurev-astro-082812-141003

Yuan, F., Peng, Q., Lu, J.-f., & Wang, J. 2000, ApJ, 537, 236, doi: 10.1086/309020

Yusef-Zadeh, F., Wardle, M., Schödel, R., et al. 2016, ApJ, 819, 60, doi: 10.3847/0004-637X/819/1/60

Zhdankin, V., Werner, G. R., Uzdensky, D. A., & Begelman, M. C. 2017, Phys. Rev. Lett., 118, 055103, doi: 10.1103/PhysRevLett.118.055103

APPENDICES

Appendix A

Importance of Variability in GRMHD simulations

In Chapter 2, we construct a static model and compare it to variable simulations. Taking this procedure in reverse, it is possible to uniquely decompose the simulations into an average piece and fluctuations. If we identify the average with our static model, and impose that it follows the static GRMHD equations, we can use the simulations to measure how much the variability back-reacts on the average solution.

A.1 Averaging

Here, we detail the theory behind the time- and azimuthal-averaging process. The conservation laws take the form of a non-linear system of partial differential equations

$$\partial_\alpha F^{\alpha i}(\vec{v}) + S^i(\vec{v}) = 0 \tag{A.1}$$

with $\vec{v} = \{\rho, p, u^\alpha, b^\alpha\}$. Written in this form, the fluxes are $F^{\alpha i} = \{\rho u^\alpha, T^{\alpha\beta}, M^{\alpha\beta}\}$, and the source terms are

$$S^i = F^{\alpha i} \frac{\partial_\alpha \sqrt{-|g|}}{\sqrt{-|g|}} + \{0, \Gamma^\beta_{\alpha\gamma} T^{\alpha\gamma}, 0\}. \tag{A.2}$$

The solutions to the GRMHD simulations are variable, and we'll assume they can be described by some average state with fluctuations. In the absence of these fluctuations, Equation A.1 becomes a set of two-dimensional differential equations, vastly simplifying the problem. The fluctuations will, however, back-react on the average state, so we first measure and classify this effect.

We decompose these fluid variables into their average values and fluctuations about this mean as

$$\vec{v}(t, r, \theta, \phi) = \langle\vec{v}\rangle(r, \theta) + \delta\vec{v}(t, r, \theta, \phi) = \frac{1}{2\pi T}\int_0^T \int_0^{2\pi} \vec{v} d\phi dt + \delta\vec{v}. \tag{A.3}$$

with angled brackets averaging over t and ϕ, and $\langle\delta\vec{v}\rangle = 0$ by construction. We can expand the fluxes and source terms in Equation A.1 about these averages

$$F^{\alpha i}(\vec{v}) = F^{\alpha i}\big|_{\langle\vec{v}\rangle} + \frac{\partial F^{\alpha i}}{\partial v^j}\bigg|_{\langle\vec{v}\rangle} \delta v^j + \frac{1}{2}\frac{\partial^2 F^{\alpha i}}{\partial v^j \partial v^k}\bigg|_{\langle\vec{v}\rangle} \delta v^j \delta v^k + \mathcal{O}\left(\delta v^j \delta v^k \delta v^l\right). \tag{A.4}$$

Note that this expansion terminates at fourth order by the nature of the GRMHD equations. The average fluxes are the fluxes of the average fluid variables plus contributions from correlated fluctuations,

$$\langle F^{\alpha i}\rangle = F^{\alpha i}\big|_{\langle\vec{v}\rangle} + \frac{1}{2}\frac{\partial^2 F^{\alpha i}}{\partial v^j \partial v^k}\bigg|_{\langle\vec{v}\rangle} \langle\delta v^j \delta v^k\rangle + \mathcal{O}\left(\langle\delta v^j \delta v^k \delta v^l\rangle\right). \tag{A.5}$$

It is precisely these extra contributions that we will measure from the GRMHD simulations.

Upon averaging Equation A.1, we find that the terms involving the fluxes are

$$\begin{aligned}\left\langle\partial_\alpha F^{\alpha i}\right\rangle &= \frac{1}{2\pi T}\int_0^{2\pi} F^{ti}\Big|_0^T d\phi + \frac{1}{2\pi T}\int_0^T F^{\phi i}\Big|_0^{2\pi} dt + \partial_r\langle F^{ri}\rangle + \partial_\theta\langle F^{r\theta i}\rangle && \text{(A.6)}\\ &\approx \partial_r\langle F^{ri}\rangle + \partial_\theta\langle F^{r\theta i}\rangle, && \text{(A.7)}\end{aligned}$$

and a similar expression for the source terms. The azimuthal integrals go to zero exactly, but the temporal terms only go to zero as $T \to \infty$ and the start and end fluid states are sufficiently similar. This is an assumption that the accretion flow is in a steady state with saturated variability and no secular evolution. We've also assumed that the metric is independent of the t and ϕ coordinates, a condition satisfied by the Kerr-Schild metric used in the simulations.

Thus, we define two types of residuals. To measure the effects of numerically integrating by summing simulation dumps, non-infinite T, and numerical floors in the GRMHD code, we define R_1^i by

$$\partial_r\langle F^{ri}\rangle + \partial_\theta\langle F^{r\theta i}\rangle + \partial_r\langle S^r\rangle + \partial_\theta\langle S^{r\theta}\rangle = R_1^i. \tag{A.8}$$

To measure the effects of the average of correlations between fluctuations, we define R_2^i as

$$\partial_r \left(F^{ri}\big|_{\langle \vec{v} \rangle} \right) + \partial_\theta \left(F^{\theta i}\big|_{\langle \vec{v} \rangle} \right) + \partial_r \left(S^{r}\big|_{\langle \vec{v} \rangle} \right) + \partial_\theta \left(S^{\theta}\big|_{\langle \vec{v} \rangle} \right) = R_2^i. \tag{A.9}$$

In addition to the differential equations, we have two analytical equations, whose residuals we append to R_1^i and R_2^i. Respectively, these are $\langle u^\alpha u_\alpha \rangle$ and $\langle b^\alpha u_\alpha \rangle$ to R_1^i, and $\langle u^\alpha \rangle \langle u_\alpha \rangle$ and $\langle b^\alpha \rangle \langle u_\alpha \rangle$ to R_2^i.

A.2 The Magnitude of Variability in GRMHD Simulations

In this section, we use the GRMHD simulations to calculate both the residuals in the averaged equations and in the equations of the averages. To interpret the results, we must first create a normalization with which to compare these residuals.

A.2.1 Normalization

Due to GRMHD simulations solving differential equations, the solutions will contain a fractional numerical error, which will have spatial structure as the underlying fluid variables change. Primarily, all primitive variables increase in magnitude close to the black hole. Therefore, we need a normalization to define what constitutes a "small" residual.

For analytical equations, we set the normalization to 1 for the residuals associated with $u^\alpha u_\alpha = -1$ and $\sqrt{\langle b^\alpha b_\alpha \rangle}$ for residuals associated with $b^\alpha u_\alpha = 0$. For differential equations of the form of Equation A.1, we choose the normalization

$$\mathcal{N}^i = \left\langle \left| \partial_r F^{ri} \right| + \left| \partial_\theta F^{\theta i} \right| + \left| \partial_\phi F^{\phi i} \right| + \left| S^i \right| \right\rangle \tag{A.10}$$

This effectively measures the typical size of the terms being balanced in the equations.

A.2.2 Testability of Simulations

Figure A.1 shows the normalized residuals of the GRMHD simulations including only the effects of the numerical integration. The quantity shown is the L2 norm over all the equations, that is

$$\mathcal{L}_2 \left(\frac{R_1^i}{\mathcal{N}^i} \right) = \sqrt{ \sum_{i \in \mathrm{Eq.}} \left(\frac{R_1^i}{\mathcal{N}^i} \right)^2 } \tag{A.11}$$

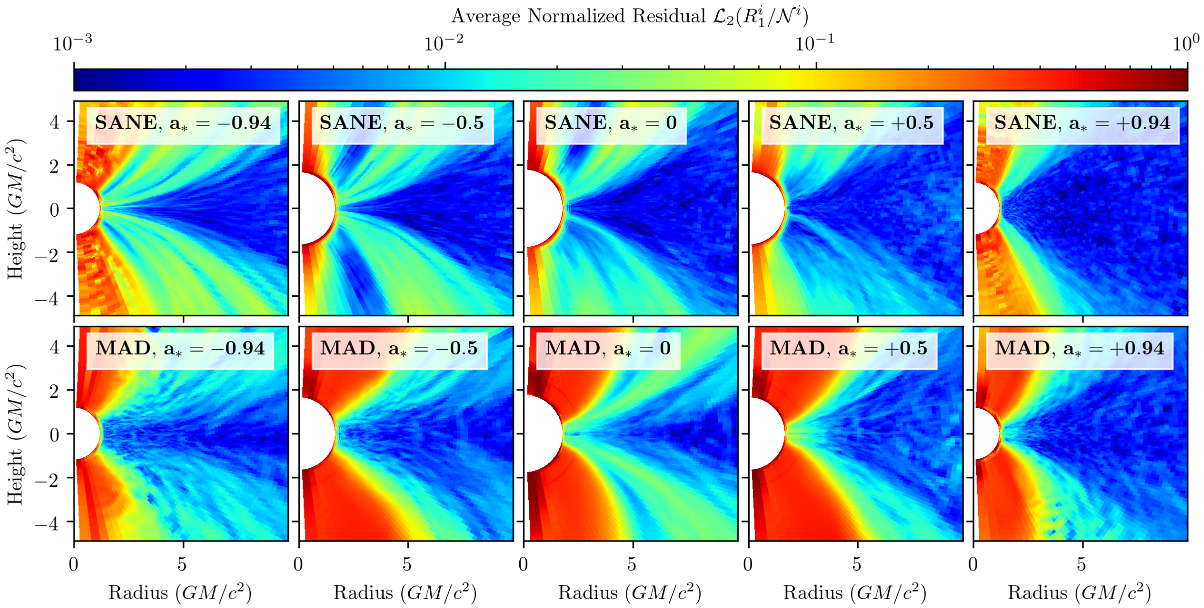

Figure A.1: Normalized residuals of the average of the GRMHD equations. The residuals are small near the midplane but become greater in the jet region. The white region in the middle left of each panel corresponds to the region outside the simulation domain, not the event horizon.

with the colorbar chosen to highlight any residuals more than 1%. The residuals have structure, with numerical effects becoming stronger away from the midplane, and often along contours that resemble the jet-corona boundary.

There are a few empirically-found reasons for non-negligible residuals. Primarily, an average quantity at each location corresponds to a summation of a subsampled red-noise signal with only 128×1000 points. For our choice of normalization, this creates a numerical floor of about 10^{-3}. Near the polar axis, the simulations have an evacuated region and are filled with a numerical density floor to avoid unphysical wave speeds. In addition to the larger grid size near the poles, the average residuals of the GRMHD equations become greater there.

Between the poles and the midplane, the MAD models accrue more error than the SANEs do solely from the continuity equation, roughly corresponding to the region where $\langle b^\alpha b_\alpha \rangle / \langle \rho \rangle \gtrsim 1$. Here, a typical dump has a very low density, so numerical floors and effects become more impactful. The SANE models have contributions from each constituent equation roughly equally.

Thus, for these GRMHD simulations, we can only expect to measure the effects of variability in the accretion disk and corona region.

A.2.3 Effects of Variability

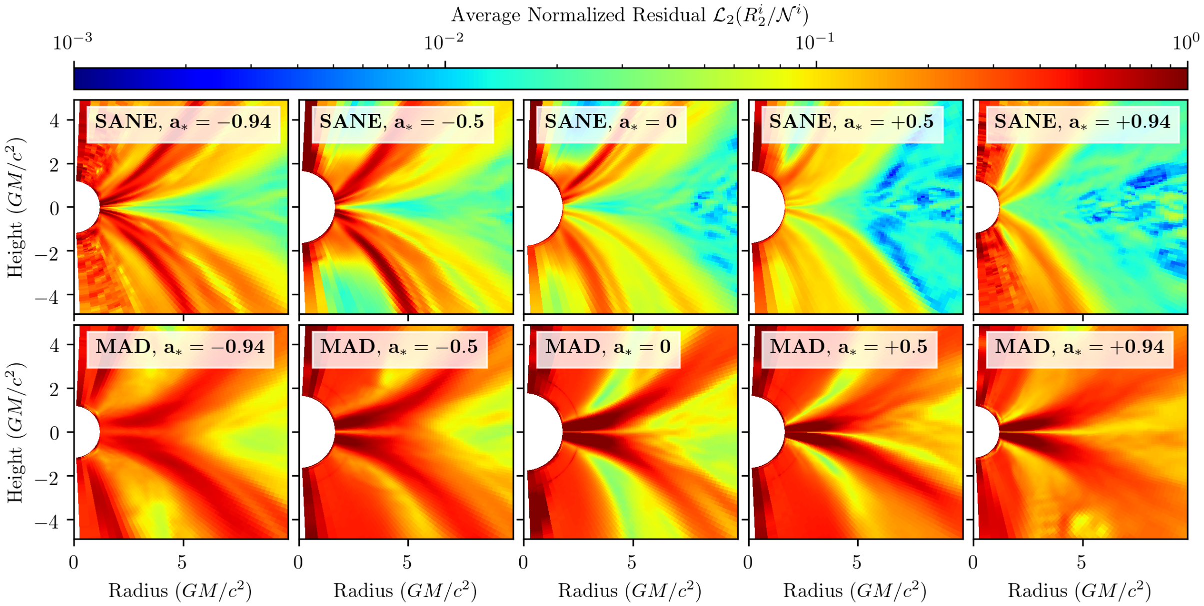

Figure A.2: Similar to Figure A.1, but showing normalized residuals of the GRMHD equations when evaluated for the average quantities. The residuals are greater but still have the smallest residuals in the midplane. MAD models have greater residuals than SANE models.

Figure A.2 shows the normalized residuals similar to Figure A.1, but now separating all the averages of products in the GRMHD equations into products of averages. The residuals are greater, with the general structure remaining the same.

At the midplane, SANE models have small residuals, implying that incorporating the variability is not important for midplane profiles. MAD models have thin disks squished by strong magnetic pressures and characterized by outward bursts of magnetized evacuated regions. Evidently, these variable effects have a significant impact in the region directly above the midplane.

Figure A.3 and Figure A.4 show the contributions to the average residuals from each constituent equation for a SANE, $a_* = -0.94$ and MAD, $a_* = 0$ simulation, respectively.

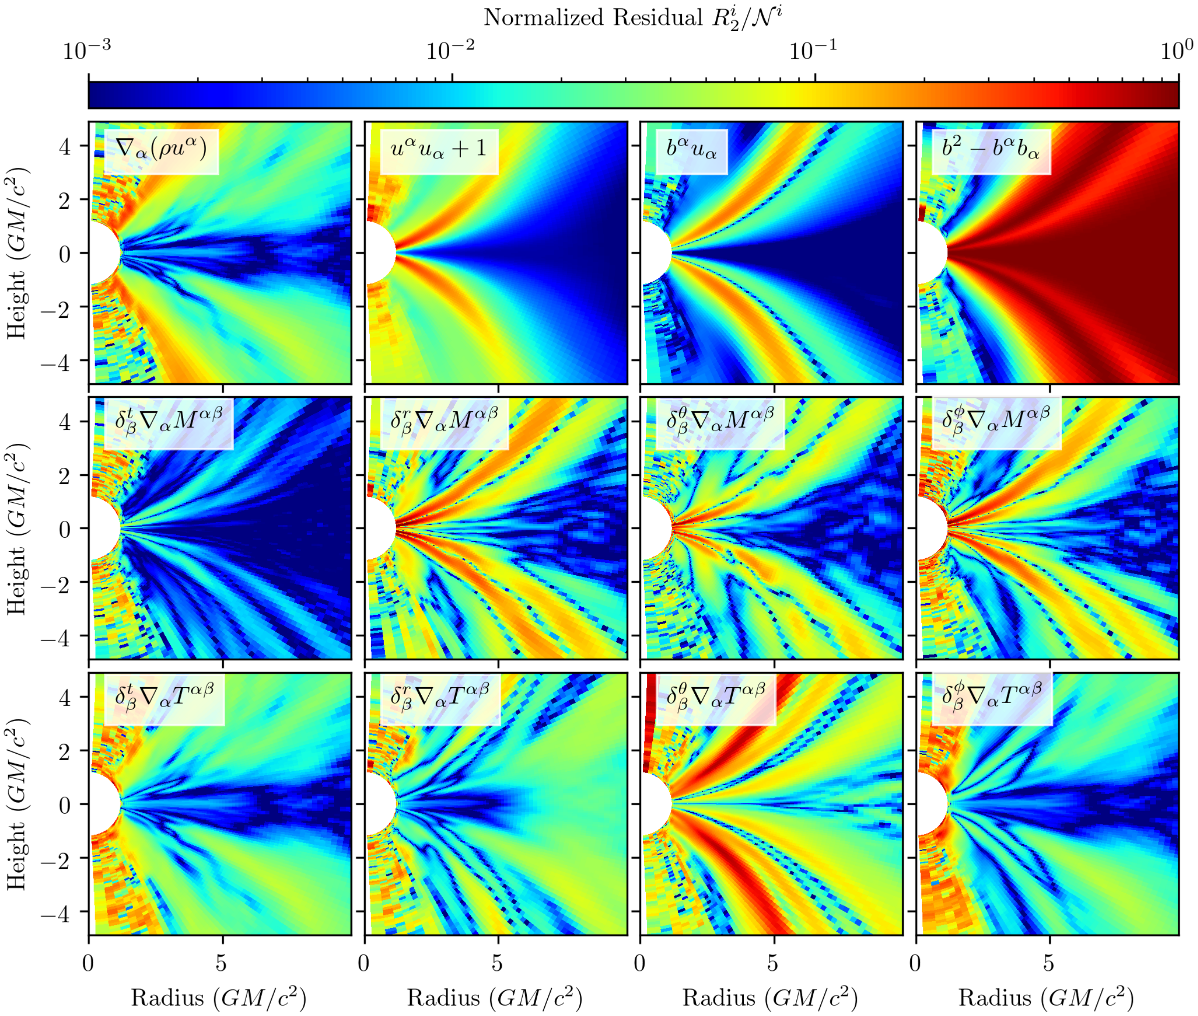

Figure A.3: Normalized residuals split up by the constituent GRMHD equation for the SANE simulation with $a_* = -0.94$. The equations in the plots signify which equation's residual is being shown. For the top right panel, we show $1 - \langle b^\alpha \rangle \langle b_\alpha \rangle / \langle b^\alpha b_\alpha \rangle$, which is not included in Figure A.2.

Shown also in these figures' top right panel is the difference between the average magnetic pressure $\langle b^\alpha b_\alpha \rangle$ and the magnetic pressure caused by the average magnetic field $\langle b^\alpha \rangle \langle b_\alpha \rangle$.

The SANE simulation has large residuals in parabolic bands in the corona region. These correspond to where the standard deviation of the variability in the 4-velocity be-

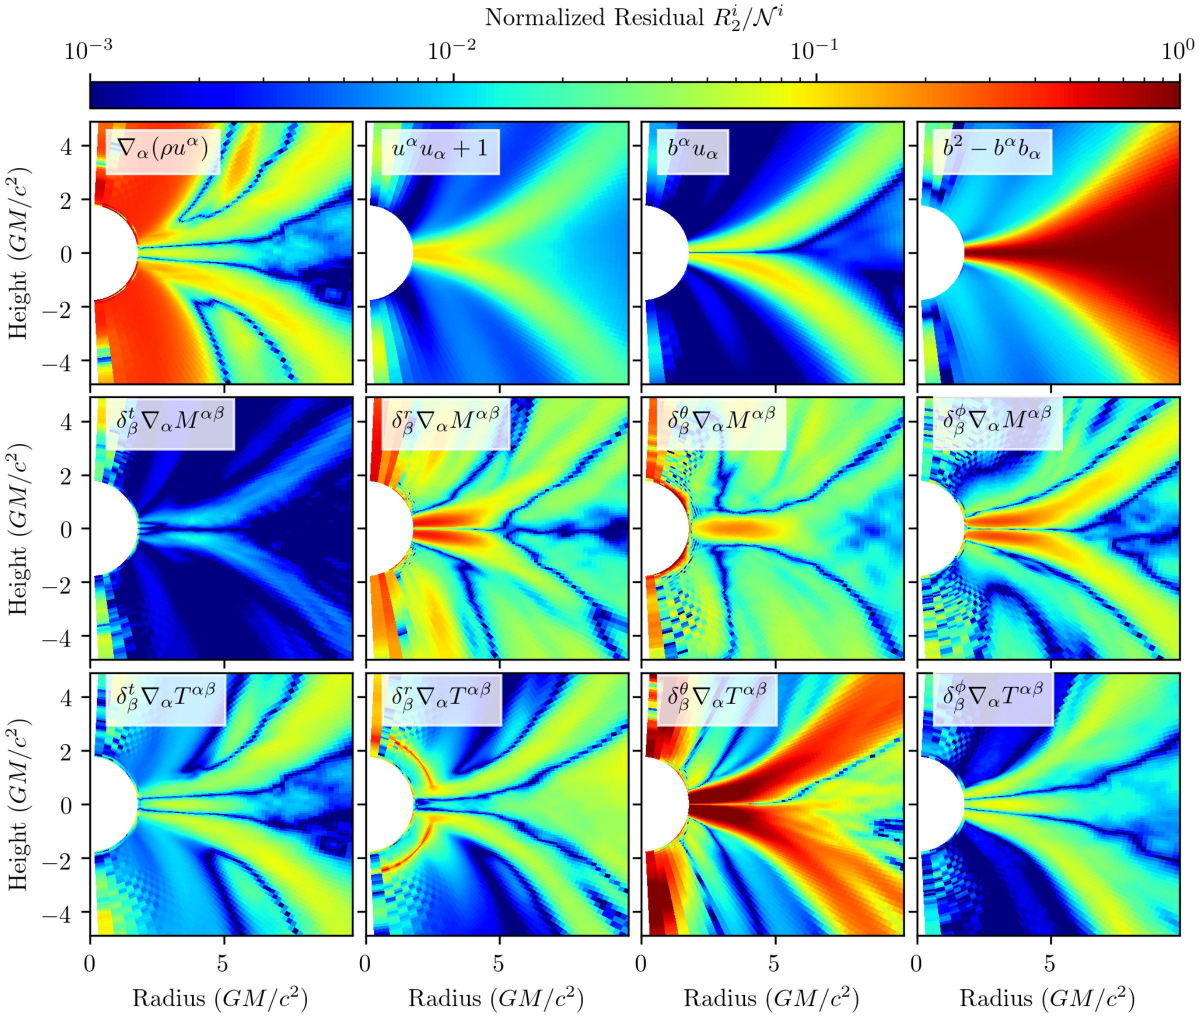

Figure A.4: Same as Figure A.3, but for a MAD simulation with $a_* = 0$.

comes comparable to its magnitude, as indicated by the large residuals in the velocity normalization conditions. Besides this, only the θ-component of the energy-momentum equations has large residuals in this region, indicating that the turbulent fluctuations are only transporting poloidal momentum. The residuals in the electromagnetic equations have more structure (i.e., they are noisier), but this is due to the relatively small contributions of the magnetic field to the accretion flow.

The largest residuals in the MAD simulation also occur from the vertical momentum balance equation. Since the disk in the MAD simulations is thinner, these large residuals extend nearly to the midplane. The MAD simulations have lower radial velocity than the SANEs, so the relative fluctuations in the velocity have a greater effect on the velocity normalization equation. The rest of the energy-momentum equations have a similar structure to the mass conservation law, hinting that the residuals are caused by variability that only transports momentum by a loss of mass. The source of the mass loss through correlated fluctuations remains unknown, however. The magnetic fields are more ordered in MADs than SANEs, so their residuals are less noisy. We leave the explanation of specific contributions of constituent equations to the overall variability for future work.

Appendix B

Limitations of Measuring Variability in GRMHD Images of Black Holes and Applications to Sgr A*

B.1 The Impact of GRMHD Numerical Approximations

B.1.1 Dependence on the Fastlight Approximation

The model sets used in this work primarily use the fastlight approximation, where the fluid output is saved at constant values of the Kerr-Schild timelike coordinate and each image is constructed using only one of these fluid outputs, effectively assuming that the background flow is constant in time as light propagates through all space. In contrast, slowlight (or finite-speed light) methods account for the evolution of the fluid by loading the changing fluid state (i.e., a sequence of constant-Kerr-Schild time snapshots) as light propagates through space. In our approach, the fluid data at an arbitrary spacetime event are synthesized by performing linear interpolation in both space (between neighboring grid zones) and time (between sequential fluid snapshots). Since the fluid velocity at the emission region has a significant component in each fluid output and the scale of the system is not much smaller than its light-crossing time, we expect an approximated fastlight image to deviate from the truth for temporal and spatial scales comparable to the gravitational scales.

To test the numerical effects of this approximation on the power spectra, we imaged a MAD simulation with $a_* = 0.94$ with fluid outputs at a short cadence of 0.5 GM/c^3

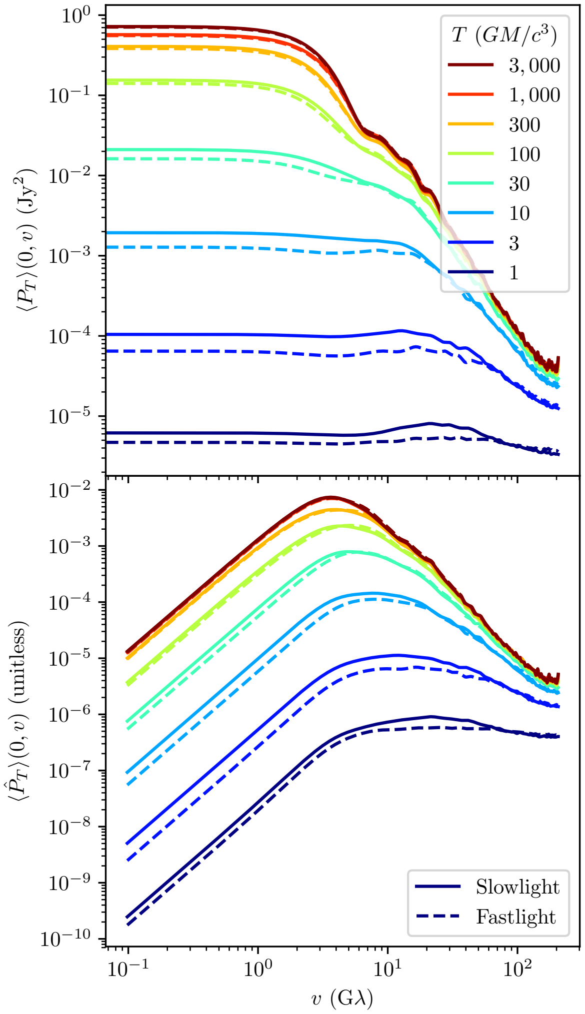

Figure B.1: Effect of the fastlight approximation on the PSDs for the same simulation as in Figure 3.5 and Figure 3.6. The top panel shows $\langle P_T \rangle$ for images in which the fluid evolves during one photon trajectory (slowlight; solid lines), and for images where the fluid is set static (fastlight; dashed lines). The bottom panel shows the same for $\langle \hat{P}_T \rangle$. On shorter timescales, the fastlight approximation suppresses variability.

(model set B). For the fastlight imaging, each fluid output creates one image for a viewer far from the black hole. For the slowlight imaging, the same viewing screen was offset in time corresponding to the time the light ray would need to travel from the black hole so that the fastlight and slowlight images would roughly span the same accretion history. To find this offset more exactly, we compared the light curves of the two simulations and maximized their correlations relative to a time shift. The light curves are dominated by the longest timescales, and thus are nearly identical with the appropriate time shift. After this calibration procedure, the two sets of images have the same cadence and nearly the same start and end point.

Figure B.1 shows the effect of using the fastlight approximation on $\langle P_T \rangle$ and $\langle \hat{P}_T \rangle$. Solid lines show the "true" slowlight quantities, and dashed lines show the fastlight approximation. For long timescales, the approximation holds, likely due to long timescale variability caused by gas farther from the black hole, which only has a small component of its velocity orthogonal to the timelike coordinate. On shorter timescale, the fastlight approximation consistently suppresses variability on all spatial scales, though the magnitude of the difference is not a particularly large one.

B.1.2 Dependence on Simulation Resolution

Numerical GRMHD simulations lay out a grid upon which the fluid is evolved. Therefore, they set a lower boundary on the spatial size of fluctuations of the fluid properties. It is not known whether physics at the unmodelled small scales could cause large-scale deviations in the fluid properties through perhaps a different saturation of the MHD turbulence. These potential different states could manifest as different variability power spectra.

In order to test the effects of finite spatial resolution, we have performed GRMHD simulations of a MAD model with $a_* = 0.94$ in three different resolutions, 96^3, 128^3, and 192^3 starting from an identical initial condition (model set E). Figure B.2 shows the variability PSDs for these three simulations. Although there are differences in the variability PSDs, there is no clear trend with resolution, hinting that the main source of the variability stems from physics on the resolved larger scales. This test does not preclude the possibility of the simulations not being sufficiently resolved, but resolving such effects likely requires far greater resolutions, and is thus outside the scope of this work.

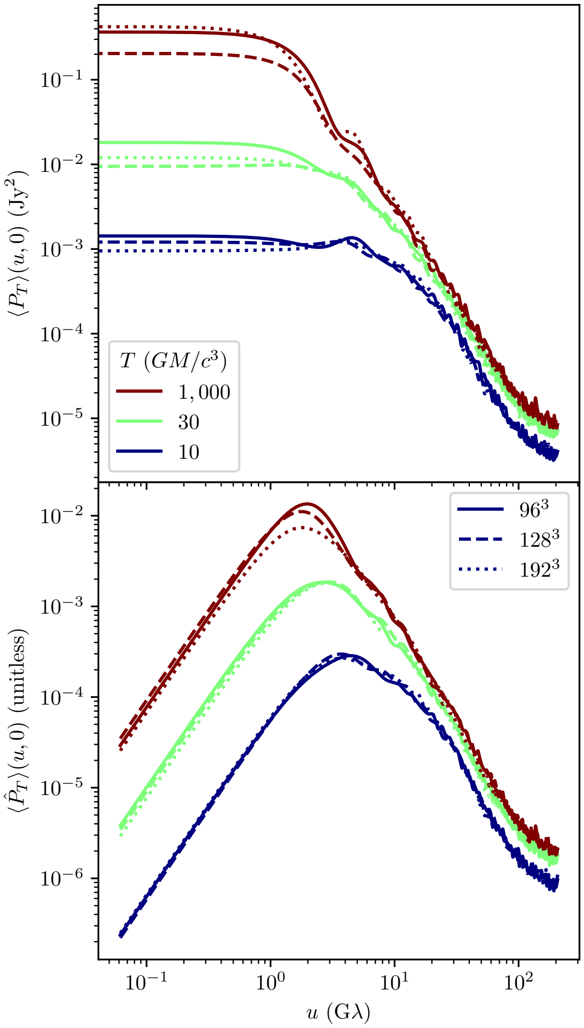

Figure B.2: Effect of varying fluid simulation resolution on the variability PSDs. In the top panel, the lines measure the average mean-subtracted PSD for the same simulation with a resolution of 96^3 (solid), 128^3 (dashed), and 192^3 (dotted). The bottom panel shows the same for the average residual PSD. For the case considered, no clear trend is seen for any of the timescales measured.

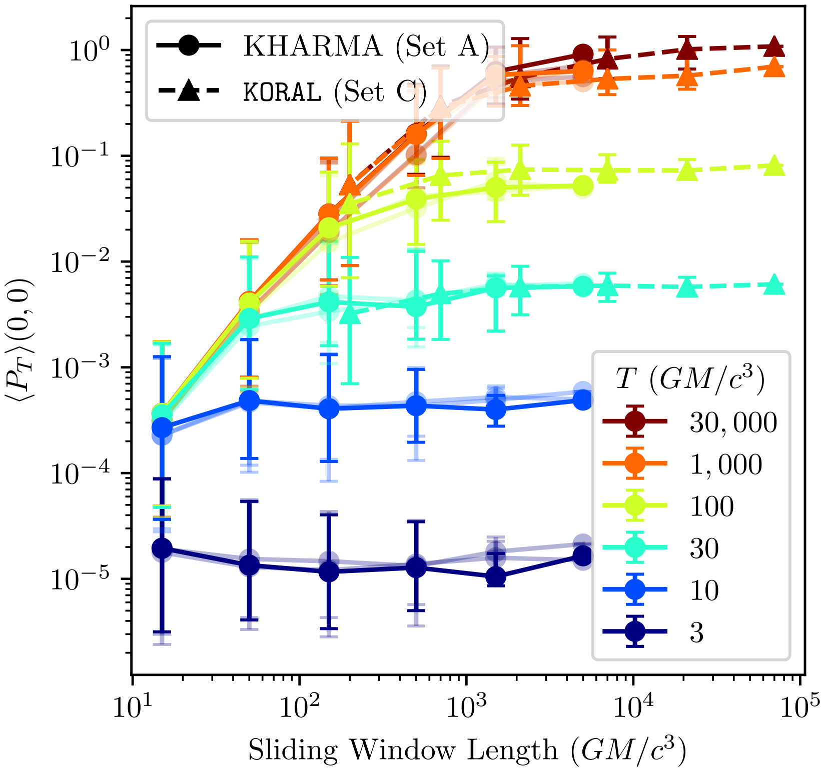

Figure B.3: Effect of simulation length on the measurement of the average mean-subtracted PSD. For a simulation of length τ, we designate a sliding window of some shorter length with a separate realization starting every frame. The points show the median of $\langle P_T \rangle(0,0)$ over all sliding windows, and the error bars show the inner 68% percentiles. From model set A (circles, solid lines), we use a MAD state with $a_* = -0.94$, $R_{\text{high}} = 10$, and $i = 10$, with each of its three segments shown. From model set C (triangles, dashed lines), we use a MAD state with $a_* = -0.9$, $R_{\text{high}} = 20$, and $i = 10$.

B.1.3 Dependence on the Turbulent Realization

For each of the simulations considered in this work, the total length of time available (τ in Section 3.3) is finite, and thus we only have access to one possible realization of turbulence. Therefore, for a given simulation length, our PSD measurement will be one random draw from a distribution of possible PSDs. Here we try to characterize any bias and the distribution width expected for such a measurement.

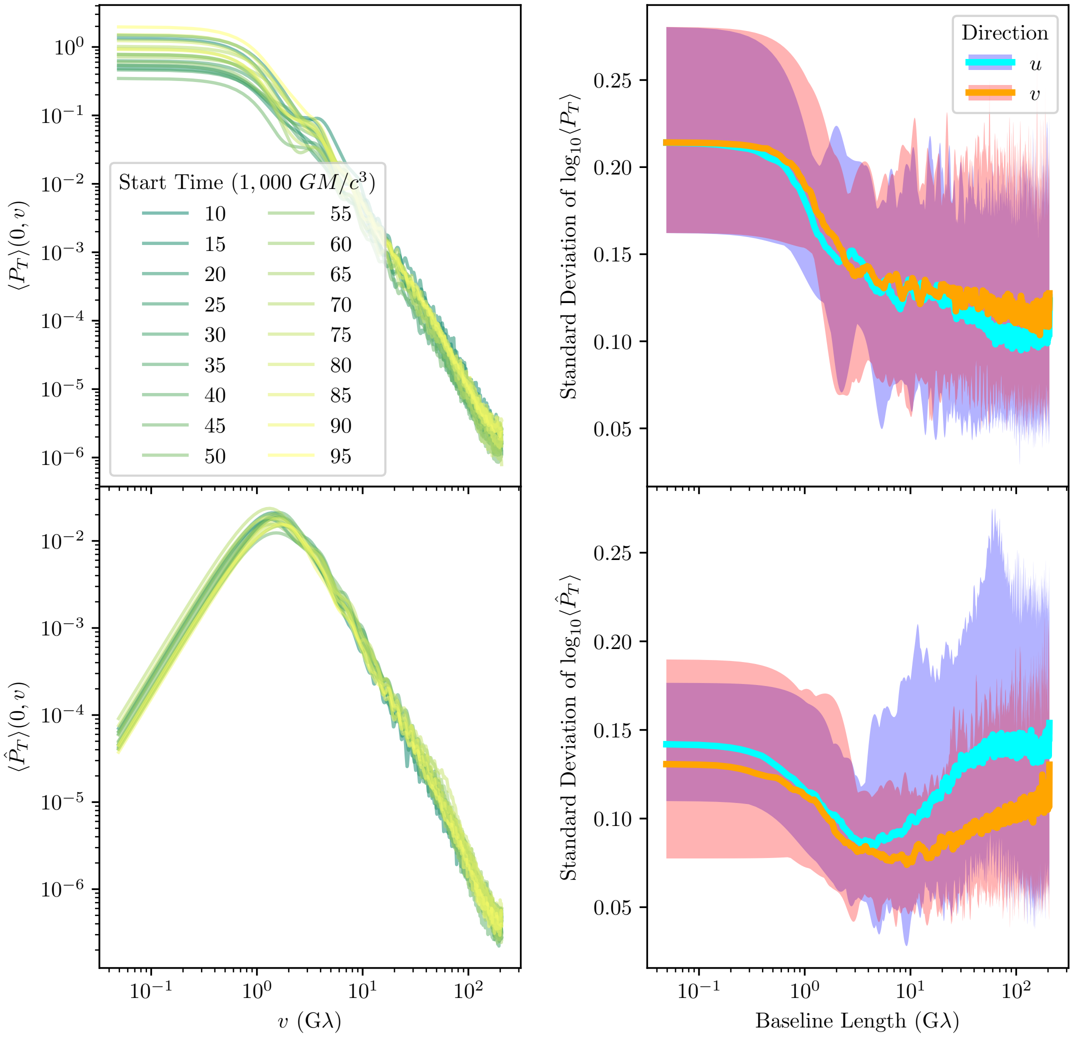

Figure B.4: Effect of simulation length on the variability PSDs. The left column shows $\langle P_T \rangle(0, v)$ (top) and $\langle \hat{P}_T \rangle(0, v)$ (bottom) for a MAD simulation from model set C with $a_* = -0.9$, $R_{\text{high}} = 20$, and $i = 10$ with $T = 30,000\ GM/c^3$. This simulation has been split up into 18 windows of length 5,000 GM/c^3. No trend exists with window start time, but a scatter exists. The right column shows a measurement of this scatter for every simulation in model set C along $v = 0$ (blue) and $u = 0$ (red). The median value over these simulations is shown along $v = 0$ (cyan) and $u = 0$ (orange) and shaded bands show the entire range over 81 simulations.

For the conclusions in this work, we primarily use simulations of length $5,000\ GM/c^3$. As a representative, we choose a simulation in the MAD state from model set A with $a_* = -0.94$, $R_{\rm high} = 10$, and $i = 10$. It contains three independent realizations. As the "true" model, we take a simulation in the MAD state from model set C, with $a_* = -0.9$, $R_{\rm high} = 20$, and $i = 10$. Using only frames with cadence $50\,GM/c^3$, it covers $70,000\,GM/c^3$, thus containing the equivalent of 14 separate realizations of the representative simulations.

From each simulation, we constructed a sliding window of some shorter time. From each sliding window, we calculate $\langle P_T\rangle(0,0)$, treating each independently. Figure B.3 shows $\langle P_T\rangle(0,0)$ for different sliding window lengths. For each sliding window length, the error bars and points show the 16, 50, and 84 percentiles over every sliding window. Note that these sliding windows are not independent, and the number of uncorrelated sliding windows is approximately the ratio of the sliding window length to τ. Thus, the measurements of the percentiles for ratios $\lesssim 10$ are not robust. Where there are enough independent sliding windows, these percentiles represent a distribution of potential measurements of $\langle P_T\rangle(0,0)$ for a simulation with shorter τ.

Building confidence, a measurement of the median of $\langle P_T\rangle(0,0)$ using the shorter simulation does fall within the measured distribution from the longer simulation, despite the two simulations having slightly different values of spin and $R_{\rm high}$. For both simulations, this median is constant over sliding window length, with a deviation occurring when the window length is comparable to the averaging time, T. This is simply because the numerical Gaussian mean-subtraction from Equation 3.7 truncates when the simulation ends, thus shortening the effective T used in the calculation of $\langle P_T\rangle(0,0)$. The distribution of $\langle P_T\rangle(0,0)$ becomes narrower with increasing sliding window length, though it is unclear if that is due to fewer samples or a stabilization of the variability on long enough timescales or both.

We can use these results to extrapolate the PSD measurements from finite to infinite τ. For these accretion flow parameters, the measurement of the distribution of $\langle P_T\rangle(0,0)$ calculated using multiple realizations of total length only $\tau \gtrsim 1,500\,GM/c^3$ contains the infinite τ limit within the distribution of these realizations. For $\tau \lesssim 1,500\,GM/c^3$, this empirical extrapolation is only valid up to some maximum value of T.

Assuming the conclusions found for $\langle P_T\rangle(0,0)$ hold for $\langle P_T\rangle(u,v)$ and $\langle \hat{P}_T\rangle(u,v)$, we now characterize how the uncertainty associated with a single turbulent realization behaves across different baseline lengths. The left column in Figure B.4 shows the same procedure as Figure B.3 for nonzero baselines. A slight difference exists, since in this test, we do not require a constant simulation cadence, and it includes images from earlier in the simulation. No clear trend exists with starting time, indicating that the variability is not secularly evolving. These PSDs combined form a distribution, whose width we characterize as a scatter by the standard deviation of the logarithm of the variability PSDs. The right column in Figure B.4 shows this modulation index for all the simulations in model set C, each containing a potentially different cadence and simulation length. The bands show the maximum range of possible scatters over all simulations. For $\langle \hat{P}_T \rangle(u, v)$, the scatter is relatively constant in baseline length and (u, v)-direction. We take this to mean that a single measurement of $\langle \hat{P}_T \rangle(u, v)$ for simulation length τ is different from any other realization of $\langle \hat{P}_T \rangle(u, v)$ by approximately a tenth of an order of magnitude (i.e. 25%). Note that for a τ small enough determined by some timescale of the simulation ($\tau \lesssim 1,500\, GM/c^3$ in Figure B.3), there will be a bias along with this distribution width. Further note that model set C only contains MAD models with $R_{\text{high}} = 20$, and other accretion flow parameters could show different distribution widths.

B.1.4 Code Comparison

GRMHD simulations include many numerical approximations and choices that vary between codes. To test the effects of these different implementations, we compare the variability PSDs from model sets A (using only inclinations less than 90 degrees) and D. These come from different fluid modeling codes (KHARMA and `BHAC`) and different GRRT codes (`ipole` and `BHOSS`), and contain 200 simulations with matching fluid and imaging parameters (see Subsection 3.2.2). Each simulation has three windows in model set A and two windows in model set D. To compare these, we introduce a difference measure,

$$\mathcal{D}(x, y) = 2 \left\langle \left| \log_{10} \frac{x}{y} \right| \right\rangle_{\text{windows}} , \tag{B.1}$$

where x and y stand in for one of the variability PSDs for two simulations, and the averaging is performed over all the possible window combinations. This measure is a natural extension of the standard deviation in the log domain in the limit of two data points, and so is analogous to that used in Figure B.4. Differences between codes could also stem from simulations segments being separate turbulence realizations, so we can compare the difference between windows in separate codes to that between windows in one code. When comparing within the same code, the averaging only includes independent window combinations.

This difference measure is shown for all simulations in Figure B.5. The red and green regions show the range of differences between separate windows in the same simulations, which arises from separate windows corresponding to different turbulence realizations.

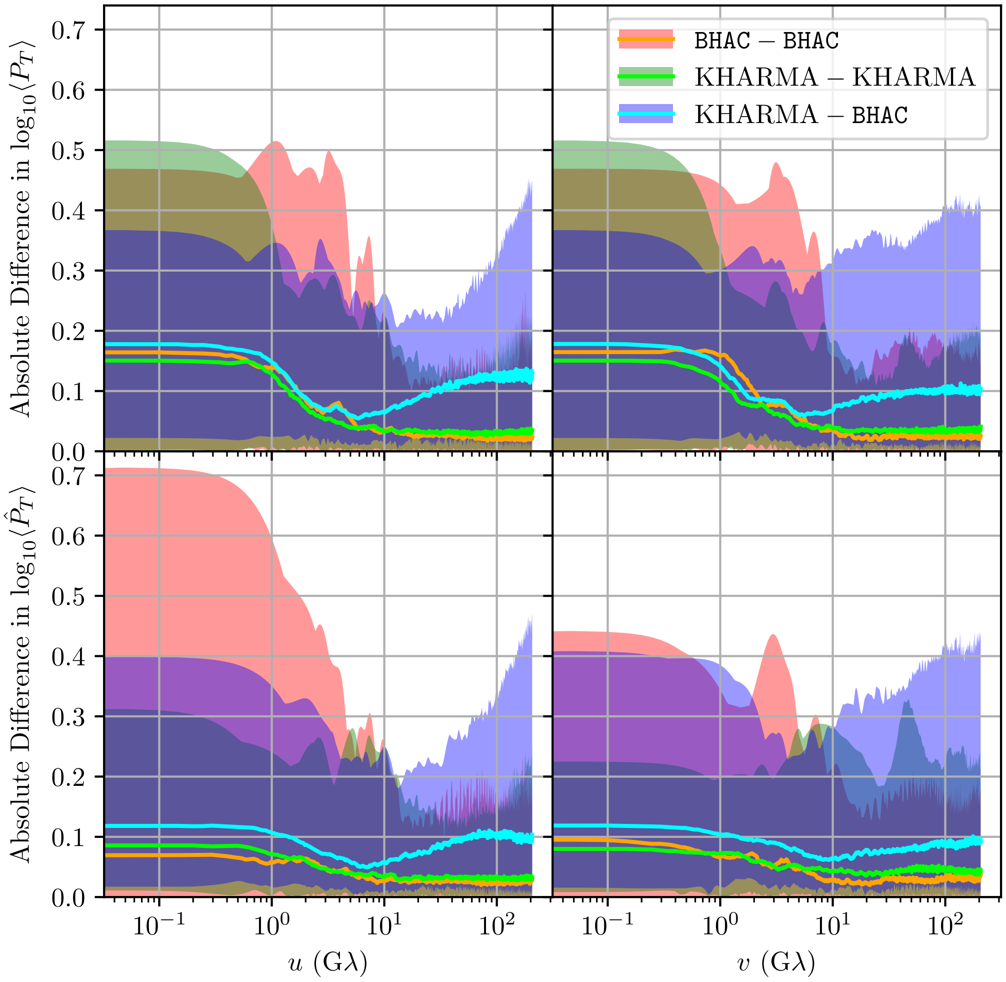

Figure B.5: Effect of code used on the variability PSDs. The quantity shown is the difference measure as defined in Equation B.1. Rows show the two variability PSDs from Section 3.3 and columns show the difference along the u and v directions. Shaded regions show the minimum and maximum of the difference measure over all simulations, and lines show the median. The red/orange (model set A) and green/lime (model set D) regions correspond to the difference between separate windows of the same simulations, while the blue/cyan regions show the difference between model sets A and D, and therefore the difference between codes.

These are similar to the regions in Figure B.4, and include a larger sampling of GRMHD parameters, but with only two or three measurements compared per simulation.

The blue regions show the range of differences between the two model sets. On short baselines $|u| \lesssim 10\mathrm{G}\lambda$, the inter-code and intra-code differences are similar, with the median inter-code difference somewhat larger. On these large scales, the choices in each code are subdominant to the uncertainty in the PSD measurements from a single realization. On longer baselines, the inter-code differences are larger. For $\langle \hat{P}_T \rangle$, the typical (median) deviation between codes is a tenth of an order of magnitude (i.e. 25%).

A comparison of the difference in $\langle \hat{P}_T \rangle$ on short baselines among these model sets can also be seen through the difference of the centroid of emission in the images in Section B.4.

The inter-code differences stem from a variety of numerical and algorithmic choices, some of which have been studied and quantified for GRMHD in Porth et al. 2019 and GRRT in Gold et al. 2020, and we leave further study of code differences for future works.

B.2 Applications to 2017 EHT Observations of Sgr A*

The formalism developed in Chapter 3 has been applied to the 2017 EHT observations of Sgr A*. Since the EHT data is already in the Fourier domain, calculating the lightcurve-normalized mean-subtracted PSD can be done directly, though not without significant challenges pertaining to the sparsity of the data (Broderick et al., 2022b). Figure B.6 shows the variability PSD as measured over 5 days of Sgr A* observations, as well as the expectations from the GRMHD image library (Paper IV). Although the simulations are on average more variable than the data, many of them are consistent with the data. This did not have to be true; the simulations predicted a universal amplitude and slope, and thus, there is some support that, despite all of their assumptions, the GRMHD simulations are producing the right type of variability.

Figure B.7 shows which GRMHD models are consistent with the observed structural variability in Sgr A* at $4\,\mathrm{G}\lambda$ as defined in Subsection 3.5.3. There is a slight favoring of face-on models and a disfavoring of $a_* = -0.94$ models. Specific trends may be understood through Figure 3.13, and the failing models are typically noisier than the data[1]. The specific details of the calculation of the structural variability can be found in Broderick et al. (2022b), Paper IV, and Paper V.

---

[1] In Paper V, the GRMHD code `H-AMR` was used alongside KHARMA and `BHAC`.

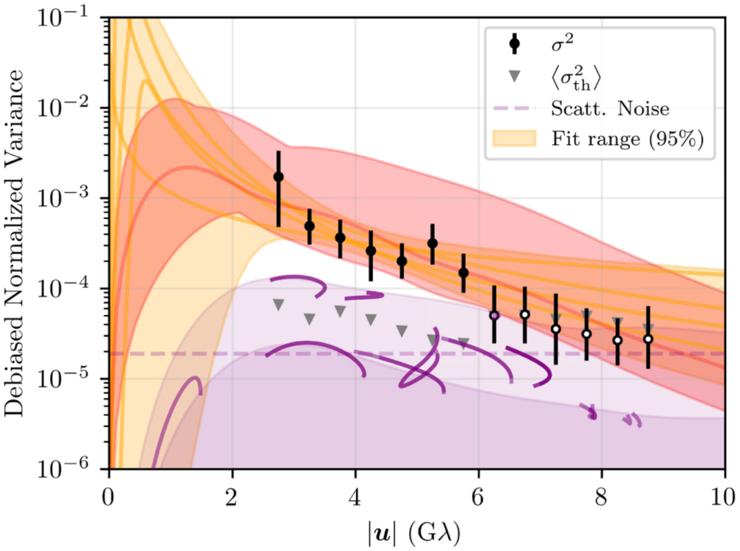

Figure B.6: Normalized structural variability in 2017 EHT observations of Sgr A*. Black points show measurements from the data, and the orange lines and bands are a broken power law fit. The red lines and bands show the range covered by GRMHD simulations in this work. Purple lines and bands represent the refractive scattering noise, and gray triangles are the thermal noise. This figure is figure 18 of Paper IV.

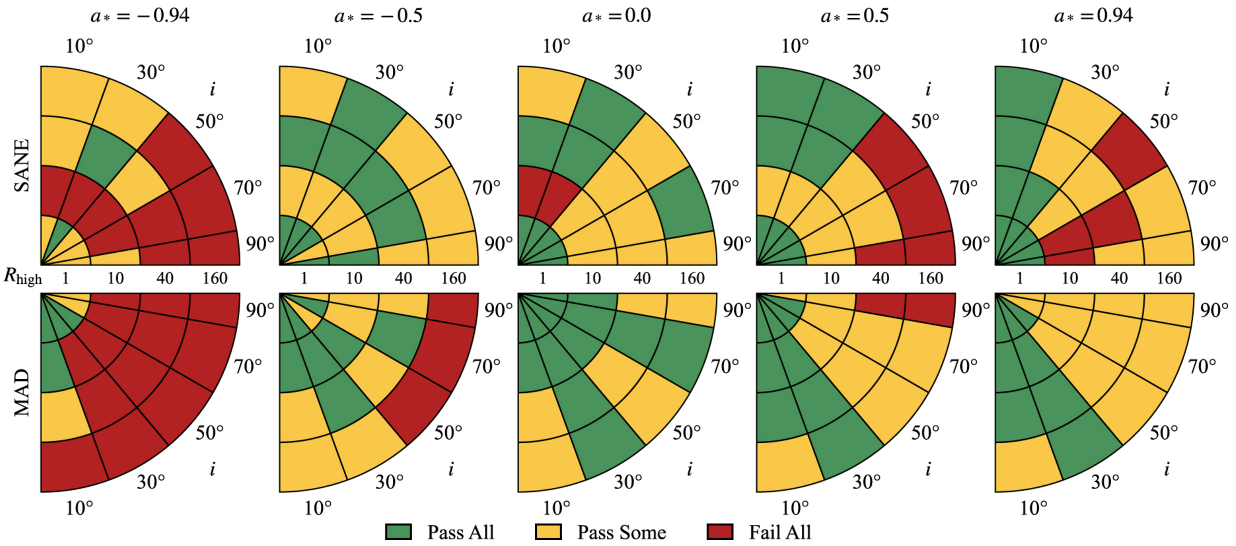

Figure B.7: GRMHD constraints from the normalized structural variability PSD at 4 Gλ and the 2017 EHT observations of Sgr A*. Green indicates that models from all codes pass, yellow that models from one or two codes fail, and red that all three fail. This figure is figure 35 of Paper V.

B.3 Short-baseline PSD and Observational Systematic Uncertaities

The short baseline expansion made in Equation 3.16 faces a number of practical complications in standard VLBI analysis applications. These include:

1. The time-averaged truth image is not known a priori, and thus typically, $\bar{\mathcal{I}}_T$ is unavailable. As a result, any approximation of the mean is not strictly constrained in the same manner as $\bar{\mathcal{I}}_T$ is to $\mathcal{I}$.

2. Each visibility is subject to individual statistical uncertainties (e.g., thermal errors) that impose a floor on the variability associated with the measurement details.

3. In the absence of an absolute phase calibration, unknown atmospheric delays add a large (many 2π) random station-based component to the phases of the measured visibilities. These phase shifts vary on the atmospheric coherence time, which at 230 GHz, can be as short as tens of seconds.

The last of these presents a significant complication for the interpretation of observational estimates of the variability PSDs.

In the absence of a separate phase reference, for VLBI arrays the individual station phase delays are typically set via self-calibration: an optimization process in which the individual station phases are estimated during image reconstruction that leverages the image priors (e.g., positivity) and the fact that the number of baselines, $N(N-1)/2$, is larger than the number of unknown station phase delays, N, when $N > 3$, rendering the problem over-constrained. Nevertheless, this procedure is formally degenerate with an overall phase shift and linear phase gradient. The former is benign, and we will not consider it further. The latter, a linear phase gradient, is equivalent to translating the source in the image domain. It is for this reason that, in the absence of a phase reference, VLBI does not have the capacity to constrain the absolute position of the source being imaged.

Additional practical limitations include additional sources of variability (e.g., interstellar scattering), station gain amplitudes, and polarization leakage, all of which will result in effective mismatches between the true intrinsic $\mathcal{V}$ and that observed. We will, however, neglect these in favor of addressing the phase shifts.

The above considerations enter into the formalism of Subsection 3.3.2 in two ways: first the conceptual severing of $\bar{\mathcal{I}}_T$ and $\mathcal{I}$, and second, the introduction of additional time-dependent phase shifts of $\mathcal{V}$. That is, the observed visibilities are

$$\mathcal{V}_{\rm obs} = \mathcal{V} e^{2\pi i \boldsymbol{u}\cdot\boldsymbol{X}} \tag{B.2}$$

for some arbitrary $\boldsymbol{X}(t)$ that describe the unknown phase shifts. As a result, the observed intensity map is shifted by $\boldsymbol{X}$, i.e., $\mathcal{I}_{\rm obs}(\boldsymbol{x}) = \mathcal{I}(\boldsymbol{x}+\boldsymbol{X})$.

We will require the first four terms (but only the first three explicitly) in the Fourier series expansion employed in Equation 3.16, i.e., after expanding

$$e^{-2\pi i \boldsymbol{u}\cdot\boldsymbol{x}} = 1 - 2\pi i \boldsymbol{u}\cdot\boldsymbol{x} - 2\pi^2 (\boldsymbol{u}\cdot\boldsymbol{x})^2 + \frac{4\pi^3}{3} i (\boldsymbol{u}\cdot\boldsymbol{x})^3 + \dots \tag{B.3}$$

with associated contributions to the complex visibility:

$$
\begin{aligned}
\mathcal{V}_0 &= \int \left[\mathcal{I}(\boldsymbol{x}+\boldsymbol{X}) - \bar{\mathcal{I}}_T(\boldsymbol{x})\right] \mathrm{d}^2x = 0 \\
\mathcal{V}_1 &= -\int 2\pi i (\boldsymbol{u}\cdot\boldsymbol{x})\left[\mathcal{I}(\boldsymbol{x}+\boldsymbol{X}) - \bar{\mathcal{I}}_T(\boldsymbol{x})\right] \mathrm{d}^2x = -2\pi i \boldsymbol{u}\cdot(\boldsymbol{\mathcal{X}} - \boldsymbol{X} - \bar{\boldsymbol{\mathcal{X}}}_T) \\
\mathcal{V}_2 &= -\int 2\pi^2 (\boldsymbol{u}\cdot\boldsymbol{x})^2\left[\mathcal{I}(\boldsymbol{x}+\boldsymbol{X}) - \bar{\mathcal{I}}_T(\boldsymbol{x})\right] \mathrm{d}^2x = -2\pi^2 \boldsymbol{u}^T\boldsymbol{C}\boldsymbol{u} + \frac{\mathcal{V}_1^2}{2} \\
\mathcal{V}_3 &= \int \frac{4\pi^3}{3} i (\boldsymbol{u}\cdot\boldsymbol{x})^3\left[\mathcal{I}(\boldsymbol{x}+\boldsymbol{X}) - \bar{\mathcal{I}}_T(\boldsymbol{x})\right] \mathrm{d}^2x
\end{aligned}
\tag{B.4}
$$

where $\boldsymbol{\mathcal{X}}$ and $\bar{\boldsymbol{\mathcal{X}}}_T$ have their meanings from Subsection 3.3.2 and $\boldsymbol{C}$ is the second moment of $\mathcal{I}(\boldsymbol{x}+\boldsymbol{X}) - \bar{\mathcal{I}}_T(\boldsymbol{x})$. Therefore, to order $|\boldsymbol{u}|^4$,

$$
\begin{aligned}
\langle \hat{P}_T \rangle &= \left\langle |\mathcal{V}_1|^2 + |\mathcal{V}_2|^2 + \mathcal{V}_1^*\mathcal{V}_3 + \mathcal{V}_3^*\mathcal{V}_1 \right\rangle \\
&= \left\langle 4\pi^2 \left[\boldsymbol{u}\cdot(\boldsymbol{\mathcal{X}} - \boldsymbol{X} - \bar{\boldsymbol{\mathcal{X}}}_T)\right]^2 + 4\pi^4 \left(\boldsymbol{u}^T\boldsymbol{C}\boldsymbol{u}\right)^2 + Q\boldsymbol{u}\cdot(\boldsymbol{\mathcal{X}} - \boldsymbol{X} - \bar{\boldsymbol{\mathcal{X}}}_T) \right\rangle
\end{aligned}
\tag{B.5}
$$

where Q is a collection of non-singular terms that are of order $|\boldsymbol{u}|^3$.

In principle, in the absence of additional information, the value of $\boldsymbol{X}$ is arbitrary. In practice, $\boldsymbol{X}$ will be selected on the basis of the prior assumptions underlying the analysis, e.g., temporal regularization or minimization of the deviations from a static model. When they are chosen such that $\boldsymbol{X} = \boldsymbol{\mathcal{X}} - \bar{\boldsymbol{\mathcal{X}}}_T$, the quadratic terms are eliminated altogether, leaving the short-baseline behavior of the average residual PSD completely determined by the quartic term, $\langle \hat{P}_T \rangle = 4\pi^4 \left(\boldsymbol{u}^T\boldsymbol{C}\boldsymbol{u}\right)^2$. Where $\boldsymbol{X} \neq \boldsymbol{\mathcal{X}} - \bar{\boldsymbol{\mathcal{X}}}_T$, the short-baseline PSD is reflective of the manner in which this is selected.

B.4 GRMHD simulation images and their centroids

Figure B.8, Figure B.9, and Figure B.10 provide average images of the GRMHD simulations in model set A corresponding to the last window and only showing inclinations of 10, 50, and 90 degrees. Overlaid on the average images are the 1-σ covariance ellipses of the centroid of emission for model sets A and D, scaled larger by a factor of 10. These are directly related to the trends discussed in Subsection 3.5.2.

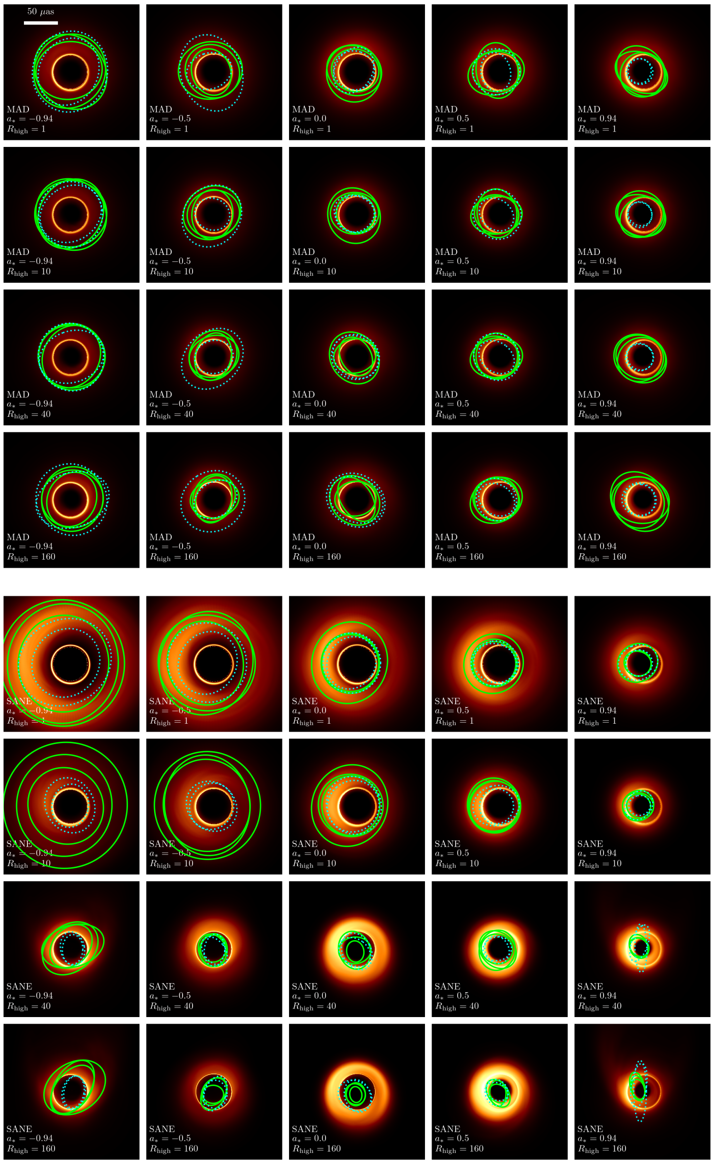

Figure B.8: Average images of the GRMHD simulations at 10-degree inclination. Ellipses show 1-sigma contours of the distribution of the centroid of emission for each window in model sets A (solid green) and D (dotted cyan), scaled up by a factor of 10 for visual clarity. Average images are for the last window of model set A only. The black hole axis is pointed up for positive spin, and down for negative spin. Color corresponds to intensity per pixel, and is not the same among images.

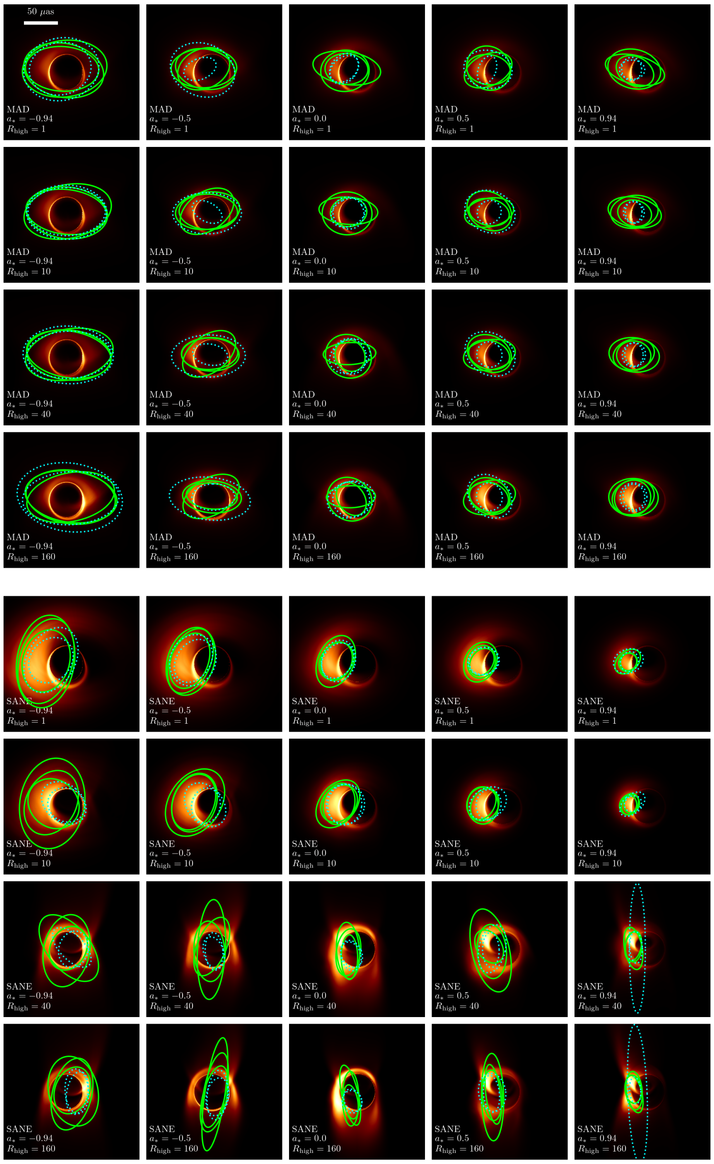

Figure B.9: Similar to Figure B.8 for 50-degree inclination.

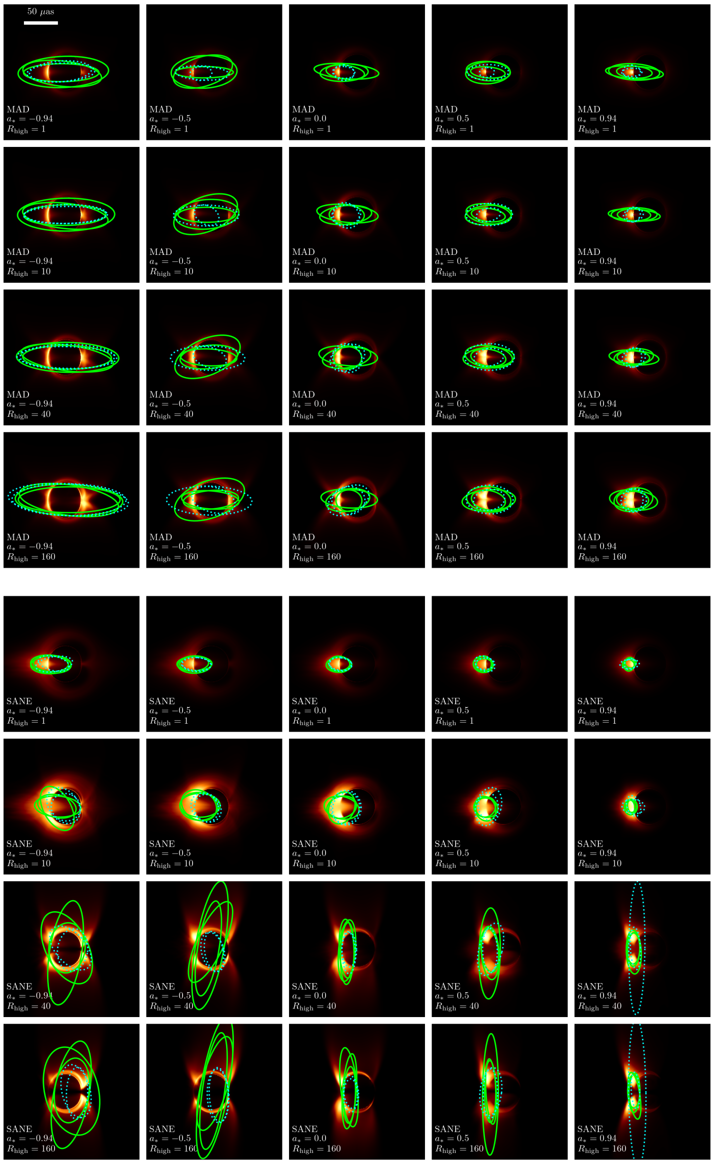

Figure B.10: Similar to Figure B.8 for 90-degree inclination.