



Quasars Are Not Light Bulbs: Testing Models of Quasar Lifetimes with the Observed Eddington Ratio Distribution

Citation

Hopkins, Philip F., and Lars Hernquist. 2009. "QUASARS ARE NOT LIGHT BULBS: TESTING MODELS OF QUASAR LIFETIMES WITH THE OBSERVED EDDINGTON RATIO DISTRIBUTION." *The Astrophysical Journal* 698 (2): 1550–69. <https://doi.org/10.1088/0004-637x/698/2/1550>.

Permanent link

<http://nrs.harvard.edu/urn-3:HUL.InstRepos:41381860>

Terms of Use

This article was downloaded from Harvard University's DASH repository, and is made available under the terms and conditions applicable to Open Access Policy Articles, as set forth at <http://nrs.harvard.edu/urn-3:HUL.InstRepos:dash.current.terms-of-use#OAP>

Share Your Story

The Harvard community has made this article openly available.
Please share how this access benefits you. [Submit a story](#).

[Accessibility](#)

QUASARS ARE NOT LIGHT-BULBS: TESTING MODELS OF QUASAR LIFETIMES WITH THE OBSERVED EDDINGTON RATIO DISTRIBUTION

PHILIP F. HOPKINS^{1,2} & LARS HERNQUIST¹

Submitted to ApJ, September 8, 2008

ABSTRACT

We use the observed distribution of Eddington ratios as a function of supermassive black hole (BH) mass to constrain models of quasar/AGN lifetimes and lightcurves. Given the observed (well constrained) AGN luminosity function, a particular model for AGN lightcurves $L(t)$ or, equivalently, the distribution of AGN lifetimes (time above a given luminosity $t(>L)$) translates directly and uniquely (without further assumptions) to a predicted distribution of Eddington ratios at each BH mass. Models for self-regulated BH growth, in which feedback produces a self-regulating “decay” or “blowout” phase after the AGN reaches some peak luminosity/BH mass and begins to expel gas and shut down accretion, make specific predictions for the lightcurves/lifetimes, distinct from e.g. the expected distribution if AGN simply shut down by gas starvation (without feedback) and very different from the prediction of simple phenomenological “light bulb” scenarios. We show that the present observations of the Eddington ratio distribution, spanning nearly 5 orders of magnitude in Eddington ratio, 3 orders of magnitude in BH mass, and redshifts $z = 0 - 1$, agree well with the predictions of self-regulated models, and rule out phenomenological “light bulb” or pure exponential models, as well as gas starvation models, at high significance ($\sim 5\sigma$). We also compare with observations of the distribution of Eddington ratios at a given AGN luminosity, and find similar good agreement (but show that these observations are much less constraining). We fit the functional form of the quasar lifetime distribution and provide these fits for use, and show how the Eddington ratio distributions place precise, tight limits on the AGN lifetimes at various luminosities, in agreement with model predictions. We compare with independent estimates of episodic lifetimes and use this to constrain the shape of the typical AGN lightcurve, and provide simple analytic fits to these for use in other analyses. Given these constraints, the average local BH must have gained its mass in no more than a couple of bright, near peak-luminosity episodes, in agreement with models of accretion triggering in interactions and mergers.

Subject headings: galaxies: evolution — cosmology: theory — galaxies: active — quasars: general

1. INTRODUCTION

Quasars and active galactic nuclei (AGN) are among the most luminous, energetic, and distant objects in the Universe. Simple integral arguments (Soltan 1982) make it clear that the supermassive black hole (BH) population was grown primarily through accretion in luminous AGN phases, and that the accretion luminosity released in these phases dominates the X-ray background and constitutes a large fraction of the bolometric energy production of the Universe. Comparison of e.g. the clustering (Croom et al. 2005; Porciani et al. 2004; Hopkins et al. 2007d) and host galaxy properties (Bahcall et al. 1997; Canalizo & Stockton 2001; Dunlop et al. 2003; Hopkins et al. 2006c; Zakamska et al. 2006, 2008) of high and low redshift AGN and galaxies demonstrates that AGN are the progenitors of modern-day spheroids.

Moreover, with the discovery of tight correlations between the masses of black holes and the velocity dispersion (Ferrarese & Merritt 2000; Gebhardt et al. 2000), masses (Magorrian et al. 1998), and perhaps most fundamentally binding energy or potential well depth (Hopkins et al. 2007c,b; Aller & Richstone 2007) of the host demonstrates a fundamental link between the growth of supermassive black holes and galaxy formation. A number of models have been developed arguing that the energy or mo-

mentum released from an accreting supermassive black hole, even if only a small fraction couples to the surrounding ISM, is sufficient to halt further accretion onto the black hole and drive away gas, self-regulating growth by shutting off the quasar³ and quenching star formation in the galaxy and therefore allowing it to redden rapidly (see e.g. Ciotti & Ostriker 1997, 2001; Silk & Rees 1998; Burkert & Silk 2001; Di Matteo et al. 2005; Hopkins et al. 2005d, 2006a; Murray et al. 2005; Sazonov et al. 2005; Springel et al. 2005b,a).

One of the most basic aspects of black hole growth, and a powerful test of these self-regulating models for AGN evolution, is the quasar/AGN lifetime. Given a sufficiently well-known lifetime distribution (lifetime as a function of e.g. luminosity, black hole mass, and other properties), the well-constrained quasar luminosity function (QLF) can be empirically translated (without invoking any models or additional assumptions) into the triggering rate of AGN as a function of e.g. luminosity, BH/host galaxy/dark matter halo mass, redshift, and other properties, as well as the active BH mass function, Eddington ratio and duty cycle distributions, and differential (mass-dependent) rate of buildup of the BH mass function.

Observations generally constrain *quasar* lifetimes to the range $\approx 10^7 - 10^8$ yr (for a review, see Martini 2004). These estimates are primarily based on demographic or integral

¹ Harvard-Smithsonian Center for Astrophysics, 60 Garden Street, Cambridge, MA 02138

² Department of Astronomy, University of California Berkeley, Berkeley, CA 94720

³ In what follows, we use the term “quasar” somewhat loosely, as a proxy for high-Eddington ratio accretion activity, rather than as a reference to specific optical properties. We use the term AGN to refer to BH accretion at all levels.

arguments which combine observations of the present-day population of supermassive black holes and accretion by the high-redshift quasar population (e.g., Soltan 1982; Haehnelt et al. 1998; Yu & Tremaine 2002; Yu & Lu 2004; Haiman et al. 2004; Marconi et al. 2004; Shankar et al. 2004), or incorporate quasars into models of galaxy evolution (e.g., Kauffmann & Haehnelt 2000; Wyithe & Loeb 2002; Di Matteo et al. 2003, 2004, 2008; Granato et al. 2004; Scannapieco & Oh 2004; Lapi et al. 2006; Hopkins et al. 2006a, 2008d; Sijacki et al. 2007) or reionization of HeII (Sokasian et al. 2002, 2003; Faucher-Giguère et al. 2008a,b; McQuinn et al. 2009). Results from clustering in quasar surveys (e.g., Porciani et al. 2004; Grazian et al. 2004; Croom et al. 2005; Myers et al. 2006, 2007; Lidz et al. 2006; Porciani & Norberg 2006; Shen et al. 2007; da Angela et al. 2008; Hopkins et al. 2007d), the proximity effect in the Ly α forest (Bajtlik et al. 1988; Haiman & Cen 2002; Yu & Lu 2005; Faucher-Giguère et al. 2008c) (but see also Lidz et al. 2007), and the transverse proximity effect (Jakobsen et al. 2003; Schirber et al. 2004; Worseck & Wisotzki 2006; Worseck et al. 2007; Gonçalves et al. 2008) similarly suggest lifetimes \sim a few $\times 10^7$ yr.

These observations, however, pertain in particular to the *quasar* lifetime – i.e. the characteristic time spent at high Eddington ratios/accretion rates, where much of the mass of a BH is accreted. Unsurprisingly, the observations suggest a lifetime similar to the Salpeter (1964) time (the *e*-folding time for Eddington-limited black hole growth) $t_S = 4.2 \times 10^7 (\epsilon_r/0.1)$ yr for accretion with radiative efficiency $\epsilon_r = L/\dot{M}_{\text{BH}}c^2 \sim 0.1$. AGN, however, are not a homogeneous population, and so the lifetime is not a single number – in general, it should be a function of luminosity and other parameters such as BH or host mass and (possibly) various physical effects. As advocated by Hopkins et al. (2005d,a), the AGN lifetime should properly be thought of as a *luminosity-dependent lifetime* distribution – i.e. some time $t(>L)$ as a function of L or differential $dt/d\log L$.

Hopkins et al. (2005d,a, 2006a,b); Hopkins & Hernquist (2006) study this luminosity-dependent lifetime/duty cycle distribution in hydrodynamic simulations and analytic models of feedback-regulated BH growth, and show that such self-regulation leads to a generic and unique predicted form for the lifetime distribution. After some initial trigger that fuels gas inflows (such as e.g. major and minor mergers or disk instabilities; for discussion see Di Matteo et al. 2005; Hopkins & Hernquist 2006; Hopkins et al. 2008c, 2009a,e, 2008a, 2009c,b; Hopkins & Hernquist 2009; Younger et al. 2008), AGN grow in approximately Eddington-limited fashion until reaching some critical mass where, if some small fraction of the radiant energy/momentum can couple to the surrounding ISM, the feedback is sufficient to halt inflows and/or expel gas and shut down future accretion. This “upper limit” to growth is essentially an Eddington limit effective at the scales where the host galaxy, rather than the BH, dominates the gravitational potential, and therefore is set not by the details of fueling mechanisms, but instead by (relatively generic) global parameters of the accretion physics such as the host galaxy mass and AGN luminosity.

That the resulting lifetime distribution is independent of fueling mechanism (i.e. is not specifically related to, for example, merger-induced fueling, but is a generic consequence of models where BH growth is self-regulated by feedback) has been demonstrated in the nearly identical lifetime dis-

tributions obtained by models of fueling in major mergers, minor mergers, flyby events, bar instabilities, and random “stochastic” encounters with nuclear molecular clouds, under similar feedback-regulated conditions (Hopkins et al. 2006a,b; Hopkins & Hernquist 2006; Johansson et al. 2009; Younger et al. 2008). These various fueling mechanisms do result in other important differences (in e.g. host properties and spectral properties of observed AGN), and they will lead to different evolution of BHs in a cosmological sense (global triggering rates and their evolution varying significantly for the mechanisms above) and produce BHs of different masses (to the extent that they produce bulges across a large range of masses, their self-regulating nature and the existence of the BH-host correlations ensures this to be the case). These are discussed in more detail in other papers (see e.g. Hopkins & Hernquist 2009), but the important point is that, for a given triggering rate at some redshift and BH mass interval (set by some cosmological or galactic processes), self-regulated models predict a similar effective Eddington limit and lightcurve.

Just as growth at the traditional Eddington limit leads to a self-similar solution for the AGN lightcurve – exponential growth – the expulsion of gas in this analogous limit leads to a self-similar lightcurve once the gas begins to be removed from the vicinity of the BH – a power-law decay of the form $L \propto t^{-(1.5-2.0)}$. In turn, this translates into a lifetime distribution $dt/d\log L$ with a characteristic faint-end (low- L) power-law like behavior: i.e. the time spent above a given luminosity/Eddington ratio (at fixed final BH mass) scales $\propto L^{-\beta}$ with a $\beta \sim 0.6$ at low- L , with a cutoff near the Eddington limit/peak luminosity of the system. The normalization is set by the characteristic timescales of the system – the Salpeter time and the (very similar) characteristic dynamical times in the central regions of the galaxy – naturally yielding a robust prediction of the observed quasar lifetime as well as the complete AGN lifetime distribution.

This is significantly different from what is assumed in commonly adopted phenomenological models, as well as other physical prescriptions. Quasars are often treated crudely as “light-bulbs” – i.e. assumed to be “on” for a fixed time (the “quasar lifetime”) at fixed luminosity or Eddington ratio, and otherwise “off.” In such a case, the lifetime/Eddington ratio distribution does not increase towards lower luminosities, but instead is strongly peaked about a characteristic high Eddington ratio/luminosity (strictly speaking, a δ -function; or in a Schechter-function parameterization with some scatter, $\beta \ll 0$). Similar results are obtained if one assumes that quasar lightcurves are pure exponentials (corresponding to growth at fixed Eddington ratio with either an instantaneous cutoff or time-mirrored exponential luminosity decay), which yields equal time spent per logarithmic interval in luminosity (i.e. a luminosity-independent lifetime, or $\beta = 0$ Schechter function).

More physically motivated but distinct models make their own predictions for the lifetime distribution. For example, if one assumes that the quasar accretion is regulated by a standard Shakura & Sunyaev (1973) thin disk, and the fuel supply is immediately removed but there is no feedback (i.e. a gas starvation scenario), one obtains a similarity solution for the accretion rate versus time that yields a lifetime distribution more akin to, but still significantly distinct from that predicted in self-regulated models (see e.g. Yu et al. 2005).

A number of indirect tests have been proposed to break the

degeneracies between these models and constrain the quasar lifetime distribution, and the observations in these scenarios have thus far supported the predictions of self-regulated models. These include: the dependence of quasar clustering on luminosity (Lidz et al. 2006; Hopkins et al. 2007d), where observations finding e.g. a weak dependence of clustering amplitude on luminosity (at fixed redshift) support lifetime models with more time spent at lower L (Adelberger & Steidel 2005; Coil et al. 2007; Myers et al. 2007; da Angela et al. 2008); the shape of the active BH mass function in various luminosity-selected samples (Hopkins et al. 2005b, 2006a, 2008d), including more massive systems at lower L (in extreme cases even being peaked) rather than tracing an identical shape to the QLF (Heckman et al. 2004; Greene & Ho 2007); the evolution in the faint-end QLF slope with redshift (Hopkins et al. 2006b), flattening (weakly) with redshift (Ueda et al. 2003; Hasinger et al. 2005; La Franca et al. 2005; Silverman et al. 2005; Hopkins et al. 2007e; Fontanot et al. 2007; Silverman et al. 2008b; Siana et al. 2008) as predicted in some self-regulated models owing to a (weak) dependence of quasar lifetime distributions on BH mass/peak luminosity; the shape of the distribution in quasar host galaxy masses/luminosities (Hamilton et al. 2002; Hopkins et al. 2006e), similar in nature to the active BH mass function as a test of lifetime models; the mass functions and clustering as a function of mass of quasar “remnants” (i.e. bulges/spheroids, given the observed $M_{\text{BH}} - \sigma$ relation; see Hopkins et al. 2006c, 2008b, 2007a; Bundy et al. 2006; Haiman et al. 2007; Shankar et al. 2009a; Yu & Lu 2008); and the relation between observed luminosity functions in different bands and active BH masses (Hopkins et al. 2005c; Shankar et al. 2009b; Merloni & Heinz 2008; Bundy et al. 2008; Yu & Lu 2008).

Although these observations are consistent with the predictions from self-regulated models, they are indirect, and as such are not able to rule out alternative interpretations (in the case of e.g. clustering or the evolution of the faint-end QLF slope) or depend on additional (albeit observationally and physically-motivated) assumptions. Moreover, many are restricted to relatively bright Seyfert/quasar populations, where the models are all similar (near these luminosities, they all predict that the entire population must be dominated by relatively massive BHs at high Eddington ratios $\sim 0.1 - 1$, and given the rapid growth at these Eddington ratios, the lifetime in this regime must be similar, comparable to the Salpeter (1964) time). Biases introduced (selecting for specific Eddington ratios, for example) will also be important in samples selected by e.g. broad emission lines or optical/UV/IR colors/spectral shape (see Hopkins et al. 2009d).

The observed Eddington ratio distributions – specifically, the complete distribution of Eddington ratios for *all* BHs of a given mass (obscured or unobscured, luminous or underluminous), however, represent a direct and powerful test of these models, with the predicted behavior at lower Eddington ratios/luminosities being strongly divergent. Given the QLF, a lifetime model directly and uniquely translates (without any additional assumptions) into a distribution of Eddington ratios at each BH mass, and vice versa. The requirements are demanding: breaking these degeneracies necessitates large samples, complete to all objects of a given BH mass, in a large volume (to constrain rare high-Eddington ratio objects), but sufficiently deep to measure even faint levels of AGN activity in those objects. Furthermore, a probe of AGN activity such as X-ray or narrow-line emission, robust to obscuration ef-

fects (or the possible disappearance of the broad-emission line region and/or thin disk at low Eddington ratios) is important both to obtain a complete census of AGN activity and to avoid biases from e.g. luminosity or Eddington ratio-dependent obscuration/dilution/SED shapes. Fortunately, with the advent of wide area spectroscopic surveys such as the SDSS, it has become possible to constrain the Eddington ratio distribution at low redshifts, over a sufficiently large dynamic range to break the degeneracies between these models, as a function of various AGN and galaxy properties.

Here, we combine a large number of observations of AGN Eddington ratio distributions as a function of BH and host galaxy mass, AGN luminosity, and redshift, in order to test these models and tightly constrain critical quantities such as the quasar lifetime as a function of luminosity, and show how the present observations are already sufficient to rule out, at high significance, a wide variety of alternative, simplified physical and phenomenological AGN lifetime/lightcurve models.

In § 2 we compare these model predictions with observations of the Eddington ratio distribution measured directly at $z = 0$ over a range of BH mass, and inferred indirectly at $z = 0 - 1$. In § 3 we similarly compare with Eddington ratio distributions measured not at fixed BH mass, but at fixed AGN luminosity, again over the range $z = 0 - 1$. In § 4 we show how these observations tightly constrain physical and phenomenological models for the quasar lifetime/accretion rate distribution, relate to a possible dependence of quasar feedback efficiency on mass, and rule out a number of alternative lifetime/lightcurve models. We show how the observations tightly constrain even general, parameterized lifetime models to a narrow range about the physical models, and can directly be converted to yield the cosmologically integrated AGN lifetime and $z = 0$ duty cycles as a function of Eddington ratio. In § 6 we translate these Eddington ratio/lifetime distribution constraints to limits on the form of the “typical” AGN lightcurve, and discuss constraints on the “episodic” quasar lifetime and how, combined with the duty cycle constraints, this can give a bound on the number of accretion episodes per AGN and shape of the typical lightcurve. In § 8 we demonstrate the constraints from these observations on how much mass present-day BHs accreted in various intervals in Eddington ratio and luminosity. Finally, in § 9 we discuss our results, the implications of the observations for a broad range of AGN properties, and the prospects for future observational tests.

For ease of comparison, we convert all observations to bolometric luminosities given the appropriate bolometric corrections from Hopkins et al. (2007e) (see also Elvis et al. 1994; Richards et al. 2006). We adopt a $\Omega_{\text{M}} = 0.3$, $\Omega_{\Lambda} = 0.7$, $H_0 = 70 \text{ km s}^{-1} \text{ Mpc}^{-1}$ cosmology and normalize all observations and models appropriately (note that this generally affects only the exact normalization of quantities here, not the qualitative conclusions, and differences are negligible within the range of cosmologies allowed by present constraints; e.g. Komatsu et al. 2009).

2. COMPARING WITH COMPLETE EDDINGTON RATIO DISTRIBUTIONS

Given the QLF $\Phi(L|z)$ and some model for the quasar lifetime or differential time at different Eddington ratios:

$$\frac{dt}{d \log L}(L|M_{\text{BH}}) \equiv \frac{d}{d \log L} [t(L' > L|M_{\text{BH}})], \quad (1)$$

it is straightforward to de-convolve and determine the Eddington ratio distribution. For example, if quasars were “light

bulbs” that shine either at some fixed high Eddington ratio

$$\lambda \equiv \frac{L}{L_{\text{Edd}}} = \frac{L_{\text{bol}}}{1.3 \times 10^{38} \text{ erg s}^{-1} (M_{\text{BH}}/M_{\odot})} \quad (2)$$

with $\lambda_{\text{on}} \sim 1$ in an “on” state for a time t_Q and $\lambda \ll 1$ in an “off” state at other times, then the implied lifetime distribution $dt/d\log L$ is a delta-function at λ_{on} . The observed QLF in such a case is, of course, a linear translation of the active BH mass function (with λ_{on} determining the re-normalization or shift in the horizontal axis, from M_{BH} to L , and the absolute value of the quasar lifetime t_Q determining the abundance/vertical axis). So, at a given M_{BH} , the Eddington ratio distribution will be a delta-function at λ_{on} with a normalization/fractional abundance at this λ_{on} determined by the number of active quasars (of the corresponding luminosity) at the observed redshift.

In general, in fact, so long as the quasar lifetime at a given L is short compared to the Hubble time (i.e. cosmological evolution in e.g. triggering rates can be ignored in the constraint), then the Eddington ratio distribution at a given M_{BH} will be $dt/d\log L$ (modulo a normalization reflecting the “on” fraction or t_Q/t_H , where t_H is the Hubble time at the given redshift) – i.e. we can trivially translate to a “duty cycle” distribution (fractional population at each L or λ):⁴

$$\frac{d\delta}{d\log \lambda} = \frac{d\delta}{d\log L} \approx \frac{1}{t_H(z)} \frac{dt}{d\log L} \quad (t \ll t_H). \quad (3)$$

By definition of the duty cycle, of course, this directly relates to the actual number density function of BHs at a given mass and luminosity or accretion rate

$$\Phi(\lambda | M_{\text{BH}}) \equiv \frac{dn(\lambda, M_{\text{BH}})}{d\log \lambda} = n(M_{\text{BH}}) \frac{d\delta}{d\log \lambda}. \quad (4)$$

The shape of the observed Eddington ratio distribution, therefore, contains information about the shape of $dt/d\log L$ independent of either normalization, and vice versa. Such observations hence provide a useful and direct probe of quasar accretion rate distributions.

If the AGN population (e.g. triggering rates or number of bright objects) is relatively constant for systems of a given mass over the redshift range of interest (i.e. if the population is still growing, at least in a statistical sense, around the observed redshifts) and the lifetime is short relative to the Hubble time, then Equation 3 is applicable, and the observed Eddington ratio distribution, independent of any other constraints, can be directly translated to the lifetime distribution. When the triggering rate evolves strongly with redshift and/or the characteristic lifetime is long compared to the Hubble time (e.g. for massive systems at low redshift, where their growth is dominated by higher-redshift periods), the Eddington ratio distribution at a given z is still uniquely predicted by a given $dt/d\log L$ model, but this must be convolved over the redshift distribution of activity (i.e. given the observed QLF and $dt/d\log L$, we integrate over time to determine the predicted $d\delta/d\log \lambda$). Even this, in practice, does

⁴ The translation between the distribution in $\log L$ and $\log \lambda$ is trivial in terms of observations at a given BH mass; in terms of converting model predictions between one and the other, although the relation is not completely trivial, we find in practice that the two are nearly equivalent, especially at Eddington ratios $\lesssim 0.2$ where most of the data with which we will compare lie, since some initial M_{BH} does not change much as a model system moves through a low Eddington ratio phase. We will therefore use the two interchangeably in this paper as is proper for the observations, but have converted all physical models to the appropriate representation.

not significantly change the direct mapping between the shape of the observed Eddington ratio distribution and $dt/d\log L$; it mainly amounts to deriving a more correct effective “duty cycle” (normalization) than multiplying by $\sim 1/t_H$. For more discussion and details of the relevant equations, we refer to Yu & Lu (2004); Hopkins et al. (2005a, 2006a,b); for now, we note that the observed QLF $\Phi(L|z)$ is well-constrained at all the luminosities and redshifts of interest (since the relevant observations are primarily at low redshifts $z \lesssim 1$), and so any given model for $dt/d\log L$ *uniquely* translates to an Eddington ratio distribution as a function of BH mass or AGN luminosity. We wish to compare these distributions to the observations.

In order to do so, we want to begin with a complete distribution of BH accretion rates at a given BH mass. It is important to do so – in models where e.g. systems of a given mass can have a broad range of Eddington ratios, the distribution of Eddington ratios at a given BH mass can be *qualitatively* very different from the distribution at a given luminosity (see § 3). Moreover, if what is desired is a distribution at all possible Eddington ratios, then there is no useful definition of an “active” AGN: what we really desire is to begin with a complete census of all BHs of a given mass (active or inactive at any level), and to measure the Eddington ratio distribution within this sample.

Fortunately, the existence of a tight correlation between host galaxy luminosity/stellar mass/velocity dispersion and BH mass makes this possible: Heckman et al. (2004) and Yu et al. (2005) select complete samples of all SDSS galaxies with a narrow range in velocity dispersion σ (and corresponding narrow range in BH mass), and then examine this sample to a limiting depth for narrow equivalent width AGN features (following the methodology and classifications in Kauffmann et al. 2003; Kewley et al. 2006). Given a bolometric correction, this allows a complete census of all activity at a given BH mass down to some well-known limit in λ .

Moreover, the use of narrow lines is helpful for three reasons, as opposed to e.g. use of optical broad lines or an optical/IR color cut in identification of AGN and determination of their bolometric luminosities. First, it can probe very faint AGN and does not introduce much bias in terms of e.g. the risk of the host light in brighter systems diluting the AGN. Second, it allows us to include the obscured/Type 2 population (the abundance of which may depend on luminosity; see e.g. Ueda et al. 2003; Simpson 2005; La Franca et al. 2005; Barger & Cowie 2005; Beckmann et al. 2006; Bassani et al. 2006; Gilli et al. 2007; Hasinger 2008). Third, it should enable us to identify quasars even in states of moderate radiative inefficiency: there is growing evidence that AGN at low Eddington ratios may transition to a radiatively inefficient state characterized by the absence of a thin disk. Such objects are still accreting but appear primarily as X-ray (because a hot corona survives this transition) and narrow-line sources, rather than broad-line or optical continuum sources (see e.g. Narayan & Yi 1994; Narayan et al. 1995, 1996; Meier 2001; Maccarone et al. 2003; Yuan & Narayan 2004; Jester 2005; McClintock & Remillard 2006; Cao & Xu 2007). We discuss these distinctions in more detail in a companion paper (Hopkins et al. 2009d), and outline how they can influence e.g. the Eddington ratio distribution determined via various selection criteria. For our purposes here, however, the data adopted either avoid these uncertainties owing to their selection/identification methodology, or cover a luminosity range where these concerns are not important (see also

Heckman et al. 2004; Kauffmann et al. 2003).

Figures 1-2 show this analysis. We plot the Eddington ratio distribution determined in Yu et al. (2005) and Heckman et al. (2004) (see also Kauffmann & Heckman 2008), and compare to the theory of Hopkins et al. (2006a,b, 2007e). We consider three versions of the estimates to show the range of theoretical uncertainty inherent in the model, but the differences are, for the most part, minimal. First, we plot the Hopkins et al. (2006b) fitted $dt/d\log L$ distribution to their typical $\sim M_*$ BH simulations, multiplied by the number density of active BHs at each mass (the effective duty cycle, discussed in § 4) and divided by the Hubble time (since we are considering the fraction of active BHs at each λ ; this is shown as a the solid red line). This is an accurate approximation to a much more complete cosmological calculation so long as the quasar lifetime is short compared to the Hubble time and provided that there is not some strong feature in the redshift history of triggering (i.e. so long as the cosmological evolution around the redshift of interest is relatively weak). At low λ , this is simply a power-law; inevitably, the point where the lifetime nears the Hubble time will be reached, and more full cosmological models will turn over.

Second, we adopt the fits from Hopkins et al. (2007e) to the observed bolometric QLF as a function of redshift, using the lifetime distributions as a function of BH mass fitted in Hopkins et al. (2006b), and integrate these over redshift. In Hopkins et al. (2007e) the authors quantify the range of fits allowed given both the uncertainty in the lifetime model and the observed QLF; we show two lines that bracket this range (the dotted and dashed lines in Figures 1-2). Again, these use the same model lightcurves (the power-law like fits to the lightcurve shape as a function of BH mass, from Hopkins et al. 2006b) and are matched to the same QLF (the data compiled in Hopkins et al. 2007e). However, there remain uncertainties in the data and degeneracies in the fit – the models shown bracket the range in the λ distribution allowed by this set of models and data. The predicted distributions turn over at low λ in both cases (albeit at slightly different λ , reflecting these degeneracies in fitting the low-luminosity population), as the lifetime nears the Hubble time (obviously, in a full cosmological model, they must turn over, so that the “duty cycle” of a given BH across all Eddington ratios integrates to unity).

The theoretical predictions all agree well, and agree remarkably well with the observations. We stress that these are *predictions*; no quantity has been fitted to the data points (both the shape and normalization of the model curves are entirely predicted by the papers above).

The theoretical curves also agree with one another. This is because, at most of the luminosities of interest, the lifetime is well below the Hubble time and the (especially low-mass) BHs are still growing as a population (see e.g. Hasinger et al. 2005), so subtleties of cosmological evolution are irrelevant (the only quantity of interest is $dt/d\log L$). Moreover at low redshifts $z \lesssim 2$ the observed QLF is well-constrained, so there is little freedom given a $dt/d\log L$ model in what the distribution can be. The only differences appear at the highest masses – this is because, in a complete cosmological calculation, most of the triggering of these systems occurred preferentially at high redshifts, so there is some pile-up at low Eddington ratios (precisely below where the condition of $t_Q \sim t_H$ begins to be satisfied, as expected).

Figure 2 also compares the Eddington ratio distribution determined in Hopkins et al. (2006d) from the combined ob-

served samples of Marchesini et al. (2004) and Ho (2002). The range in BH mass in this sample is less narrowly constrained, $M_{\text{BH}} \sim 10^{7.5} - 10^{8.5} M_{\odot}$, but the data survey extremely faint systems in the radio and X-rays, and allow us to extend the observed Eddington ratio distribution from the already deep $\lambda \sim 10^{-4} - 10^{-3}$ in the Yu et al. (2005) and Heckman et al. (2004) samples by another two orders of magnitude. (Note in Hopkins et al. (2006d) the authors consider possible corrections for varying radiative efficiencies and bolometric corrections at these luminosities, but the resulting change is for our purposes here within the error bars shown.) Again, the agreement is good over the entire range.

At higher redshifts, analogous measurements are not, at present, available. However, we can compare with an alternative, albeit indirect, observational estimator. It has been argued that the combination of radio, X-ray, and optical luminosities can be used to constrain the Eddington ratios of BHs, in a manner like that well-established for accreting X-ray binaries. Indeed, it is increasingly established that radio-loudness of AGN appears to be a function (on average) of Eddington ratio (Narayan et al. 1995, 1996; Falcke & Biermann 1996; Meier 2001; Ho 2002; Merloni et al. 2003; Marchesini et al. 2004; Maccarone et al. 2003; Falcke et al. 2004; Merloni & Heinz 2007; Greene et al. 2006, for a review see Fender et al. (2007)), so if this is true at higher redshift it can be used as at least a statistical estimator of Eddington ratio. Merloni & Heinz (2008) adopt these observations, combined with the measured X-ray, radio, and optical luminosity distributions of observed quasars, to constrain the Eddington ratio distribution as a bivariate function of BH mass and redshift. Their methodology allows for intrinsic scatter in these correlations, so it should be reasonably robust as long as there is some physical relationship between Eddington ratio and X-ray-radio-optical spectral shape over a wide baseline in Eddington ratio (from $\sim 10^{-5} - 1$).

Figure 3 compares their inferred Eddington ratio distributions (A. Merloni, private communication; for details see Merloni & Heinz 2009) to the same predictions, at $z = 0$ and $z = 1$ (at higher redshifts, Merloni & Heinz (2008, 2009) do infer Eddington ratio distributions, but the radio luminosity function is not directly measured and the X-ray luminosity function is increasingly unconstrained as well, so this relies on extrapolation of the low-redshift trends, and is not a direct measurement/estimate). Note that the kink in the Merloni & Heinz (2008) Eddington ratio distributions around $\lambda \sim 10^{-2}$ is sensitive to the particular assumptions about the form of the QLF shape and SED shape as a function of luminosity; other attempts to infer this distribution have, however, seen similar features (e.g. Marchesini et al. 2004), so it may reflect a real change in accretion properties not modeled here. In any case, the inferred distributions, where they overlap with the measurements from the observations above, agree reasonably well, giving some confidence in this methodology, and the agreement with theoretical predictions is good.

As discussed in § 1, although the particular models shown were fitted from simulations of galaxy mergers, the results are comparable regardless of fueling mechanism, given a similar self-regulation from local AGN feedback. In fact, the predictions and their agreement with observations in Figures 1-3 rely on *no* information regarding fueling mechanisms – they simply follow from assuming a given lightcurve shape and matching the observed quasar luminosity function.

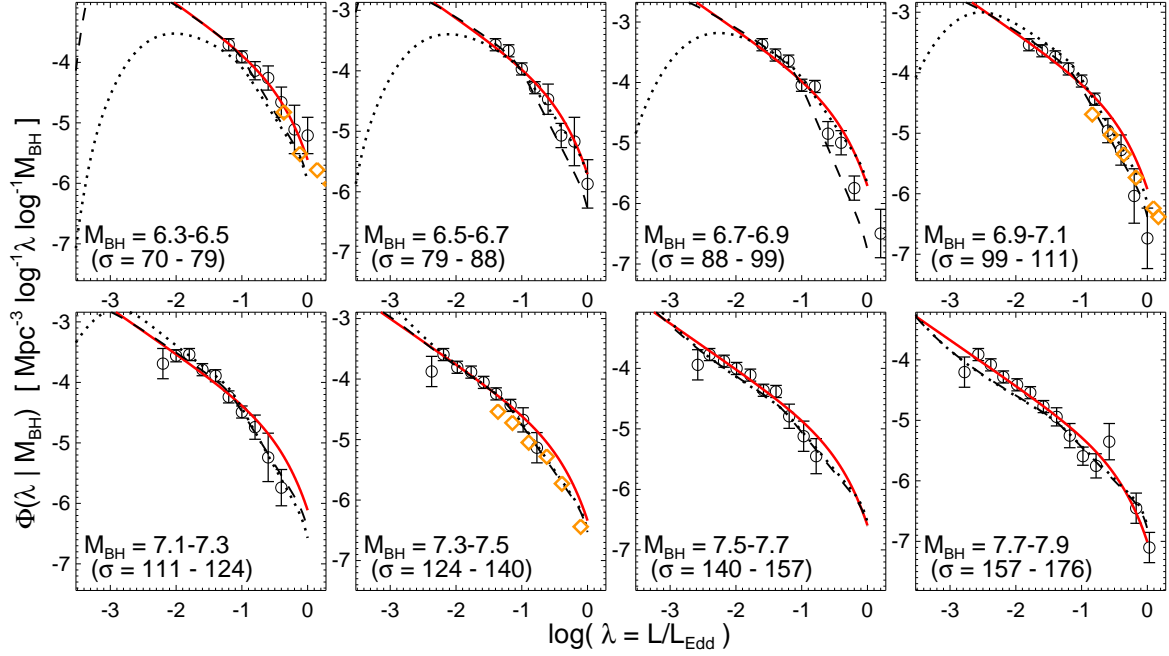


FIG. 1.— Distribution of Eddington ratios at a given BH mass (at $z < 0.2$). We compare the observed distribution from SDSS narrow-line objects (Yu et al. 2005; Heckman et al. 2004, black circles with error bars and orange diamonds, respectively) to the distribution predicted by the lightcurve/lifetime models in Hopkins et al. (2006a,b, 2007e, lines). Solid red line is the simplest model prediction for a population of single triggers with $t_Q \ll t_H$; black lines integrate over a complete cosmological history of triggering events in the model (constrained to match the observed AGN luminosity functions). Dashed and dotted lines bracket the model uncertainty. The lines are *predictions* – there are *no* free parameters fitted or tuned to match the observed λ distributions.

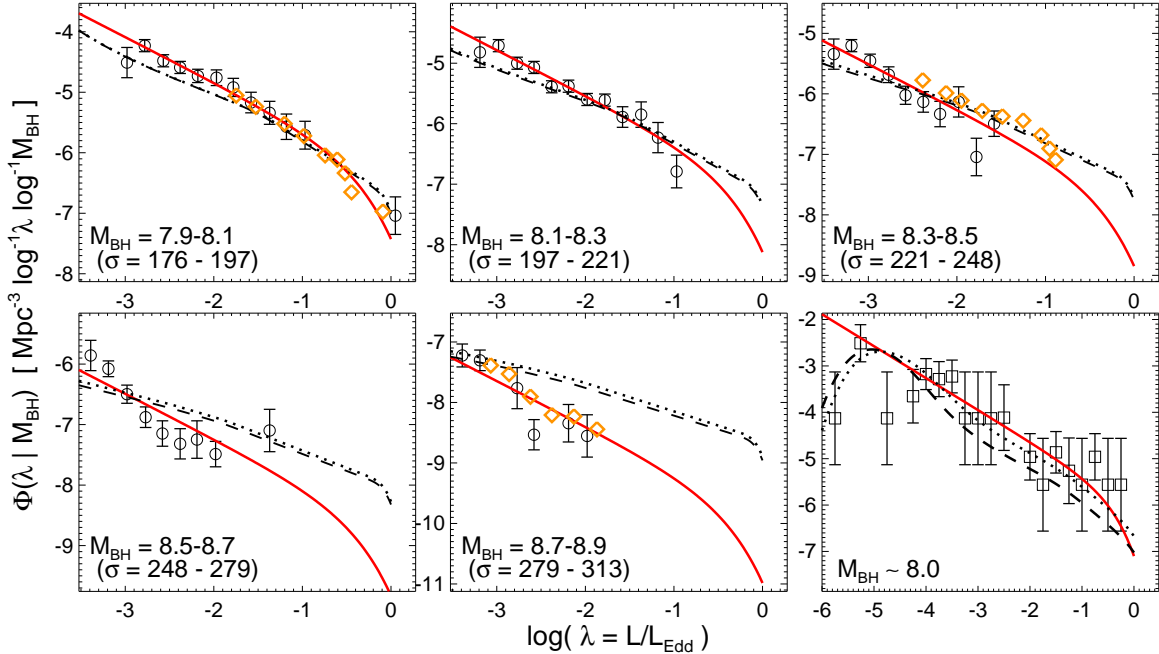


FIG. 2.— As Figure 1, continued to higher BH masses. The lower right panel shows the Eddington ratio distribution calculated in Hopkins et al. (2006d) from the observed samples of Marchesini et al. (2004) and Ho (2002), with a less tightly constrained BH mass range ($M_{\text{BH}} \sim 10^{7.5} - 10^{8.5} M_{\odot}$) but extending to lower Eddington ratios. Corrections for various changes in radiative efficiency with Eddington ratio change the observed points and models within the plotted error bars (see Hopkins et al. 2006d). Note that the apparent discrepancy in the models at the highest masses comes from extrapolation to unobservably small space densities.

3. EDDINGTON RATIO DISTRIBUTIONS AS A FUNCTION OF LUMINOSITY

At a given luminosity (as opposed to a given BH mass), the differences in the Eddington ratio distribution between various models are greatly suppressed. The reasons for this are obvious: at a given L and M_{BH} , there is only a narrow range

of λ , and given the declining number density of high-mass BHs, at high- L , one will increasingly be limited to the population of near-Eddington systems. However, over a sufficiently large baseline in L , differences are apparent, and (especially at high redshifts) luminosity-limited samples are more easily constructed than M_{BH} -limited samples. We therefore consider




FIG. 3.— As Figure 1, but comparing the model predictions with the observationally inferred Eddington ratio distribution from the combination of observed X-ray, radio, and optical luminosities in Merloni & Heinz (2008, 2009) and the statistical correlation between e.g. radio-loudness and accretion rates (for a review see Fender et al. 2007). The results are shown over the plotted mass bins, at each labeled redshift. The sharp kinks in the Merloni & Heinz (2009) distributions are somewhat model-dependent.




FIG. 4.— Predicted distribution of Eddington ratios at $z \sim 1$ in a narrow range of bolometric luminosity. Note that, because of the cut at a given luminosity (as opposed to reflecting all λ at a given M_{BH} as in Figure 1), the distributions are much more narrow and turn over at low L_{bol} . The distributions are reasonably approximated as log-normal, with some weakly luminosity-dependent skewness.

the Eddington ratio distribution predicted as a function of luminosity at different redshifts.

Figures 4 & 5 shows the results. Figure 4 demonstrates that at a given luminosity (unlike at a given BH mass), the Eddington ratio distribution is expected to be something like a lognormal distribution (more so if typical observational errors are included), albeit with some non-negligible skewness (as seen in e.g. Fine et al. 2008).⁵ This is because, at extremely low λ , arbitrarily high- M_{BH} BHs would be implied (in a bin

⁵ The skewness originates because the distribution of host BHs is not flat in mass, but increases to lower masses (following the Schechter galaxy/spheroid mass function). With respect to the median Eddington ratio/BH mass contributing to the observed population at a given luminosity, one therefore expects (if Eddington ratio distributions at a given BH mass do not change rapidly with mass) that there will be a somewhat larger population of low-

of fixed L), but the possible population of such systems is vanishing. We therefore find it convenient to approximate the predicted distributions as lognormal, and quantify the median λ and 1σ dispersion (technically based on the IPV width to prevent bias from outliers or skewness) in the lognormal as a function of luminosity, at redshifts $z = 0 - 2$.

We compare these with several observational estimates. In Fine et al. (2008), the authors consider the Type 1 quasar population near $z \approx 1$ in the 2dF survey, and estimate the distribution of BH masses in narrow bins of luminosity employing the commonly adopted virial BH mass estimators (based on the broad-line widths and the radius-luminosity relations inferred from reverberation mapping of nearby AGN; see e.g. Vestergaard & Peterson 2006, and references therein). This allows them to consider the distribution of BH masses via this proxy as a function of luminosity down to near Seyfert luminosities. Because the width of the distribution can be determined without relying on the (systematically still uncertain) absolute normalizations of these calibrators, the authors decline to estimate absolute Eddington ratios (although a rough estimate suggests they lie between $\sim 0.1 - 1$, as predicted here). Kollmeier et al. (2006) use the same technique over a narrow luminosity range of Type 1 AGN in the AGES survey.

At lower luminosities, these indicators are less useful (and the observations suggest the population is both more obscured and diluted by host galaxy light, making the virial mass estimators inaccessible). However, X-ray observations can probe Type 2 objects in this regime, where the optical light is dominated by the host galaxy and therefore a host galaxy stellar mass (and corresponding BH mass, adopting the observed $M_{\text{BH}} - M_*$ relation from Marconi & Hunt 2003) can be estimated. It is well-established that in this regime the optical luminosity/stellar mass of the galaxy is approximately con-

stant for low luminosities, and that the population is dominated by low-mass BHs at high Eddington ratio rather than high-mass BHs (an exponentially vanishing population) at low Eddington ratio.

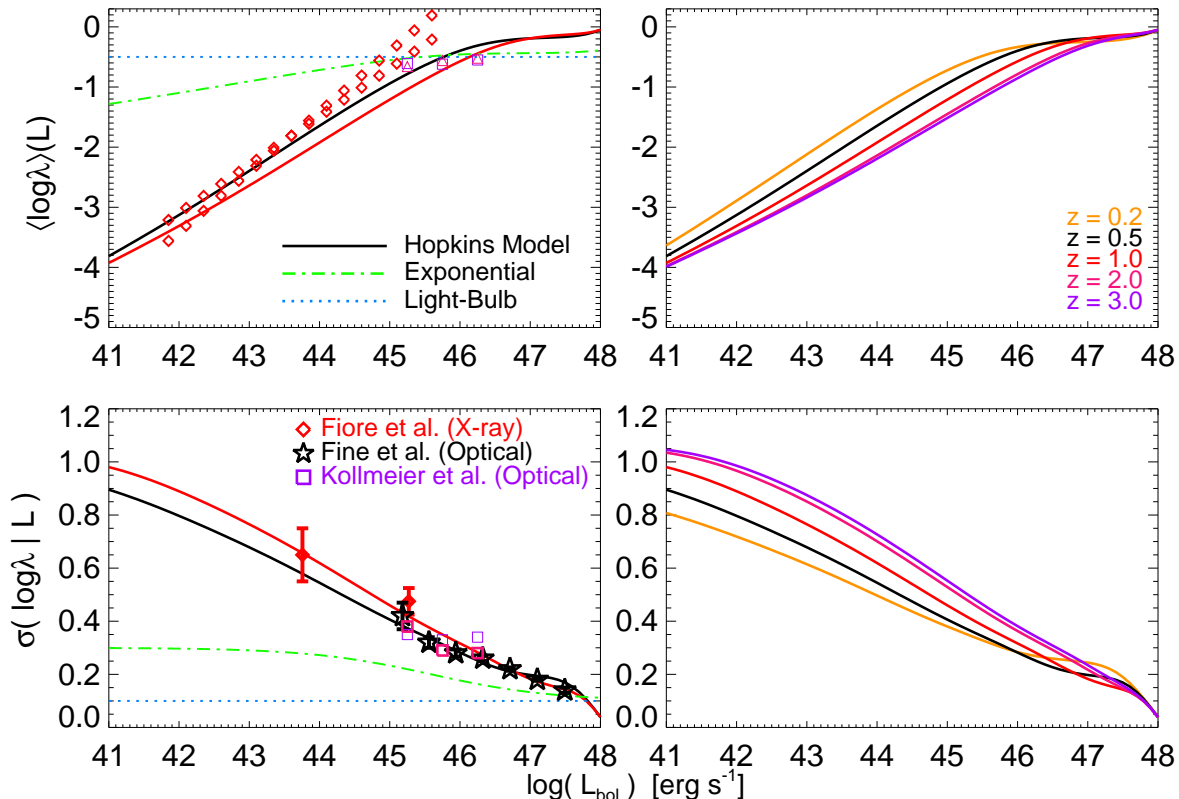


FIG. 5.— Median Eddington ratio λ and 1σ dispersion in Eddington ratios at a given narrow range in bolometric AGN luminosity (as in Figure 4, assuming a lognormal distribution). We compare the Hopkins et al. (2006a,b, 2007e) model predictions (solid lines; colors denote redshift as labeled) with the observed distributions inferred from the distribution of X-ray to host luminosities in optically obscured AGN (Fiore et al. 2003; Hasinger 2008; Hickox et al. 2007, 2009, red diamonds) (translating host luminosity to average BH mass given the observed correlations), and from the distribution fitted to broad-line optical samples in Fine et al. (2008) from the 2dF and Kollmeier et al. (2006) from AGES, using the optical virial (line-width) BH mass estimators. We also compare with the predictions for pure exponential AGN lightcurves (constant Eddington ratio or exponential decay in L after some peak) and “light-bulb” models (where AGN are “on” or “off” with a mass-independent narrow λ distribution when on), forced to obey the (necessary) constraint of matching the observed AGN luminosity functions. The Hopkins et al. (2006b) models agree well with the observations, the other models are ruled out at low- L . High- L optical samples are not ideal for breaking the degeneracies between these models, although with sufficiently large dynamic range such as that in Fine et al. (2008) the distinction can be seen.

stant while the X-ray luminosity changes, implying that at lower- L the X-ray luminosity function becomes increasingly a sequence in Eddington ratio.

We plot the implied Eddington ratios and distribution in Eddington ratios as a function of luminosity from the sample of Hasinger (2008), where the optical R -band luminosity is converted to a stellar mass based on the age and mass-dependent observed mean M/L ratios in Bell & de Jong (2000) and Bell et al. (2003). We have re-calculated these comparisons using e.g. the samples of Fiore et al. (2003) and Hickox et al. (2007, 2009) and obtain the same result (various multiwavelength surveys have reached similar conclusions regarding this correlation), and find that changing the assumed host M/L within uncertainties makes little difference (for more discussion, see Hopkins et al. 2009d). For convenience to the comparison here, we convert all the observed AGN luminosities to bolometric luminosities using the bolometric corrections in Hopkins et al. (2007e) (using instead those in Elvis et al. 1994; Marconi et al. 2004; Richards et al. 2006, makes no difference).

In addition to the theoretical predictions from the models in § 2, we contrast the results from a simple light-bulb AGN lifetime (in which both the mean Eddington ratio and dispersion are constant) and a pure exponential model (in which $dt/d\log L$ is constant at $\lambda \ll 1$, which, when the sample is cut by luminosity, does introduce some dependence on L ,

but much weaker than that predicted by the models in § 2). The observations clearly prefer the steeper dependence of the more realistic lifetime models. Note that a large baseline in luminosity is needed to see the difference at high significance – at least $\sim 2 - 3$ orders of magnitude in L below L_* (ideally more like $\sim 4 - 5$ orders of magnitude). This explains e.g. the weak dependence seen in the sample of Kollmeier et al. (2006), who note the weak dependence of mean Eddington ratio and width of the distribution on L , but have a relatively narrow range in L and are concentrated near and above $\sim L_*$. Their observations are in fact entirely consistent with the model predictions and other observations (over larger baselines) that do see such a dependence.

4. COMPARISON WITH MODELS: WHAT DEGENERACIES ARE BROKEN?

We now ask how unique these predictions are: in other words, can the observations distinguish between different models for the Eddington ratio/quasar lifetime distribution?

Figure 6 compares the data and several simple models for the lifetime distribution. There are a number of commonly adopted forms for quasar lifetimes and lightcurves in the literature, including the light-bulb, pure exponential growth (corresponding to growth at fixed Eddington ratio and then either instantaneous decline or similar decay), and self-regulated decay models. The range of possibilities can be generally ap-




FIG. 6.— As Figure 1, but comparing the consequences of different lightcurve/lifetime models for the observed Eddington ratio distribution (generally parameterized by Equation 6, with power-law like slope $-\beta$ at low λ). We compare our previous model examples (red solid, with median $\beta \sim 0.6$) and the exponential (green dot-dashed, $\beta = 0$) and light-bulb (blue dotted, $\beta \ll 0$, reflecting a log-normal or δ -function distribution) models. We also show the effective physical upper limit (dark blue dashed, $\beta = 1$; $\beta > 1$ would result in BH growth that is divergent towards low- λ). The models are all normalized to give the same number of objects at $\lambda \gtrsim 0.1$; this is required in order to match the observed quasar luminosity function where it is well-constrained ($L \gtrsim 10^{45} \text{ erg s}^{-1}$). Even freeing this normalization (i.e. allowing the AGN luminosity function to be significantly different from that observed), however, still does not allow the exponential or light-bulb models to match the observations.

proximated by a Schechter fitting function:

$$\frac{dt}{d \log L} = t_0 \left(\frac{L}{L_{\text{peak}}} \right)^{-\beta} \exp(-L/L_{\text{peak}}) \quad (5)$$

where $L_{\text{peak}} \equiv \eta L_{\text{Edd}}$ with some $\eta \sim 1$. Equivalently, we can fit:

$$\frac{dt}{d \log \lambda} (M_{\text{BH}}) = t_0 \left(\frac{\lambda}{\eta} \right)^{-\beta} \exp(-\lambda/\eta). \quad (6)$$

This allows for the fact that there must be some physical (and observed) cutoff in the Eddington ratio distribution above $\lambda = 1$, but permits an arbitrary power-law like behavior at lower Eddington ratios (which can approximate any near-similarity solution for self-regulated lightcurve decay, or a power-law-like spectrum of triggering activity, as well as a rapid/exponential rise fall and even sufficiently rapid rise similar to a light-bulb). In practice, we find that fitting an equation of the form of Equation 6 to the data in e.g. Figures 1-3 provides a useful and statistically good description of the data.

In all mass intervals, we find a similar best-fit $\eta \approx 0.2 - 0.4$ – the exact value reflects the particular choice of functional form, but the general value describes the observed and expected cutoff in the λ distribution at $\lambda \sim 1$. We discuss variations in η below, but find that it cannot vary widely: much lower η exponentially suppresses the number of bright sources, and much higher η implies no Eddington limit, both in conflict with AGN luminosity functions and direct observations of the distribution of accretion rates in bright, broad-line systems (see e.g. Kollmeier et al. 2006; Greene & Ho 2007). In what follows, we find statistically identical results fitting the observations with a free η or fixed $\eta = 0.4$ (the degeneracy between e.g. η and β , within the range allowed by observations, is not strong).

As a consequence, the shape of the distribution is primarily contained in the slope β . Various lightcurve models make

differing specific predictions for this slope:

- **Light-Bulb Models:** Strictly speaking, in such a model $dt/d \log L$ is a delta function at the characteristic λ_{on} ; allowing for some finite width or measurement errors in λ , this can be approximated as e.g. a lognormal distribution or a Schechter distribution with $\beta \ll 0$ (i.e. a large negative β).

- **Exponential Models:** For pure exponential models (e.g. accretion at a constant Eddington ratio or exponential decay), one obtains $dt/d \log L \approx \text{constant}$, or $\beta = 0$.

- **Maximal low- λ Accretion:** At the opposite extreme, there is a physical limit $\beta < 1$; because the fractional contribution to BH mass growth from a range in $\log(\lambda)$ goes roughly as $\lambda \times dt/d \log L$. For $\beta \geq 1$, the total growth is both formally divergent and (even if there is some cutoff at low- λ) weighted towards the lowest- λ values. Constraints from the observed BH mass function and the Soltan (1982) argument (see e.g. Yu & Tremaine 2002; Salucci et al. 1999; Shankar et al. 2004; Marconi et al. 2004), as well as other indirect constraints (Hopkins et al. 2006d) imply that most BH growth cannot occur in extremely low Eddington ratio states.

- **Self-Regulated Models:** In feedback-regulated models, the energy coupled to the ISM which halts quasar accretion leads to a nearly self-similar power-law like decay of the quasar lightcurve, $L \propto t^{-1/\beta}$. Hydrodynamic simulations of quasars in galaxy mergers (Hopkins et al. 2006b) suggest values $\beta \sim 0.6$ for typical $\sim L_*$ galaxies, with a weak mass dependence. Analytic calculations in Hopkins & Hernquist (2006) demonstrate that for a range of assumptions regarding the fueling mechanisms (for e.g. stochastic or secular AGN fueling mechanisms), feedback coupling mechanism, timescale, efficiency, and gas properties around the BH, a range of $\beta \approx 0.3 - 0.8$ is possible; but simulations of these non-merger scenarios (Younger et al. 2008) suggest a similar, relatively narrow range of β independent of fueling mecha-

nism.

- **Isolated Accretion Disk/Gas Starvation Models:** Yu et al. (2005); Yu & Lu (2008) show that a similar power-law decay is expected for a thin α -disk (Shakura & Sunyaev 1973) abruptly cut off from any future fuel supply, when there is no feedback. The disk therefore slowly starves as gas exhausts by a combination of accretion and (possibly) star formation. The solution should be mass and host galaxy independent, with $\beta \approx 0.80 - 0.84$. The exact solution depends (weakly) in detail on e.g. how the opacity and viscosity of the disk vary as a function of density and temperature, but the authors show that the entire set of solutions (corresponding to the allowed range for observationally reasonable α -disks) lie within this narrow interval in β . In Figure 6, we omit this model for clarity, but the result lies between the self-regulated model and the physically maximal ($\beta = 1$) model.

- **BHs Trace Stellar Evolution:** It is possible that many low-luminosity systems are fueled by stellar mass loss from aging nuclear stellar populations. If the BH simply linearly traced this evolution (i.e. grew in this fueling-limited background with no feedback), the late-time accretion rate evolution would follow the stellar mass loss rate, giving a similar power-law solution with a steeper $\beta \approx 0.9 - 1.0$ (e.g. Norman & Scoville 1988; Leitherer et al. 1999; Bruzual & Charlot 2003; Ciotti & Ostriker 2007).

Fitting each BH mass bin in Figures 1-2 to an arbitrary function of the form in Equation 6 (free t_0 , η , β), we obtain the constraints on β shown in Figure 7. In fact, the constraints can be made stronger. Any fitted λ distribution, convolved with the BH mass function (implicit in Figures 1-2, or taken from observations following e.g. Marconi et al. 2004) must reproduce the observed AGN luminosity functions. In practice, this “anchors” the number density at high $\lambda \gtrsim 0.1$ (which dominate the QLF). Re-fitting the observed λ distribution, with the fit constrained to also reproduce the observed QLF (taken here from the compilation of a large number of observations in Hopkins et al. 2007e, but the results are not sensitive to the specific choice), we obtain the constraints in the top panel of Figure 7 (note that the errors shown in this case are not independent).

Parametrically, if we assume β is independent of M_{BH} (or consider a range near $\sim L_*$), we obtain $\beta \approx 0.6 \pm 0.05$, in good agreement with the theoretical predictions from the simulations of Hopkins et al. (2006b). If we allow for a dependence of β on M_{BH} , parameterized for convenience as

$$\beta = \beta_{7.5} + \beta' \log(M_{\text{BH}}/10^{7.5} M_{\odot}), \quad (7)$$

we obtain $\beta_{7.5} = 0.63 \pm 0.04$ and $\beta' = -0.29 \pm 0.08$ if we include the constraints from the observed QLF, and $\beta_{7.5} = 0.63 \pm 0.04$ and $\beta' = -0.11 \pm 0.09$ if we do not (the observed λ distributions themselves are insufficient to determine if β depends weakly on BH mass).

Figure 8 summarizes the predictions of the different models above. The observations prefer a narrow range of slopes: it is not possible to match the observed Eddington ratio distribution in light-bulb or exponential models (ruled out at $\gg 5\sigma$), and simple stellar wind models are ruled out at $\sim 4\sigma$.

The isolated accretion disk model fares somewhat better: Yu et al. (2005) show that it is consistent with the λ distribution observed for low-mass BHs, a result we confirm here, but it is inconsistent with the observations of typical $\sim L_*$ and more massive BHs at $> 4\sigma$. Yu & Lu (2008) show that such a model does improve upon the light-bulb model in simulta-

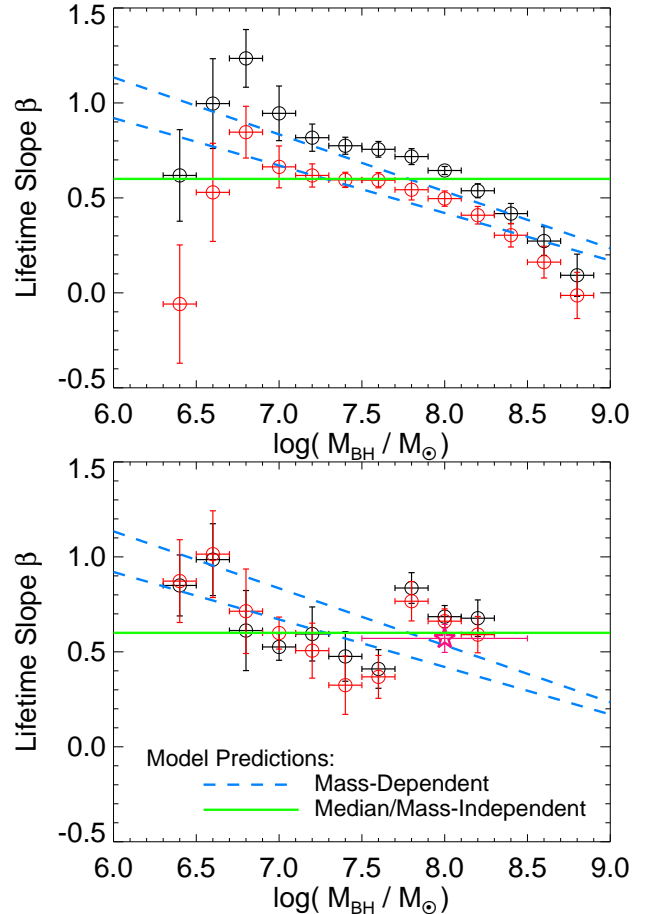


FIG. 7.— Slope β of the lifetime/lightcurve/Eddington ratio distribution fitted to the observations in Figures 1-2 with a general Schechter-function parameterization (Equation 6). *Top:* Points are the maximum likelihood fit results for each mass bin, where the fits are constrained such that the observed quasar luminosity function must be reproduced. We repeat our fits using the quoted error bars in the observed Eddington ratio distributions (black) and allowing for an additional ~ 0.3 dex intrinsic uncertainty (red). We compare with the model predictions from Hopkins et al. (2006b), determined from simulations allowing for a mass-dependent typical lifetime distribution or assuming a constant (mass-independent) lifetime distribution ($\beta \approx 0.6$). *Bottom:* Fits repeated, but removing the constraint that the observed luminosity function must be reproduced (magenta star is the fit to the observations in the lower-right panel of Figure 2, which span a wider mass and λ interval). The results are consistent, but less constrained. The results agree well with the model predictions, and if a match to the observed luminosity function is required, specifically favor the full mass-dependent model. In either case, an exponential ($\beta = 0$) or light-bulb ($\beta \ll 0$) model is ruled out at high significance.

neously reproducing the observed QLF and BH mass function. However, as noted before, the QLF is primarily a tracer of high- λ activity, and given that such activity dominates the accretion history of BHs (Soltan 1982), these observational constraints alone cannot break the degeneracy between similar, but slightly different lightcurve models $L \propto t^{-1/\beta}$ ($\beta \sim 0.6$ and $\beta \sim 0.8$ corresponding to the self-regulated and isolated accretion disk models, respectively). Hopkins et al. (2006a, 2007e, 2008d,b) demonstrate that similar solutions for the QLF and BH mass function exist for the self-regulated model (unsurprisingly, the two yield very similar growth histories and $\lambda \gtrsim 0.1$ activity). As a consequence the low- λ distribution represents an important means to break these degeneracies, and disfavor the isolated accretion disk models. We do emphasize, however, that while the isolated disk model is

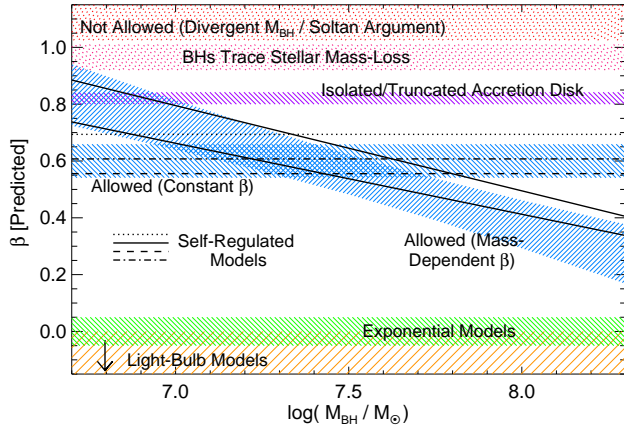


FIG. 8.— Illustration of how the observational constraints on β can discriminate between different physical and phenomenological quasar lifetime models. Blue shaded range shows the allowed slope β from the fits to the observations taken cumulatively, either assuming it is invariant (independent of M_{BH}) or allowing it to depend on M_{BH} as in Figure 7. Black lines compare different models for the self-regulating growth seen in hydrodynamic simulations (where e.g. the gas density profiles, ISM equation of state and phase breakdown, and prescriptions for feedback and BH accretion are varied; for details see Hopkins et al. 2006b; Hopkins & Hernquist 2006); these predictions agree well with the observations and are relatively independent of the detailed assumptions. If we include the constraint that the QLF must be reproduced simultaneously, the observations rule out mass-independent β (corresponding to slightly more simplified analytic self-regulated models, described in the text). The β corresponding to pure exponential or phenomenological “light-bulb” models are also shown, as is the region ruled out by the Soltan (1982) and other physical arguments. An alternative physical model, of a thin α -disk where the entire fuel supply is instantaneously provided from larger radii then cut off (i.e. large-scale gas inflows and feedback are excluded), is shown (this yields $\beta \approx 0.8$ as the accretion disk consumes gas; see Yu et al. 2005). Although a better description than the simplified phenomenological models, the observations can rule out such strictly “isolated” accretion disk exhaustion at $\gtrsim 3\sigma$ (the constraints are even stronger, $\sim 5\sigma$ at high masses, if we also force the solution to match the observed QLF). Likewise, assuming BH accretion simply traces stellar mass loss from evolving stellar populations (without feedback to regulate growth more efficiently) is ruled out ($\beta \approx 0.9 - 1.0$).

ruled out at high formal significance, the absolute difference between it and self-regulated models is not large. That modest difference may reflect the inclusion of a small, but physically important term, for example outflows or winds from the accretion disk that cause it to gas-exhaust slightly more steeply in time than otherwise.

Feedback-regulated models agree well with observations. Distinguishing between feedback-regulated models, however, is difficult. In Figure 8 we compare several sub-classes: all follow a similar behavior and reflect the basic scalings in Silk & Rees (1998). The models are derived in Hopkins & Hernquist (2006) treating feedback-driven outflows as similarity solutions for expanding winds/blastwaves; the solutions adopt various equations of state and make different assumptions for the feedback coupling (whether e.g. feedback instantaneously couples to the ISM or continues to couple continuously throughout the event). Others (Menci et al. 2003; Granato et al. 2004) have obtained similar conclusions in different feedback models. We also compare the fits to the hydrodynamic simulations in Hopkins et al. (2006b) and Younger et al. (2008) – the simulations give similar results for typical spheroids, but break the strict (analytically assumed) self-similarity: they predict a weak mass dependence (more massive galaxies being more gas poor and bulge-dominated, and having somewhat more violent resonant angular momentum transport and therefore more sharply peaked

quasar/starburst activity as a consequence). The mass dependence amounts to a predicted $\beta' \approx -0.25$ in Equation 7. The observations favor the self-regulated class of models, and may weakly favor the mass dependence predicted in simulations.

Some space of self-regulated models is ruled out. For example, Hopkins & Hernquist (2006) argue that if accretion grows primarily out of cold instabilities “falling out of” the expanding feedback-driven blastwave (in an adiabatic gas with Bondi-like spherical accretion), then very steep $\beta \sim 31/19$ at early times is expected (only decaying to the $\beta \sim 0.5 - 0.6$ seen in simulations in the late-time limit). Such a steep β (shallow lightcurve decay) is not allowed.

5. CONSTRAINTS ON AGN LIFETIMES

These fits also allow us to quantify the quasar lifetime. There are, in fact, three different “quasar lifetimes” to which we could refer, which are constrained to varying degrees by the observations.

(1) **“Effective” Lifetimes:** This is the lifetime defined in terms of the duty cycle – i.e. we define the effective quasar lifetime (at a given Eddington ratio λ at a given BH mass M_{BH}) by the duty cycle $d\delta(\lambda|M_{\text{BH}})/d\log\lambda$. Specifically, we invert Equation 3 to obtain

$$\frac{dt_{\text{eff}}(\lambda|M_{\text{BH}})}{d\log\lambda} \equiv t_H(z) \frac{d\delta}{d\log\lambda}, \quad (8)$$

or in terms of the integrated duty cycle above λ ($\delta(>\lambda)$), we have $t_{\text{eff}}(>\lambda) \equiv t_H(z) \delta(>\lambda)$. This is directly determined by the observed Eddington ratio distributions at any redshift z , and is what is often referred to as the implied “AGN lifetime” from various observed statistics.⁶

(2) **“Integrated” Lifetimes:** Recall, Equation 3 is only approximate. If one could view a BH or appropriate BH sub-population of final (remnant $z = 0$) mass M_{BH} over its complete history, then the *true* “integrated” lifetime or duty cycle at a given Eddington ratio or luminosity would be given by the appropriate integral

$$t_{\text{int}}(>\lambda) = \int \delta(>\lambda|z) \frac{dt}{dz} dz. \quad (9)$$

It is immediately clear by comparison with Equation 8 that the integrated lifetime is equal to the effective lifetime if the duty cycle $\delta(>\lambda|z)$ is relatively constant over the redshift range of interest. In practice, so long as $\delta(>\lambda|z)$ has evolved relatively weakly between the time when most BHs of mass M_{BH} were formed and the observed redshift, one will obtain $t_{\text{eff}} \approx t_{\text{int}}$. This is seen to be the situation for relatively low-mass BHs corresponding to e.g. low-luminosity broad-line AGN (which both are observed to have relatively constant number densities and duty cycles and are still growing, so their observed growth corresponds to that during their “growth epoch”; see e.g. Ueda et al. 2003; Hasinger et al. 2005; Hopkins et al. 2007e; Shankar et al. 2009b, and references therein). For more massive BHs, however, where δ is significantly lower today than at high redshift, the two will not be the same.

⁶ In standard parlance, t_{eff} is defined in terms of the Hubble time t_H . However, strictly speaking t_{eff} has physical meaning (i.e. represents the time systems spend “on” at some λ at a given redshift) only if two conditions are met: $t_{\text{eff}} \ll t_H$ (the duty cycle $\delta \ll 1$) and host properties/trigging rates are varying on a timescale $\gg t_{\text{eff}}$ (of order t_H). Both of these are satisfied at low redshift, for all but the lowest- λ populations, but may not be true at high redshifts (imagine, for example, all hosts are “just formed” at some high redshift and are also all “on” – then $t_{\text{eff}} = t_H$, but the real lifetime can be significantly shorter).

Nevertheless, the integrated lifetime is completely determined by the $z = 0$ Eddington ratio distributions for most reasonable physical models. The reason is that this is an integral constraint, and the total accretion needs to sum appropriately to produce a $z = 0$ BH of mass M_{BH} . This is trivial to see if we temporarily consider the differential lifetime in terms of luminosity (rather than Eddington ratio) for a BH of final mass M_{BH} : if we have some reasonable approximation to the *shape* of the lifetime function (Equation 5) – then the total M_{BH} must be given by the integral:

$$M_{\text{BH}} = \int L dt = \int L \frac{dt}{d \log L} d \log L. \quad (10)$$

In detail, we should truncate the lifetime function in Equation 5 at some minimum L where the integrated lifetime $t \rightarrow t_H$ – we do so but note in practice for the typical $\beta \approx 0.6$ this lower L is sufficiently small that the integral has already converged. If we assume $\beta(M_{\text{BH}})$ and η are relatively constant (and whatever physics sets them appears, at least in simulations, to be local to the AGN and redshift independent), then solving this equation yields a normalization t_0 appropriate for determining the *integrated* quasar lifetime above each L (for a more detailed derivation, see Hopkins et al. 2006a).

It is straightforward (although the continuity equations become somewhat more cumbersome and have to be solved numerically) to rewrite this derivation in terms of the Eddington ratio distribution (Equation 6), and the solution is similar⁷. Performing this exercise, we can approximate the integrated lifetimes at each logarithmic range in λ by the same Equation 6 (with $\eta \approx 0.4$), but with an “integrated” normalization t_0 whose constrained value can be roughly approximated as $1.26(1 - 1.8[\beta - 0.6]) \times 10^8 \text{ yr}$ (near a constant $\sim 10^8 \text{ yr}$, i.e. $\approx 2 - 3$ Salpeter times, appropriate for growing a BH by an order of magnitude in mass). Integrating Equation 6 from some minimum λ to $\lambda = 1$ with this normalization, we obtain the required integrated lifetime in each λ range.

(3) “Episodic” Lifetimes: This is the lifetime of an individual AGN/quasar “event” or some effective width in each peak in the quasar lightcurve. If AGN are excited to high Eddington ratio by some sort of trigger, then the time they spend in a given range of L or λ as a consequence of just that trigger is the appropriate episodic lifetime. The duty cycle at a given redshift is (to lowest order) the product of the triggering rate and episodic lifetime: for a fixed episodic lifetime, if the triggering rate is doubled, the effective lifetime doubles as well. Recall, it is the duty cycle (volume-averaged fraction) of BHs at a given Eddington ratio/luminosity that is observationally constrained by what we consider here. Therefore, based on the observations discussed thus far, we can only place an upper limit on the episodic lifetime (it cannot, of course, be longer than the integrated lifetime). We discuss the degeneracies this implies and possibilities for direct constraints and complementary lower limits to the episodic lifetime in § 6.

Figure 9 compares the effective and integrated lifetimes determined from our fitting as a function of BH mass and minimum λ . The effective lifetime, being essentially a duty cycle,

⁷ In doing so, one is implicitly referring to the time spent at a given Eddington ratio *near the final BH mass* as observed at the appropriate redshift. There is no way to determine from the observations here, for example, if a $10^9 M_{\odot}$ BH spent many Salpeter times growing in near-Eddington limited fashion from a small $\sim 1 M_{\odot}$ seed at very high redshifts, or formed directly as a $\sim 10^5 M_{\odot}$ seed BH. More rigorously we could write this as $t > L/L_{\text{Edd}}(M_{\text{BH}, t})$. However in terms of the Eddington ratio distribution that would be measured at any epoch in a relatively narrow range of BH masses, the two are identical.

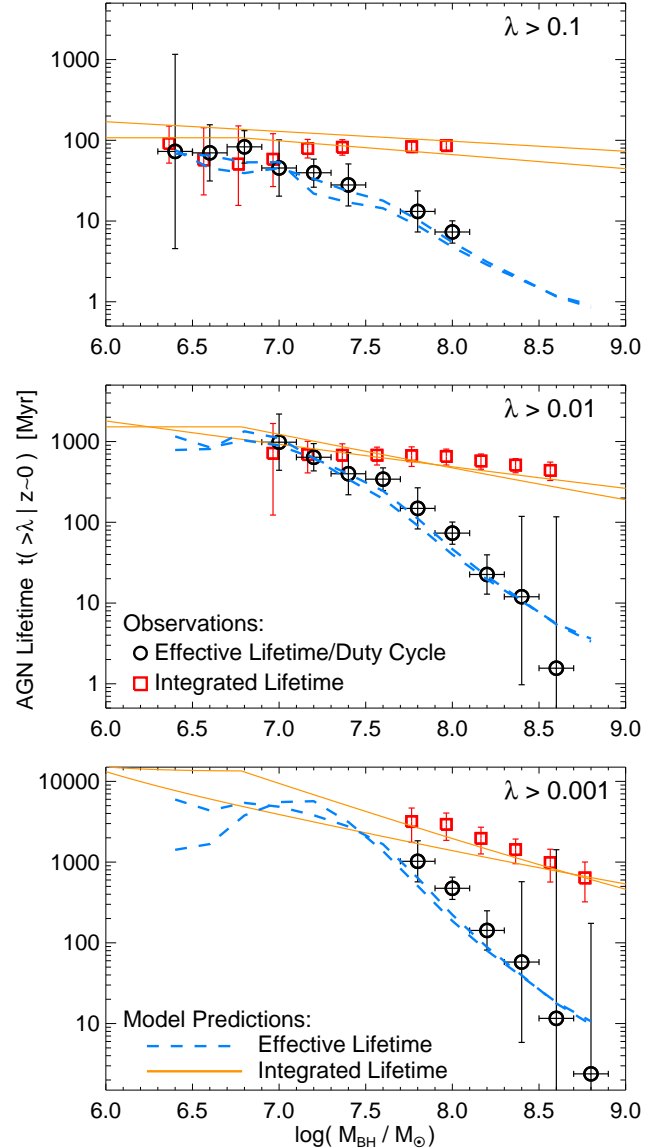


FIG. 9.— Effective quasar lifetime (black circles) above a given Eddington ratio λ , i.e. the duty cycle $\delta(> \lambda)$ times the Hubble time, fitted as a function of BH mass to the data in Figures 1-2. The data are at low redshift, so this should be thought of as the average lifetime for objects whose luminosity function is similar to that observed at low- z (see text). Dashed blue lines compare the model predictions, integrating over the cosmological history of triggering. The integrated quasar lifetime (red squares) determined by integrating over the best-fit lifetime distribution to the observations in each bin of M_{BH} (applying the appropriate mass conservation equation) is shown for comparison. Dotted orange lines are the predicted integrated quasar lifetimes – reflecting the total time at each λ (integrated over any arbitrary cosmological history) required by continuity to produce a BH at $z = 0$ with the observed mass. The two lifetimes are similar for BHs at lower masses $M_{\text{BH}} \lesssim 10^{7.5} M_{\odot}$ (at which masses observations suggest the duty cycle/volume density of quasars has been relatively constant since $z \sim 1 - 2$), but different at the highest masses (where there has been a steep drop-off in quasar activity since $z \sim 2$; i.e. the systems did most of their growth preferentially at high redshifts and have lower average integrated lifetimes at higher λ today).

can be determined directly from the Eddington ratio distributions where they cover the necessary dynamic range. The results are as expected – roughly constant at low BH masses, with a steep fall in duty cycles/effective lifetimes at high- M_{BH} . The integrated lifetimes are determined from the fitted lifetime distributions (Equation 6), given the continuity

requirement. They are also as expected – for $\lambda \gtrsim 0.1$, for example, this gives an expectation that most BHs spent \sim a few Salpeter times at high Eddington ratio, a requirement in any self-consistent Eddington-limited model. They are much more weakly mass-dependent – since λ is dimensionless and BH growth at fixed λ exponential, we would expect that (to lowest order) systems of different masses should have spent similar time at λ over the course of their evolution. At high λ , this is true by definition; as noted above the time at a given λ refers to time near that λ near the final BH mass – if the Eddington limit is applicable, all systems must spend a similar time in this regime.

Comparing the integrated and effective lifetimes yields an effective ratio of duty cycles:

$$\frac{t_{\text{eff}}}{t_{\text{int}}} = \frac{\delta(> \lambda | z_{\text{obs}})}{\langle \delta(> \lambda | z > z_{\text{obs}}) \rangle}, \quad (11)$$

where $\langle \delta(> \lambda | z > z_{\text{obs}}) \rangle$ is an appropriate weighted average δ (weighting over the time integral in Equation 9 and by the relative fraction of systems with mass M_{BH} at z_{obs} formed around each z). Roughly speaking, this yields the ratio of the duty cycle at the observed redshift z_{obs} to that during the epoch in which most BHs of the given mass were actively growing. For local low-mass BHs, this ratio is near unity; for high-mass BHs, it falls steeply to somewhere between $\sim 1 - 10\%$ of its high-redshift value, consistent with the now-standard picture of “downsizing” in BH growth.

Note that the lifetimes in Figure 9 are not directly analogous to those in previous works (e.g. Yu & Tremaine 2002; Salucci et al. 1999; Shankar et al. 2004; Marconi et al. 2004), because those assume a simplified light-bulb model, rather than a continuous distribution of accretion rates.

Also, these models usually quote the total time at high Eddington ratio at all BH masses since some adopted initial condition (not the time at some Eddington ratio in a narrow range of BH mass, constrained directly by observations here); as such, the comparison requires matching the boundary conditions for seed BHs in the light-bulb models.⁸ Also, δ here is the duty cycle above some λ , and is not directly comparable to the “AGN fraction” often defined observationally.

For example, Heckman et al. (2004) point out (as seen here in the same data) that $\delta(> \lambda)$ decreases strongly with M_{BH} ; however, Kauffmann et al. (2003) and Kewley et al. (2006) argue from the same data that the “AGN fraction” increases with galaxy mass. This apparent contradiction owes to a number of well-known trends: bulge-to-disk ratios and velocity dispersions drop rapidly in low-mass galaxies, so the span of BH mass and galaxy mass are not the same at low masses; specific star formation rates also increase rapidly in low-mass galaxies, so for the same fractional Eddington luminosity low-mass systems will be less AGN-dominated; and most such samples are AGN luminosity-limited, and so probe more massive hosts to lower Eddington ratios. Modeling these selection effects is non-trivial and quite sensitive to the specific selection effects and wavelengths of each observed sample – we therefore do not attempt a direct comparison here. We do note, however, that the comparison of the Heckman et al.

⁸ Specifically, we show in Figure 9 the average time spent in a given λ range for individual BHs with a given $z = 0$ mass. Some other definitions adopted in the literature reflect the *total* time for all BHs that are in or passed through a given BH mass interval (e.g. massive $z = 0$ BHs when they were, earlier, at this mass). The latter lifetime definition will be higher at lower masses, since it includes the time local, more massive systems spent “on” at earlier times getting to their present-day masses.

(2004); Kauffmann & Heckman (2008) and Kauffmann et al. (2003); Kewley et al. (2006) samples indicates the consistency of these different indicators; in Hopkins et al. (2008d) the authors attempt to model similar selection effects to compare the self-regulated lightcurves considered here with AGN fractions in the samples observed by Kauffmann et al. (2003) at low redshift and Erb et al. (2006); Kriek et al. (2007) at high redshift (see also Silverman et al. 2008a, and references therein), and find consistent results.

6. TRANSLATION TO AGN LIGHTCURVES: EPISODIC VERSUS INTEGRATED LIFETIMES

Given the quasar lifetime distribution, we can ask how this relates to some “average” quasar lightcurve $L(t | M_{\text{BH}})$. If the average lightcurve were monotonic, this would be trivial: inverting $dt/d\log L$ and integrating defines $L(t)$. For the Schechter function in Equation 6 this is numerically straightforward but tedious; we find a statistically identical answer with the convenient analytic representation of the lightcurve

$$\lambda = \left[1 + (t/t_Q)^{1/2} \right]^{-2/\beta} \quad (12)$$

where β is the same fitted to the lifetime distribution in Equation 6 and

$$t_Q = t_Q^{\text{max}} \equiv t_0 \frac{\eta^\beta}{\beta \ln 10} \quad (13)$$

relative to the t_0 fitted in Equation 6 (the zero-point in t in Equation 12 is obviously arbitrary and here is fixed to the time when the lightcurve is at maximum). For the median $\beta \approx 0.6$, this gives $t_Q \approx 5.3 \times 10^7$ yr, similar to the Salpeter time and characteristic dynamical times in the central regions of galaxies. It is straightforward to see that the lifetime distribution yielded by this lightcurve is statistically a good match to the Schechter function fits and observations in § 2, and it is convenient for analytic models of quasar evolution. If $t_Q \rightarrow t_Q/2$ and $t \rightarrow |t|$, it is trivial to treat this as time-symmetric about some peak, and a nearly identical lifetime distribution is obtained if one assumes exponential growth up to a peak luminosity followed by a “decay phase” given by Equation 12.

Equation 12 is a good match to the decline or blowout phase of quasar evolution in feedback-regulated simulations (Hopkins et al. 2006b). The assumption of monotonicity need not be exact, provided that the (average) fractional amplitude of variation in the quasar luminosity is small (factor $\lesssim 2 - 3$) on timescales shorter than the integrated quasar lifetime at the luminosity of interest. At some point as the quasar lifetime approaches the Hubble time this is almost certainly not true; however, for the short, high- L periods this is reasonable.

What if, however, there were two such episodes (each compressed in duration and separated by a cosmological interval $\gtrsim \text{Gyr}$)? Figure 10 shows a few such models (arbitrarily re-normalized to produce the same final BH mass). First, a monotonic or single event case (for convenience, we assume the rise and fall about each “event” are symmetric). Second, a superposition of three such events, each with shorter duration so that the sum $dt/d\log L$ is conserved. Third, a superposition of a couple of bright events and a larger number of lower-peak luminosity events, with a power spectrum following $dt/d\log L$. And fourth, a superposition of such events with effectively short episodic lifetimes ($t_Q \ll t_{\text{eff}}$), or equivalently large-amplitude variability on short timescales added according to a power-law spectrum with slope β .

All of these model lightcurves reproduce the same lifetime distribution $dt/d\log L$, and given $dt/d\log L$ and the observed




FIG. 10.— Examples of different quasar lightcurves that match the observed quasar lifetime distribution (Equation 6). The models are distinguished by different *episodic* quasar lifetimes and frequency of triggering. The integrated quasar lifetime/duty cycle and average post-peak decay are the same (required to match the observed λ distributions) – the events decay after each peak following Equation 12. Different observations, for example probing the transverse proximity effect, are required to break the degeneracies between these models. Constraints from these observations at present appear to rule out the model in the lower right panel, but the others, with an episodic lifetime $\sim 0.3 - 1$ times the effective/integrated lifetime, are allowed.

Eddington ratio distribution, the observed QLF will necessarily be the same. Moreover, they can all be normalized to give the same final BH mass (although for this final BH mass, they may correspond to different initial BH masses; but this is much smaller than the final mass and so not constrained by the observations considered). Measurements of the Eddington ratio distribution cannot uniquely determine the average lightcurve; only $dt/d\log L$.

In physical and observable terms, this is because the Eddington ratio distribution and corresponding $dt/d\log L$ yield only the *effective* and *integrated* quasar lifetimes – i.e. the duty cycle and total lifetime – at some L , not the *episodic* lifetime. Averaged over some redshift interval, the time $t(L)$ determined from $dt/d\log L(L)$ is the *total* time that an object will be active at that L in the interval. Whether that time came continuously – in a single episode – or in some large number of shorter episodes, cannot be inferred from the observations considered thus far.

The episodic lifetime is of particular interest because it relates to the triggering rate of quasars. If quasars are triggered at a rate \dot{n} and have an episodic lifetime (i.e. a duration around luminosity L per peak or triggering episode) of t_i ($\sim t_Q$ if the event is similar to the description of Equation 12), then the number density observed is $n = \dot{n}t_i$. If we know the observed number density and the effective lifetime t_{eff} , then attempting to infer the triggering rate we can only obtain

$$\dot{n} \approx \frac{n}{t_{\text{eff}}} \left(\frac{t_{\text{eff}}}{t_i} \right), \quad (14)$$

i.e. we are limited by some guess for the ratio t_i/t_{eff} . The statement that $t_i = t_{\text{eff}}$ is equivalent to the assumption that every object at a given L has had only one triggering episode in recent time. But if t_i were smaller than t_{eff} by some ratio N ($t_i = t_{\text{eff}}/N$), then the implied triggering rate for the same number of observed systems is higher by a factor N .

Assuming a physical model for what triggers quasars, we can use observations of quantities such as e.g. the galaxy merger rate to determine \dot{n} , and doing so for the case of galaxy

mergers gives the result that $t_i \sim t_{\text{eff}}$ (i.e. there are typically ~ 1 triggers per unit Hubble time). If we wish to determine \dot{n} without reference to a specific quasar fueling model, then we need the episodic lifetime t_i .

The observed rate of evolution of the quasar luminosity function places an upper limit $t_i \lesssim 10^9$ yr at high luminosities (Martini 2004), but this is already larger than the integrated lifetime at these luminosities. Lower limits can be derived from the sizes of narrow line regions (Bennett et al. 2002), but these are short ($t_i > 3 \times 10^4$ yr), so are not especially constraining ($t_i \ll t_{\text{eff}}$ is unlikely, as it requires extremely high rates of triggering, on timescales faster than the relevant dynamical times). Indirect measures relying on matching the observed BH mass function and QLF (Yu & Tremaine 2002; Salucci et al. 1999; Shankar et al. 2004; Marconi et al. 2004) and halo occupation models which compare clustering data (Porciani et al. 2004; Croom et al. 2005; Fine et al. 2006; Porciani & Norberg 2006; da Angela et al. 2008) constrain only the effective lifetimes.

One probe of episodic lifetimes is the transverse proximity effect – if individual episodes are short, the sizes of ionized bubbles around quasars should be smaller than if individual episodes are long. Preliminary constraints from non-detections (see e.g. Schirber et al. 2004; Martini 2004, for a review) suggest lifetimes $t_i > 10^7$ yr (with a strong limit $t_i \gg 10^6$ yr) at high luminosities ($\lambda \gtrsim 0.1$) where the observations are possible. More recently, potential detections (Jakobsen et al. 2003; Gonçalves et al. 2008) and indirect proximity effects (Worseck & Wisotzki 2006; Worseck et al. 2007) in a few objects suggest $t_i \approx 2.5 - 5 \times 10^7$ yr (almost exactly $t_i = t_{\text{eff}} \sim t_Q^{\text{max}}$ for $\lambda \gtrsim 0.1 - 0.2$). Similar analysis of the Gunn-Peterson effect around high-redshift ($z > 6$) quasars suggests comparable episodic lifetimes (Bajtlik et al. 1988; Haiman & Cen 2002; Yu & Lu 2005), but is more subject to uncertainties in the structure of the surrounding gas (Lidz et al. 2007).

The lengths of relativistic jets and radio lobes also imply episodic lifetimes, but it is not clear that the lifetime of radio loud activity or large-scale jet formation is the same as (or even correlated with) the lifetime for specific bolometrically luminous activity. These observations (see e.g. Scheuer 1995; Blundell et al. 1999) suggest lower-limits for $t_i \gtrsim$ a few 10^8 yr, but it is important to note that the observed systems with large jets in these samples are primarily at lower Eddington ratio ~ 0.01 , so this is comparable to or a factor of a few smaller than the total lifetime at these luminosities. For more luminous FR II sources, a lifetime similar to that from the transverse proximity effect, $\approx 2 \times 10^7$ yr, has been estimated (Bird et al. 2008) (although see also Reynolds & Begelman 1997; Merloni & Heinz 2007).

These observations, although tentative, seem to suggest an episodic lifetime similar to the total lifetime, in the range $t_i/t_{\text{eff}} \approx 0.3 - 1$ over reasonably high luminosity ranges $\lambda \sim 0.01 - 1$. This is consistent with simulations and cosmological models, and similar to the characteristic timescales of the problem (the Salpeter time for the growth of the BH, 4.2×10^7 yr, and characteristic dynamical/free-fall times in galaxy centers, $\sim 10^7 - 10^8$ yr). Improved constraints would permit the extension of this comparison to lower Eddington ratios and enable observations to uniquely determine the *triggering rates* of AGN as a function of BH mass, luminosity, and redshift.

Most of the observations considered here are at low redshift. Some redshift evolution in Eddington ratio distributions *must* occur, as indicated in Figure 9: $t_{\text{eff}} \ll t_{\text{int}}$ for massive BHs at low redshift, so there must have been some point at higher redshift where the duty-cycle of high-mass BHs was greater (when they accreted most of their mass, around $z \sim 2$). Constraints on this evolution can be obtained from the evolution of the quasar luminosity function – however, as discussed above, the QLF is primarily sensitive (especially at high redshifts, where the faint end is less well-constrained) to high luminosities $\lambda \gtrsim 0.1$; likewise integral constraints from the shape of the BH mass function (and the requirement that the integrated luminosity yield the appropriate final BH mass) primarily relate to these Eddington ratios, where most mass is gained.

As a consequence, the QLF constrains only the appropriate combination of β , t_0 , and η in Equation 6 such that the duty cycle at high Eddington ratio is reproduced. As noted earlier, the duty cycle is equivalent to the effective lifetime t_{eff} ($t_{\text{eff}}(\lambda > 0.1)/t_H(z)$ being the high- λ duty cycle). If the lifetime/duty cycle is parameterized in the form of Equation 6, then it is straightforward to show that this duty cycle of interest is

$$t_{\text{eff}}(\lambda > 0.1) \approx t_0 \eta \exp(-0.1/\eta), \quad (15)$$

essentially independent of the slope parameter β (that affecting the lifetime only at lower λ) and only weakly dependent on η ; for the reasonable physical range $\eta \sim 0.4 - 1$, t_{eff} changes by a factor $\sim 2 - 3$ (η cannot decrease much below this with redshift, or the number density of bright objects, observed to rise, would be exponentially suppressed). Various observations and synthesis models have been used to constrain this duty cycle – since this is just the high- λ duty cycle, “light bulb” and more sophisticated models yield nearly identical results here (see e.g. Haehnelt et al. 1998; Yu & Tremaine 2002; Yu & Lu 2004; Haiman et al. 2004; Marconi et al. 2004; Shankar et al. 2004). From the fits of these synthesis models to the QLF, or from constructing an analogous simple model (assuming the lightcurves here, but fitting to the observed QLF from Hopkins et al. (2007e) and integrating the continuity equations as in the synthesis models above), it can be seen that the duty cycle at fixed BH mass increases with redshift (at least from $z \sim 0 - 2$) in approximate power-law fashion, i.e. as

$$t_{\text{eff}}(\lambda > 0.1) \propto (1+z)^\alpha \quad (16)$$

with a maximum $t_{\text{eff}} = t_{\text{int}} \approx 10^8$ yr for this λ (obviously t_{eff} cannot increase beyond this point – and indeed this reproduces the “flattening” in duty cycles at $z \gtrsim 2$).

Clearly, the duty cycle of high-mass systems must increase more quickly with redshift than that of low-mass systems. From the comparison with observations in the various synthesis models above, we note that the needed evolution can be crudely approximated as

$$\alpha = \ln \left[1 + \frac{M_{\text{BH}}}{10^7 M_\odot} \right]. \quad (17)$$

This is not a rigorous derivation; it simply provides a useful interpolation formula that approximately reproduces the QLF evolution. At lowest order, this power-law increase in the duty cycle of high- λ activity simply reflects the evolution in the observed number density of quasars at a given $L = L_{\text{Edd}}(M_{\text{BH}})$ (see e.g. Hasinger et al. 2005). Second order corrections come from e.g. the evolving number density of

BHs, but these depend only weakly on the lightcurve model in this λ range.

Given this, the simplest possible model is indeed consistent with the observations: a model in which quasar lightcurves (β and t_0 , in Equation 12) are redshift-independent. Effectively this means β and η in Equation 6 are fixed, as is the episodic lifetime of quasars – all physics local to AGN evolution are redshift-independent, only the triggering rate (hence the duty cycle and corresponding t_{eff} and t_0 in Equation 6) evolve. This is demonstrated in models that adopt this assumption (Hopkins et al. 2006a, 2007e, 2008d,b), but the form of the constraints in Equation 16 makes it implicit.

It is clearly important to obtain direct constraints on the Eddington ratio distributions at high redshift, particularly constraints on the low- λ population, in order to break the degeneracies between this redshift-independent model and one in which quasar lightcurves evolve. If there is some evolution, it may indicate a difference in fueling or feedback modes: one could imagine different lightcurves resulting if the bright, high-redshift population is fueled by violent mergers where the low-redshift population is fueled by stochastic mechanisms, bar-induced inflows, or minor mergers; or if feedback is not important in some populations – at high redshift, feedback may act efficiently as the systems of interest are more bulge-dominated (leading to more sharply peaked lightcurves), whereas at low redshift, with activity in primarily low- M_{BH} systems with large, gas-rich disks, feedback may be ineffective at expelling gas content and suppressing inflows on large scales.

There are some constraints on β and t_0 that can be obtained from low-redshift observations. As discussed above, if a population is still growing (i.e. there is no sharp feature in the redshift history of triggering/growth) then the λ distribution today simply reflects typical low-redshift lightcurves. However, if there are distinct differences in the redshift history of triggering – if one population (at the same BH mass) is “quenched” (ceases growth/new triggering) at a different time than another – then the resulting Eddington ratio distribution (for the same quasar lightcurves) will not be the same at $z = 0$.

Figure 11 demonstrates this with a very simple toy model. Assume that AGN lightcurves are universal and redshift-independent, given by Equation 12 with $t_0 = 5 \times 10^7$ yr and $\beta = 0.6$. Objects have luminosity $L = 0$ until “triggered” at some time, then follow this simple lightcurve decay (t in Equation 12 is the time since the trigger). The probability of a “trigger” as a function of time we arbitrarily parameterize as a Gaussian in cosmic time, rising from high redshift ($z \sim 6$) to some peak at the “quenching redshift” z_{Quench} and then declining to $z = 0$. Specifically,

$$P(\text{trigger}|z) \propto \exp \left\{ \frac{t_H(z)}{t_H(z=6)} - \left[\frac{t_H(z)}{t_H(z_{\text{Quench}})} \right]^2 \right\}. \quad (18)$$

This is roughly chosen to correspond to the shape of the evolution in the observed AGN luminosity density, but we emphasize that it is just for illustrative purposes. Figure 11 shows the resulting $z = 0$ distribution in λ , for a set of Monte Carlo populations evolved in this toy model, but with a different quenching redshift z_{Quench} for each.

Populations that quench early (high z_{Quench}) have all decayed for a long time, to pile up at low λ and form a power-law distribution in λ above some very low minimum λ . (The power-law *must* turn over at some sufficiently low λ , such that the duty cycle integrated over all Eddington ratios is unity; the specific turnover around $\lambda \sim 10^{-4}$ in Figure 11 reflects the




FIG. 11. — *Top*: Distribution of Eddington ratios (arbitrary units) predicted for populations with the same lightcurve (Equation 12, with $\beta = 0.6$) but different toy model triggering histories. Triggering rises with time in the same manner until a redshift z_{Quench} , when the systems are shut down (a roughly Gaussian rise/fall motivated by observations). For all still-active/young systems (low z_{Quench}), the distributions asymptote to a single distribution given by the ratio of $t(>L)/t_H$ for a single lightcurve/event (Equations 1-3; with a turnover once $t(>L) \sim t_H$). For quenched/inactive systems (high z_{Quench}), the distributions “pile up” in the long power-law tail of decayed luminosities since the epoch of triggering. *Bottom*: Observed distributions from Kauffmann & Heckman (2008) at fixed BH mass ($\sim 10^7 - 10^8 M_{\odot}$) as a function of stellar population age (the parameter D_n ; $D_n \lesssim 1.5 - 1.6$ corresponding to star-forming/still active systems, $D_n \sim 2$ corresponding to “red and dead” systems with stellar population ages $\sim t_H$). The observed distribution reflects the predicted trend: a single (redshift-independent) lightcurve is consistent with the observed dependence of λ on age/star-forming classification. The “log-normal” behavior in young populations simply reflects the universal nature of the lightcurve and inevitable turnover when $t_{\text{eff}}(>\lambda) \sim t_H$; allowed to decay after quenching, this becomes the power-law-like tail in old populations.

specific lightcurve and triggering functional forms assumed here, together with the age of the Universe. Adjusting the late-time lightcurve behavior or early-time triggering rates can shift this to lower λ .)

Populations that quench late (low z_{Quench}) – i.e. approach the limit of continuous growth still occurring today – show a power-law behavior at high- λ that directly traces the quasar lightcurve (as discussed above, with duty cycles simply given by $t_{\text{eff}} \sim t_{\text{int}}$), with a turnover where $t_{\text{eff}} \rightarrow t_H$. The resulting distribution looks roughly log-normal, and asymptotes to the same distribution for all low- z_{Quench} (all still-active populations).

Note that the details in Figure 11 are somewhat arbitrary – we have just chosen a toy model triggering history to highlight the dependence on models with strong features in that history; it is easy to construct others. It is possible to tune

these histories such that, for example, the high- z_{Quench} predictions continue their power-law like behavior to lower λ or shift the “turnover” in the lognormal (low- z_{Quench}) regime. But the *qualitative* results are insensitive to the precise parameterization of the triggering history, provided there is a significant feature/shutdown after a given time.

Recently, Kauffmann & Heckman (2008) expanded upon the Eddington ratio distributions measured by Heckman et al. (2004) and Yu et al. (2005), and quantified the λ distribution as a function of stellar population age (specifically the observable parameter D_n , which is a tracer of activity: $D_n \lesssim 1.5$ systems being still active or $\lesssim 3$ Gyr old, in mean stellar population age, whereas $D_n \rightarrow 2$ systems are quenched with ages $\sim t_H$). Figure 11 compares the λ distribution measured as a function of stellar population properties (for BHs with fixed mass $10^7 - 10^8 M_{\odot}$) at $z = 0$. The trends are very similar to those predicted (and it is not hard to imagine more detailed star formation history models, more precisely tuned to yield quantitative agreement).

This is already interesting: Kauffmann & Heckman (2008) interpret the observed trend as an indicator of two independent accretion modes (a “power-law” mode, in old systems, and a “log-normal” mode, in young ones). This could still be the case, but the comparison here demonstrates that the observed trend is also the natural expectation of a model in which all AGN lightcurves are identical, but there are simply differences in the triggering rate corresponding to when different galaxies/BH populations were “quenched” or slowed their growth. The uniformity of the Eddington ratio distribution in the “log-normal” regime is not surprising in this case – it simply reflects the fact that quasar lightcurves are similar and that all populations in this regime of stellar population age are “still active” (i.e. are continuously growing, in a population-averaged sense; they do not have a sharp feature of higher-redshift activity that dominated their growth); the “power-law” regime is simply the “log-normal” population allowed to decay to lower luminosities as triggering rates decline with redshift.

Moreover, these observed trends place significant constraints on lightcurve evolution. Figure 12 considers the same model predictions, but in models where the lightcurve parameters in Equation 12 evolve with the “triggering redshift.” For convenience we parameterize the evolution simply as $\beta \propto (1+z)^{\beta'_z}$ and $t_Q \propto (1+z)^{t'_z}$, where z here refers to the redshift where a given trigger occurs. For $\beta'_z = -0.5$ or $t'_z = -0.5$, the predictions are similar to the no-evolution case, consistent with the observations. Again, we stress that the exact mapping to observations depends on the precise redshift evolution of the triggering distribution, and more importantly on how star formation histories in detail evolve relative to AGN triggering histories, needed to predict observed quantities such as D_n – however, as long as the observed qualitative features are intact, it is always possible to find models where those assumptions are adjusted to give a more precise quantitative match to the observations. However, for $\beta'_z = 0.2$ or $t'_z = 0.5$, the results are qualitatively much different – they do not resemble the observations. The sense of the implied evolution in this case would be that lightcurves become more shallow/extended with redshift: as a consequence, quenched systems would not decay sufficiently relative to observations. In these cases, where the features of the predicted distributions are qualitatively different from the observations, we find that no tuning of the triggering rate distributions is able to




FIG. 12.— As Figure 11, but in models where the median lightcurve (parameterized by lifetime slope β and episodic lifetime t_Q ; Equation 12) evolves with redshift. *Top*: Lightcurves become more shallow/extended with increasing redshift ($\beta \propto (1+z)^{0.2}$, $t_Q \propto (1+z)^{0.5}$). It is not possible to match the observed trends from Figure 11 with such evolution. *Bottom*: Lightcurves become more sharply peaked with increasing redshift ($\beta \propto (1+z)^{-0.5}$, $t_Q \propto (1+z)^{-0.5}$). Such evolution is consistent with the observed trends (and $z \sim 1$ observations in Figure 3). The trend of λ distribution with age/triggering history puts significant constraints on lightcurve evolution: evolution towards more “quiescent” models (including stellar wind or isolated accretion disk modes) is disallowed; no evolution or weak evolution towards more violent/efficient feedback modes are consistent with observations.

“fix” the disagreement with observations.

This comparison constrains the allowed range:

$$\begin{aligned} -0.5 &\leq \beta'_z \leq 0.05 \\ -0.7 &\leq t'_z \leq 0.25. \end{aligned} \quad (19)$$

These are not error bars; rather, they represent the allowed range in which solutions exist that can reproduce the observed trends from Kauffmann & Heckman (2008). Within this range, the results are also consistent with the observationally inferred Eddington ratio distributions from Merloni & Heinz (2009) at $z = 1$. The constraints are non-trivial: redshift evolution must be relatively mild, and the form of evolution allowed is such that lightcurves become more sharply peaked at higher redshifts – the sense that might be expected if triggering events are more violent and/or feedback is more efficient. Evolution in the opposite sense (evolution towards the predictions of the stellar wind fueling or isolated accretion disk models discussed in § 4, or towards less efficient feedback) is ruled out. Together with the constraints above regarding evolution in the duty cycle, $\delta \equiv t_{\text{eff}}/t_H \sim t_Q \times dN/dt$, the constraints on evolution in t_Q imply corresponding constraints on the triggering rate of independent AGN events, dN/dt .

8. IMPLICATIONS FOR THE INTEGRATED GROWTH OF BLACK HOLES

Given the observationally constrained AGN lifetime distribution, Figure 13 uses it to infer how different Eddington ratios contribute to integrated black hole growth. Essentially this amounts to assuming that the shape of the Eddington ratio distribution is well-described by Equation 6, and integrating over the time at each Eddington ratio to determine the fractional contribution to the final BH mass from each range in




FIG. 13.— Contribution from different Eddington ratios to the integrated BH growth, given the observationally constrained typical quasar lifetime distribution (Equation 6). *Top*: Differential contribution to the integrated BH mass from each logarithmic interval in Eddington ratio λ , shown with a linear (*left*) and logarithmic (*right*) y-axis scale. *Bottom*: Cumulative contribution from all Eddington ratios less than the given λ . From the observations, the BH growth is dominated by moderate/large Eddington ratios ~ 0.2 . The contribution from low Eddington ratios that may be radiatively inefficient ($\lambda \lesssim 0.01$) is small ($\sim 20\%$), but non-negligible, in agreement with various independent constraints (see Hopkins et al. 2006d).

λ . Recall from § 4, the total mass accreted must add to the total BH mass, and so in fractional terms the absolute normalization of the lifetimes can be factored out. Also as noted above, it makes little difference (because little mass growth is contributed by such low accretion rates) whether we truncate the distribution at some sufficiently low λ or integrate down to $\lambda \rightarrow 0$. For simplicity, we consider the median best-fit lifetime distribution with $\beta \approx 0.6$, the dependence on mass that may be present is sufficiently weak that the results are similar in the observed range.

Figure 13 shows the fractional contribution to the final BH mass from each logarithmic interval in λ (as well as the cumulative contribution from all Eddington ratios $< \lambda$). Given the observations, BH growth is dominated by moderate/large Eddington ratios ~ 0.2 . The contribution from low Eddington ratios that may be radiatively inefficient ($\lambda \lesssim 0.01$) is small ($\sim 20\%$), but it is worth noting that it is not entirely negligible, and this fraction is sufficient that such populations could be a significant contributor to the growth in some low-luminosity AGN populations (for more detailed comparison, we refer to Hopkins et al. 2009d). These expectations are in good agreement with various independent constraints, integral arguments (Soltan 1982), and models for the fueling of AGN and buildup of the BH mass function (for a review of these constraints and discussion of the contribution of radiatively inefficient sources, see Hopkins et al. 2006d, and references therein).

For this reason, it is easy to construct models matching both the observed QLF and the $z = 0$ BH mass function with a “universal” lightcurve within the constraints developed here. For the specific lightcurve models discussed here, Hopkins et al. (2007e) show this explicitly (see their Figure 10, comparing the predicted BH mass function from fitting the compiled QLF data with the same set of light curves as those considered here). Similar conclusions are reached by Yu & Lu (2008).

9. DISCUSSION AND CONCLUSIONS

9.1. Overview

We have compared observations of the Eddington ratio distribution as a function of BH mass, redshift, and luminosity with various theories. We find good agreement between the observations and the predictions of the self-regulated feedback models for BH growth and evolution from Hopkins et al. (2005d,a,b, 2006a,b); Hopkins & Hernquist (2006). The agreement covers the entire observed dynamic range, with observations extending down 5 – 6 orders of magnitude in Eddington ratio ($L/L_{\text{Edd}} \sim 10^{-6} - 1$), over three orders of magnitude in BH mass ($M_{\text{BH}} \sim 10^6 - 10^9 M_{\odot}$), corresponding to ~ 8 orders of magnitude in luminosity, and (given indirect observational constraints) from redshifts $z \sim 0 - 1$. The models are not fitted to these observations – there are no free parameters adjustable, and the range of uncertainty between different versions of the models constrained to reproduce the observed quasar luminosity function is small (within the observational error bars).

The lifetime – already as constrained in a purely empirical sense from the observations, is clearly *not* a delta function – i.e. quasars are not “light bulbs” – rather, it is a smooth, continuous, and relatively steep function of luminosity/Eddington ratio, with more time spent at lower luminosities/Eddington ratios.

We also show agreement with observed Eddington ratio distributions as a function of AGN luminosity, but these are less constraining, because of the inherent Eddington ratio limits implied by the selection (wavelength-dependent effects further complicate comparisons; discussed in Hopkins et al. 2009d).

9.2. Implied AGN Lifetimes and Growth Histories

The observationally implied lifetime distribution can be generally parameterized (in model-independent fashion) as a Schechter function, with a characteristic normalization lifetime, a turnover at Eddington ratios near ~ 1 (reflecting a physical limit around the Eddington limit), and a faint-end slope β (such that the luminosity-dependent lifetime $t(> \lambda) \propto \lambda^{-\beta}$ at small λ). The observations favor a best-fit $\beta \approx 0.6 \pm 0.05$ (in agreement with feedback-regulated models, discussed below) for typical $\sim L_*$ galaxies and BHs. Combined with constraints from the quasar luminosity function, there is evidence for weak BH mass-dependence of β , as predicted by hydrodynamic simulations (Hopkins et al. 2006b).

The observations directly yield quasar duty cycles and lifetimes as a function of Eddington ratio and BH mass. In terms of the integrated time at a given Eddington ratio interval (around some final BH mass), the “quasar” lifetime – i.e. lifetime at high Eddington ratios $\gtrsim 0.1$, is similar to that expected from the Salpeter (1964) time and other observational constraints, $\sim 10^8$ yr. However, the time at lower Eddington ratios rises rapidly – with $t(\lambda \gtrsim 0.01) \sim 0.5 - 1$ Gyr and $t(\lambda \gtrsim 0.001) \sim 1 - 5$ Gyr.

This allows us to quantify the fractional contribution to present-day BH mass from various intervals in λ ; growth is dominated by large $\lambda \sim 0.2$, with a small – but non-negligible – contribution $\sim 20\%$ from accretion at low Eddington ratios $\lambda \lesssim 0.01$, in agreement with integral arguments (Soltan 1982) and various independent constraints (Hopkins et al. 2006d, and references therein).

Comparison of these integral constraints with the $z = 0$ observations (comparing “integrated” and “effective” AGN lifetimes/duty cycles) reflects the increasingly established AGN downsizing trend: the low-mass BH population is still growing today, the high-mass population has shut down since its

earlier epoch of peak activity.

9.3. Constraints on Models of Lightcurves and Lifetimes

There is enormous constraining power in the Eddington ratio distribution at low λ in BH mass-limited samples. Already, the observations at $z = 0$ are sufficient to limit the Eddington ratio/quasar lifetime distribution to a narrow range around the theoretical predictions from recent hydrodynamic simulations incorporating a self-consistent model for accretion and feedback. Those models predict that the self-regulated nature of BH accretion should lead to a relatively self-similar decay phase in AGN luminosity or Eddington ratio (regardless of e.g. triggering mechanisms, or the exact details of feedback physics), $L \propto t^{-(1.5-2.0)}$, which gives rise to the observed power-law-like faint behavior in the lifetime distribution with predicted slope $\beta \approx 0.6$.

The observations strongly rule out (at $> 5\sigma$ significance) simplified models for quasar lightcurves, including: “light-bulb” models in which quasars turn “on” for some time at a fixed or relatively narrow range of accretion rates, or “exponential” models in which AGN grow at fixed Eddington ratio and then “shut off” rapidly.

Models where BH accretion simply traces stellar mass loss are also ruled out ($\gtrsim 4\sigma$). Stellar mass loss may still be a fuel source; however, the observations argue that some process must further regulate AGN activity, shutting down accretion more efficiently than the simple slow starvation expected if BH growth directly traced stellar mass loss.

At $\sim 3 - 5\sigma$ significance, the observations also rule out isolated accretion disk models (Yu et al. 2005, and references therein) – i.e. accretion disks fueled rapidly but then cut off from a future gas supply, without feedback (such that they evolve by gas exhaustion). These are a considerable improvement on the light bulb model, but are still ruled out formally – with the sense again that some process must shut down growth more efficiently (possibly the addition of even mild outflows to the isolated disk solution) – especially for high-mass BH populations.

Local observations of the Eddington ratio distribution alone do not strongly constrain the mass or redshift dependence of these properties. Combining the observations here with constraints from the quasar luminosity function, observations do favor a weak mass-dependence in the shape of the lifetime distribution, with the sense that β decreases with increasing BH mass ($\beta \sim 0.6 - 0.2 \{\log[M_{\text{BH}}/10^7 M_{\odot}]\}$). This is equivalent to the statement that more massive systems “shut off” more abruptly than low-mass systems; a trend predicted by hydrodynamic models, as quasar feedback becomes relatively more dominant (and bulge-to-disk ratios increase while global disk gas fractions decrease) in more massive systems.

Recently, Kauffmann & Heckman (2008) have shown that the λ distribution depends on the stellar population age of the system. We show that the observed trends naturally arise from a *single* lightcurve of the form constrained here, as a consequence of different triggering histories. Systems which are still active/growing (i.e. have not “quenched” their gas supply, BH growth, or star formation) remain at higher accretion rates reflecting the median lightcurve directly, while systems which quenched at some earlier epoch have decayed down to the lower λ power-law-like tail of the lightcurve distribution. The existence of these trends, and observational constraints on the Eddington ratio distribution at $z = 1$ (Merloni & Heinz 2009), set some constraints on the evolution of typical lightcurves with redshift.

We find that β and the episodic quasar lifetime t_Q cannot evolve strongly (parameterized as $\beta \propto (1+z)^{\beta'_z}$ and $t_Q \propto (1+z)^{t'_z}$; the allowed range is $-0.5 \leq \beta'_z \leq 0.05$, $-0.7 \leq t'_z \leq 0.25$). Evolution towards a stellar wind or isolated accretion disk solution is strongly ruled out – rather, the sense of (mild) evolution allowed is towards more sharply peaked lightcurves at higher redshift, which may be expected if fueling is more violent and/or feedback is more efficient in this regime. The evolution of duty cycles (and correspondingly “effective” AGN lifetimes) at the high-luminosity end ($\lambda \gtrsim 0.1$) is directly constrained by evolution of the quasar luminosity function – this limits the combination of t_0 and η (but not the faint-end slope β), and (given some constraint on the evolution of episodic lifetimes t_Q) the triggering rate of quasars to evolve according to Equations 16-17, increasing with redshift more rapidly in higher-mass BHs.

These constraints on the AGN can be transformed into constraints on the average AGN lightcurve, suggesting a characteristic power-law like decay ($L \propto t^{-(1.5-2.0)}$) similar to that predicted in models, after rapid growth to some peak luminosity. The observations of the Eddington ratio distribution tightly limit the shape of such lightcurves, but only weakly constrain the characteristic timescale for a single such “event” (i.e. the width in time of a single “peak”) – the episodic AGN lifetime – by setting an upper limit.

Observational constraints from e.g. the proximity effect in bright quasars can independently set lower limits on the episodic lifetime, however, and suggest that (at least over the observed range for the proximity effect, namely quasar activity with $\lambda \sim 0.1 - 1$) the episodic lifetime is comparable to the integrated lifetime, in the range $t_Q \sim 0.3 - 1.0 t_{\text{int}}$. This implies that *the typical massive BH has experienced no more than a couple of bright, high Eddington ratio quasar episodes while near its current ($z = 0$) mass*. This is in agreement with predictions from cosmological models that associate the brightest quasar activity with fueling in violent major mergers (Hopkins et al. 2008d), for which there are only a couple of events expected since high redshift. Improved constraints on the episodic lifetime will allow the observed quasar luminosity function to be more robustly translated into the distribution of AGN triggering rates as a function of BH and host mass. Extending the observational constraints on the episodic lifetime to lower luminosities/Eddington ratios is also important, as one might expect that although a typical object only experiences a couple of triggers to near-peak activity, it could have many more triggers for lower-level activity.

9.4. Other Tests and Future Work

Our findings agree with other (less direct) independent constraints. Recently, for example, Yu & Lu (2008) showed that the joint evolution of AGN luminosity functions with redshift favors similar lightcurves. These results (from the QLF and/or BH mass function and integral/continuity arguments) provide independent support for the constraints here, but are primarily sensitive to luminous (high- λ) behavior and cannot distinguish between similar lightcurves (with slightly different β), such as the isolated accretion disk or feedback-regulated predictions. The lightcurve shape can also be probed by the dependence of AGN clustering on luminosity and shape of the QLF or “active” BH mass function in deep samples; the observations at present appear to favor feedback-regulated models over simplified light-bulb or exponential lightcurve models (see Adelberger & Steidel 2005; Myers et al. 2007; Greene & Ho 2007; da Angela et al. 2008), but again deeper observations are needed to distinguish between the proposed physically motivated models.

Extending the observations of the Eddington ratio distribution to higher redshift will greatly improve these constraints as well as limit possible redshift evolution in quasar light curves. More massive black holes will be closer to their peak growth at higher redshifts, allowing observations to probe the times of greatest interest. At $z \sim 1 - 3$, the $\lambda \gtrsim 0.1$ end of the distribution, in broad-line luminous quasars, is already constrained by the observed QLF and application of the virial BH mass estimators (Kollmeier et al. 2006; Fine et al. 2008). Combined with these constraints, smaller volume but deep redshift surveys can be used to construct samples which are complete to a given BH/bulge mass, and similar narrow-line searches in these hosts could limit the Eddington ratio distribution at much lower ratios. If coverage is sufficient, X-ray data can be used in the same manner. At $z \geq 2$, complete, large-volume spectroscopic or X-ray samples are not available at present. Here, indirect tests will remain important for the near future, but improved constraints on what (if any) evolution is seen at lower redshifts in the shape of the lifetime distributions will considerably inform future models and observational efforts.

We thank Josh Younger, T. J. Cox, and Guinevere Kauffmann for helpful discussions. This work was supported in part by NSF grants ACI 96-19019, AST 00-71019, AST 02-06299, and AST 03-07690, and NASA ATP grants NAG5-12140, NAG5-13292, and NAG5-13381. Support for PFH was provided by the Miller Institute for Basic Research in Science, University of California Berkeley.

REFERENCES

- Adelberger, K. L., & Steidel, C. C. 2005, *ApJ*, 630, 50
 Aller, M. C., & Richstone, D. O. 2007, *ApJ*, 665, 120
 Bahcall, J. N., Kirhakos, S., Saxe, D. H., & Schneider, D. P. 1997, *ApJ*, 479, 642
 Bajtlik, S., Duncan, R. C., & Ostriker, J. P. 1988, *ApJ*, 327, 570
 Barger, A. J., & Cowie, L. L. 2005, *ApJ*, 635, 115
 Bassani, L., et al. 2006, *ApJ*, 636, L65
 Beckmann, V., Soldi, S., Shrader, C. R., Gehrels, N., & Produit, N. 2006, *ApJ*, 652, 126
 Bell, E. F., & de Jong, R. S. 2000, *MNRAS*, 312, 497
 Bell, E. F., McIntosh, D. H., Katz, N., & Weinberg, M. D. 2003, *ApJS*, 149, 289
 Bennert, N., Falcke, H., Schulz, H., Wilson, A. S., & Wills, B. J. 2002, *ApJ*, 574, L105
 Bird, J., Martini, P., & Kaiser, C. 2008, *ApJ*, 676, 147
 Blundell, K. M., Rawlings, S., & Willott, C. J. 1999, *AJ*, 117, 677
 Bruzual, G., & Charlot, S. 2003, *MNRAS*, 344, 1000
 Bundy, K., et al. 2006, *ApJ*, 651, 120
 —. 2008, *ApJ*, 681, 931
 Burkert, A., & Silk, J. 2001, *ApJ*, 554, L151
 Canalizo, G., & Stockton, A. 2001, *ApJ*, 555, 719
 Cao, X., & Xu, Y.-D. 2007, *MNRAS*, 377, 425
 Ciotti, L., & Ostriker, J. P. 1997, *ApJ*, 487, L105+
 —. 2001, *ApJ*, 551, 131
 —. 2007, *ApJ*, 665, 1038
 Coil, A. L., Hennawi, J. F., Newman, J. A., Cooper, M. C., & Davis, M. 2007, *ApJ*, 654, 115
 Croom, S. M., Boyle, B. J., Shanks, T., Smith, R. J., Miller, L., Outram, P. J., Loaring, N. S., Hoyle, F., & da Angela, J. 2005, *MNRAS*, 356, 415
 da Angela, J., et al. 2008, *MNRAS*, 383, 565

- Di Matteo, T., Colberg, J., Springel, V., Hernquist, L., & Sijacki, D. 2008, *ApJ*, 676, 33
- Di Matteo, T., Croft, R. A. C., Springel, V., & Hernquist, L. 2003, *ApJ*, 593, 56
- 2004, *ApJ*, 610, 80
- Di Matteo, T., Springel, V., & Hernquist, L. 2005, *Nature*, 433, 604
- Dunlop, J. S., McLure, R. J., Kukula, M. J., Baum, S. A., O’Dea, C. P., & Hughes, D. H. 2003, *MNRAS*, 340, 1095
- Elvis, M., et al. 1994, *ApJS*, 95, 1
- Erb, D. K., Steidel, C. C., Shapley, A. E., Pettini, M., Reddy, N. A., & Adelberger, K. L. 2006, *ApJ*, 646, 107
- Falcke, H., & Biermann, P. L. 1996, *A&A*, 308, 321
- Falcke, H., K rding, E., & Markoff, S. 2004, *A&A*, 414, 895
- Faucher-Gigu re, C.-A., Lidz, A., Hernquist, L., & Zaldarriaga, M. 2008a, *ApJ*, 682, L9
- 2008b, *ApJ*, 688, 85
- Faucher-Gigu re, C.-A., Lidz, A., Zaldarriaga, M., & Hernquist, L. 2008c, *ApJ*, 673, 39
- Fender, R., Koerding, E., Belloni, T., Uttley, P., McHardy, I., & Tzioumis, T. 2007, *ArXiv e-prints*, in press, arXiv:0706.3838, 706
- Ferrarese, L., & Merritt, D. 2000, *ApJ*, 539, L9
- Fine, S., et al. 2006, *MNRAS*, 373, 613
- 2008, *MNRAS*, 390, 1413
- Fiore, F., et al. 2003, *A&A*, 409, 79
- Fontanot, F., Cristiani, S., Monaco, P., Nonino, M., Vanzella, E., Brandt, W. N., Grazian, A., & Mao, J. 2007, *A&A*, 461, 39
- Gebhardt, K., et al. 2000, *ApJ*, 539, L13
- Gilli, R., Comastri, A., & Hasinger, G. 2007, *A&A*, 463, 79
- Goncalves, T. S., Steidel, C. C., & Pettini, M. 2008, *ApJ*, 676, 816
- Granato, G. L., De Zotti, G., Silva, L., Bressan, A., & Danese, L. 2004, *ApJ*, 600, 580
- Grazian, A., Negrello, M., Moscardini, L., Cristiani, S., Haehnelt, M. G., Matarrese, S., Omizzolo, A., & Vanzella, E. 2004, *AJ*, 127, 592
- Greene, J. E., & Ho, L. C. 2007, *ApJ*, 667, 131
- Greene, J. E., Ho, L. C., & Ulvestad, J. S. 2006, *ApJ*, 636, 56
- Haehnelt, M. G., Natarajan, P., & Rees, M. J. 1998, *MNRAS*, 300, 817
- Haiman, Z., & Cen, R. 2002, *ApJ*, 578, 702
- Haiman, Z., Ciotti, L., & Ostriker, J. P. 2004, *ApJ*, 606, 763
- Haiman, Z., Jimenez, R., & Bernardi, M. 2007, *ApJ*, 658, 721
- Hamilton, T. S., Casertano, S., & Turnshek, D. A. 2002, *ApJ*, 576, 61
- Hasinger, G. 2008, *A&A*, 490, 905
- Hasinger, G., Miyaji, T., & Schmidt, M. 2005, *A&A*, 441, 417
- Heckman, T. M., Kauffmann, G., Brinchmann, J., Charlot, S., Tremonti, C., & White, S. D. M. 2004, *ApJ*, 613, 109
- Hickox, R. C., et al. 2007, *ApJ*, 671, 1365
- 2009, *ApJ*, 696, 891
- Ho, L. C. 2002, *ApJ*, 564, 120
- Hopkins, P. F., Bundy, K., Hernquist, L., & Ellis, R. S. 2007a, *ApJ*, 659, 976
- Hopkins, P. F., Cox, T. J., Dutta, S. N., Hernquist, L., Kormendy, J., & Lauer, T. R. 2009a, *ApJS*, 181, 135
- Hopkins, P. F., Cox, T. J., & Hernquist, L. 2008a, *ApJ*, 689, 17
- Hopkins, P. F., Cox, T. J., Kere s, D., & Hernquist, L. 2008b, *ApJS*, 175, 390
- Hopkins, P. F., Cox, T. J., Younger, J. D., & Hernquist, L. 2009b, *ApJ*, 691, 1168
- Hopkins, P. F., & Hernquist, L. 2006, *ApJS*, 166, 1
- 2009, *ApJ*, 694, 599
- Hopkins, P. F., Hernquist, L., Cox, T. J., Di Matteo, T., Martini, P., Robertson, B., & Springel, V. 2005a, *ApJ*, 630, 705
- Hopkins, P. F., Hernquist, L., Cox, T. J., Di Matteo, T., Robertson, B., & Springel, V. 2005b, *ApJ*, 630, 716
- 2005c, *ApJ*, 632, 81
- 2006a, *ApJS*, 163, 1
- Hopkins, P. F., Hernquist, L., Cox, T. J., Dutta, S. N., & Rothberg, B. 2008c, *ApJ*, 679, 156
- Hopkins, P. F., Hernquist, L., Cox, T. J., & Kere s, D. 2008d, *ApJS*, 175, 356
- Hopkins, P. F., Hernquist, L., Cox, T. J., Kere s, D., & Wuyts, S. 2009c, *ApJ*, 691, 1424
- Hopkins, P. F., Hernquist, L., Cox, T. J., Robertson, B., Di Matteo, T., & Springel, V. 2006b, *ApJ*, 639, 700
- Hopkins, P. F., Hernquist, L., Cox, T. J., Robertson, B., & Krause, E. 2007b, *ApJ*, 669, 45
- 2007c, *ApJ*, 669, 67
- Hopkins, P. F., Hernquist, L., Cox, T. J., Robertson, B., & Springel, V. 2006c, *ApJS*, 163, 50
- Hopkins, P. F., Hernquist, L., Martini, P., Cox, T. J., Robertson, B., Di Matteo, T., & Springel, V. 2005d, *ApJ*, 625, L71
- Hopkins, P. F., Hickox, R., Quataert, E., & Hernquist, L. 2009d, *MNRAS*, in press, arXiv:0901.2936 [astro-ph]
- Hopkins, P. F., Lauer, T. R., Cox, T. J., Hernquist, L., & Kormendy, J. 2009e, *ApJS*, 181, 486
- Hopkins, P. F., Lidz, A., Hernquist, L., Coil, A. L., Myers, A. D., Cox, T. J., & Spergel, D. N. 2007d, *ApJ*, 662, 110
- Hopkins, P. F., Narayan, R., & Hernquist, L. 2006d, *ApJ*, 643, 641
- Hopkins, P. F., Richards, G. T., & Hernquist, L. 2007e, *ApJ*, 654, 731
- Hopkins, P. F., Somerville, R. S., Hernquist, L., Cox, T. J., Robertson, B., & Li, Y. 2006e, *ApJ*, 652, 864
- Jakobsen, P., Jansen, R. A., Wagner, S., & Reimers, D. 2003, *A&A*, 397, 891
- Jester, S. 2005, *ApJ*, 625, 667
- Johansson, P. H., Naab, T., & Burkert, A. 2009, *ApJ*, 690, 802
- Kauffmann, G., & Haehnelt, M. 2000, *MNRAS*, 311, 576
- Kauffmann, G., & Heckman, T. M. 2008, *MNRAS*, in press, arXiv:0812.1224
- Kauffmann, G., et al. 2003, *MNRAS*, 346, 1055
- Kewley, L. J., Groves, B., Kauffmann, G., & Heckman, T. 2006, *MNRAS*, 372, 961
- Kollmeier, J. A., et al. 2006, *ApJ*, 648, 128
- Komatsu, E., et al. 2009, *ApJS*, 180, 330
- Kriek, M., et al. 2007, *ApJ*, 669, 776
- La Franca, F., et al. 2005, *ApJ*, 635, 864
- Lapi, A., Shankar, F., Mao, J., Granato, G. L., Silva, L., De Zotti, G., & Danese, L. 2006, *ApJ*, 650, 42
- Leitherer, C., et al. 1999, *ApJS*, 123, 3
- Lidz, A., Hopkins, P. F., Cox, T. J., Hernquist, L., & Robertson, B. 2006, *ApJ*, 641, 41
- Lidz, A., McQuinn, M., Zaldarriaga, M., Hernquist, L., & Dutta, S. 2007, *ApJ*, 670, 39
- Maccarone, T. J., Gallo, E., & Fender, R. 2003, *MNRAS*, 345, L19
- Magorrian, J., et al. 1998, *AJ*, 115, 2285
- Marchesini, D., Celotti, A., & Ferrarese, L. 2004, *MNRAS*, 351, 733
- Marconi, A., & Hunt, L. K. 2003, *ApJ*, 589, L21
- Marconi, A., Risaliti, G., Gilli, R., Hunt, L. K., Maiolino, R., & Salvati, M. 2004, *MNRAS*, 351, 169
- Martini, P. 2004, in *Coevolution of Black Holes and Galaxies*, ed. L. C. C. U. P. Ho, 169
- McClintock, J. E., & Remillard, R. A. 2006, *Black hole binaries (Compact stellar X-ray sources)*, 157–213
- McQuinn, M., Lidz, A., Zaldarriaga, M., Hernquist, L., Hopkins, P. F., Dutta, S., & Faucher-Gigu re, C.-A. 2009, *ApJ*, 694, 842
- Meier, D. L. 2001, *ApJ*, 548, L9
- Menci, N., Cavaliere, A., Fontana, A., Giallongo, E., Poli, F., & Vittorini, V. 2003, *ApJ*, 587, L63
- Merloni, A., & Heinz, S. 2007, *MNRAS*, 381, 589
- 2008, *MNRAS*, 388, 1011
- 2009, *MNRAS*, in prep
- Merloni, A., Heinz, S., & di Matteo, T. 2003, *MNRAS*, 345, 1057
- Murray, N., Quataert, E., & Thompson, T. A. 2005, *ApJ*, 618, 569
- Myers, A. D., Brunner, R. J., Nichol, R. C., Richards, G. T., Schneider, D. P., & Bahcall, N. A. 2007, *ApJ*, 658, 85
- Myers, A. D., et al. 2006, *ApJ*, 638, 622
- Narayan, R., & Yi, I. 1994, *ApJ*, 428, L13
- Narayan, R., Yi, I., & Mahadevan, R. 1995, *Nature*, 374, 623
- 1996, *A&AS*, 120, C287+
- Norman, C., & Scoville, N. 1988, *ApJ*, 332, 124
- Porciani, C., Magliocchetti, M., & Norberg, P. 2004, *MNRAS*, 355, 1010
- Porciani, C., & Norberg, P. 2006, *MNRAS*, 371, 1824
- Reynolds, C. S., & Begelman, M. C. 1997, *ApJ*, 487, L135+
- Richards, G. T., et al. 2006, *ApJS*, 166, 470
- Salpeter, E. E. 1964, *ApJ*, 140, 796
- Salucci, P., Szuszkiewicz, E., Monaco, P., & Danese, L. 1999, *MNRAS*, 307, 637
- Sazonov, S. Y., Ostriker, J. P., Ciotti, L., & Sunyaev, R. A. 2005, *MNRAS*, 358, 168
- Scannapieco, E., & Oh, S. P. 2004, *ApJ*, 608, 62
- Scheuer, P. A. G. 1995, *MNRAS*, 277, 331
- Schirber, M., Miralda-Escud , J., & McDonald, P. 2004, *ApJ*, 610, 105
- Shakura, N. I., & Sunyaev, R. A. 1973, *A&A*, 24, 337
- Shankar, F., Bernardi, M., & Haiman, Z. 2009a, *ApJ*, 694, 867
- Shankar, F., Salucci, P., Granato, G. L., De Zotti, G., & Danese, L. 2004, *MNRAS*, 354, 1020
- Shankar, F., Weinberg, D. H., & Miralda-Escud , J. 2009b, *ApJ*, 690, 20
- Shen, Y., et al. 2007, *AJ*, 133, 2222
- Siana, B., et al. 2008, *ApJ*, 675, 49
- Sijacki, D., Springel, V., di Matteo, T., & Hernquist, L. 2007, *MNRAS*, 380, 877
- Silk, J., & Rees, M. J. 1998, *A&A*, 331, L1
- Silverman, J. D., et al. 2005, *ApJ*, 624, 630
- 2008a, *ApJ*, 675, 1025

- . 2008b, *ApJ*, 679, 118
Simpson, C. 2005, *MNRAS*, 360, 565
Sokasian, A., Abel, T., & Hernquist, L. 2002, *MNRAS*, 332, 601
—. 2003, *MNRAS*, 340, 473
Soltan, A. 1982, *MNRAS*, 200, 115
Springel, V., Di Matteo, T., & Hernquist, L. 2005a, *ApJ*, 620, L79
—. 2005b, *MNRAS*, 361, 776
Ueda, Y., Akiyama, M., Ohta, K., & Miyaji, T. 2003, *ApJ*, 598, 886
Vestergaard, M., & Peterson, B. M. 2006, *ApJ*, 641, 689
Worseck, G., Fechner, C., Wisotzki, L., & Dall'Aglio, A. 2007, *A&A*, 473, 805
Worseck, G., & Wisotzki, L. 2006, *A&A*, 450, 495
Wyithe, J. S. B., & Loeb, A. 2002, *ApJ*, 581, 886
Younger, J. D., Hopkins, P. F., Cox, T. J., & Hernquist, L. 2008, *ApJ*, 686, 815
Yu, Q., & Lu, Y. 2004, *ApJ*, 602, 603
—. 2005, *ApJ*, 620, 31
—. 2008, *ApJ*, 689, 732
Yu, Q., Lu, Y., & Kauffmann, G. 2005, *ApJ*, 634, 901
Yu, Q., & Tremaine, S. 2002, *MNRAS*, 335, 965
Yuan, F., & Narayan, R. 2004, *ApJ*, 612, 724
Zakamska, N. L., Gómez, L., Strauss, M. A., & Krolik, J. H. 2008, *AJ*, 136, 1607
Zakamska, N. L., et al. 2006, *AJ*, 132, 1496