

The origin of fast molecular outflows in quasars: molecule formation in AGN-driven galactic winds

Alexander J. Richings[★] and Claude-André Faucher-Giguère

Center for Interdisciplinary Exploration and Research in Astrophysics (CIERA) and Department of Physics and Astronomy, Northwestern University, 2145 Sheridan Road, Evanston, IL 60208, USA

Accepted 2017 November 15. Received 2017 November 15; in original form 2017 June 9

ABSTRACT

We explore the origin of fast molecular outflows that have been observed in active galactic nuclei (AGNs). Previous numerical studies have shown that it is difficult to create such an outflow by accelerating existing molecular clouds in the host galaxy, as the clouds will be destroyed before they can reach the high velocities that are observed. In this work, we consider an alternative scenario where molecules form *in situ* within the AGN outflow. We present a series of hydro-chemical simulations of an isotropic AGN wind interacting with a uniform medium. We follow the time-dependent chemistry of 157 species, including 20 molecules, to determine whether molecules can form rapidly enough to produce the observed molecular outflows. We find H₂ outflow rates up to 140 M_⊙ yr⁻¹, which is sensitive to density, AGN luminosity, and metallicity. We compute emission and absorption lines of CO, OH, and warm (a few hundred K) H₂ from the simulations in post-processing. The CO-derived outflow rates and OH absorption strengths at solar metallicity agree with observations, although the maximum line-of-sight velocities from the model CO spectra are a factor ≈2 lower than is observed. We derive a CO (1–0) to H₂ conversion factor of $\alpha_{\text{CO}(1-0)} = 0.13 \text{ M}_{\odot} (\text{K km s}^{-1} \text{ pc}^2)^{-1}$, 6 times lower than is commonly assumed in observations of such systems. We find strong emission from the mid-infrared lines of H₂. The mass of H₂ traced by this infrared emission is within a few per cent of the total H₂ mass. This H₂ emission may be observable by *JamesWebbSpaceTelescope*.

Key words: astrochemistry – molecular processes – ISM: molecules – quasars: absorption lines – quasars: emission lines.

1 INTRODUCTION

Recent observations of galaxies hosting luminous active galactic nuclei (AGNs) have found fast outflows of molecular gas. They have been detected in several molecular lines, including CO (Feruglio et al. 2010, 2015; Cicone et al. 2012; Sun et al. 2014), OH (Fischer et al. 2010; Veilleux et al. 2013; Spoon et al. 2013; Stone et al. 2016), warm H₂ emission (Rupke & Veilleux 2013a), and dense tracers such as HCN, HNC, and HCO⁺ (Aalto et al. 2012, 2015).

These outflows are spatially extended on kpc scales, with typical outflow rates of ∼100–1000 M_⊙ yr⁻¹ and velocities up to ∼1000 km s⁻¹. While molecular outflow rates from starburst-dominated galaxies are typically comparable to the star formation rate, those in AGN-dominated systems can significantly exceed the star formation rate of the host galaxy (e.g. Sturm et al. 2011; Cicone et al. 2014; Fiore et al. 2017). Furthermore, the velocity of the outflow is found to correlate with the AGN luminosity, with

the fastest outflows only found in the most luminous quasars. This suggests that the fast molecular outflows in these systems are being driven by the AGN. Observations of molecular outflows have also derived gas depletion time-scales of ∼10⁶–10⁸ yr, which has been taken as evidence for negative feedback from AGN rapidly quenching star formation in the host galaxy (Sturm et al. 2011; Cicone et al. 2014).

Galaxy-scale outflows driven by AGN have also been observed in both the neutral and ionized phases (e.g. Moe et al. 2009; Rupke & Veilleux 2011; Greene, Zakamska & Smith 2012; Harrison et al. 2012, 2014; Liu et al. 2013). Interestingly, the mass outflow rates in the different phases of AGN-driven galaxy-scale outflows appear roughly comparable (e.g. Faucher-Giguère, Quataert & Murray 2012; Fiore et al. 2017). Furthermore, luminous quasars appear to drive such galaxy-scale, cool outflows at all redshifts, from the local Universe out to $z > 6$ (Maiolino et al. 2012; Cicone et al. 2015).

Measurements of molecular outflows can be used to deduce the energetics of the AGN wind. González-Alfonso et al. (2017) measured OH lines in 14 local Ultra Luminous InfraRed Galaxies

[★] E-mail: a.j.richings@northwestern.edu

(ULIRGs) with clear evidence for outflows. By modelling these lines with radiative transfer models, they calculated the energetics of the outflows. In most sources, they found momentum rates of $\approx 2 - 5 L_{\text{IR}}/c$. However, some showed higher momentum boosts, up to ≈ 20 , which may imply an energy-conserving phase of expansion (Faucher-Giguère & Quataert 2012, FGQ12 hereafter; Zubovas & King 2012; Costa, Sijacki & Haehnelt 2014; Tombesi et al. 2015; Stern et al. 2016). Cicone et al. (2014) also found momentum boosts of ≈ 20 in their AGN-dominated systems.

Understanding the energetics of AGN winds is important for determining how AGNs interact with their host galaxy. It has been shown that feedback from AGN can play a role in quenching star formation in the most massive galaxies, and is needed by galaxy formation models to reproduce the massive end of the observed galaxy stellar mass function (Springel, Di Matteo & Hernquist 2005; Hopkins et al. 2008; Booth & Schaye 2009; Bower, Benson & Crain 2012; Crain et al. 2015; Tremmel et al. 2017; Bower et al. 2017), and to reproduce the colours of massive galaxies (Gabor & Davé 2012; Henriques et al. 2015; Feldmann et al. 2017; Trayford et al. 2016). Note that the AGN feedback that operates in the most massive haloes (i.e. galaxy clusters) is likely to be driven by relativistic jets (e.g. Croton et al. 2006), rather than the type of AGN outflows that we study in this paper.

It has also been suggested that AGN feedback can shape the observed scaling relations between supermassive black holes (SMBHs) and their host galaxies (Silk & Rees 1998; Wyithe & Loeb 2003; Di Matteo, Springel & Hernquist 2005; Murray et al. 2005), although Anglés-Alcázar, Özel & Davé (2013); Anglés-Alcázar et al. (2015, 2017) found that the SMBH scaling relations in their cosmological simulations arose primarily due to feeding of gas on to the SMBH via gravitational torques, with only a weak dependence on feedback.

While there is growing evidence for fast molecular outflows driven by AGN, it is still not understood how these molecular outflows are formed. Narayanan et al. (2006) and Narayanan et al. (2008) explored the effects of galactic winds (both star formation- and AGN-driven) on molecular gas emission from CO using hydrodynamic simulations of galaxy mergers, and they showed that the resulting molecular outflows in their simulations produced detectable CO emission. However, their simulations simply assumed that half of the cold neutral gas was molecular, and did not model the non-equilibrium chemistry of CO.

The origin of the observed fast molecular outflows thus remains unclear. One possibility is that the hot, high-velocity wind launched from regions close to the AGN accelerates existing molecular clouds in the host galaxy. For example, Gaspari & Sadowski (2017) considered the entrainment of hot ($\sim 10^7$ K), warm ($\sim 10^4$ – 10^5 K), and cold ($\lesssim 100$ K) gas from the host galaxy by an AGN wind in their model for AGN feeding and feedback. They argued that such entrainment would produce a multiphase outflow, with velocities and outflow rates in the hot/warm/cold phases of a few $\times 10^3/10^3/500$ km s $^{-1}$ and 10/100/several 100 M $_{\odot}$ yr $^{-1}$, respectively, in good agreement with observed multiphase outflows. However, numerical studies of the interaction of a hot, fast wind with a cold cloud find that the cloud will tend to be rapidly destroyed by the hot wind, before it can be accelerated to the high velocities that have been observed (Klein, McKee & Colella 1994; Scannapieco & Brüggén 2015; Brüggén & Scannapieco 2016; Schneider & Robertson 2017; Zhang et al. 2017). Furthermore, even if a dense molecular cloud is long-lived, it is difficult to accelerate it significantly due to its small cross-section.

Alternatively, the molecules may have formed *in situ*, within the material swept up by the AGN-driven outflow. It has been shown that galactic winds can efficiently cool under certain conditions (Wang 1995; Silich, Tenorio-Tagle & Muñoz-Tuñón 2003; Silich, Tenorio-Tagle & Muñoz-Tuñón 2004; Tenorio-Tagle et al. 2007; Thompson et al. 2016; Scannapieco 2017), which may explain the presence of cool gas in observed galactic outflows (e.g. Martin et al. 2015). If this cooling gas can form molecules rapidly enough, it could also produce the observed fast molecular outflows. Neufeld & Dalgarno (1989) also demonstrated that molecules can form in the cooling gas behind a dissociative shock, although they explored a very different physical regime, with velocities 60–100 km s $^{-1}$ and pre-shock densities 10^{4-6} cm $^{-3}$.

Ferrara & Scannapieco (2016) recently explored this alternative scenario for forming molecular gas within AGN outflows. They ran a series of hydrodynamic simulations of a 1 kpc patch of the shell of gas swept up from the ambient ISM by an AGN outflow, including a treatment for dust destruction via sputtering. They found that, while the gas can cool to $\sim 10^4$ K, it remains as a single phase, rather than forming a two-phase medium, and clumps of dense gas are smoothed out by pressure gradients. They also found that dust grains are rapidly destroyed, within $\sim 10^4$ yr. They therefore concluded that clumps of molecular gas cannot condense out of the AGN outflow material. However, their study only focused on the formation of cold clumps and the survivability of dust grains, and did not explicitly follow the time-dependent molecular chemistry. Furthermore, their simulations only followed a patch of the swept-up shell in the AGN outflow, rather than the full AGN wind structure and its interaction with the ambient ISM, which will miss the effects of the bulk motions of the outflowing material. Additionally, they only considered one ambient medium density and one forward shock velocity, so it remains unclear whether different conditions may be more conducive to the formation of molecular clumps. Therefore, this scenario requires further investigation before we can discount it altogether.

In this paper we investigate *in situ* molecule formation in AGN winds using a series of hydro-chemical simulations of an isotropic AGN wind interacting with a uniform ambient ISM. These simulations include a treatment for the time-dependent chemistry of the gas, which allows us to follow the formation and destruction of molecules, including H $_2$ and commonly observed molecules such as CO, OH, and HCO $^+$. The chemical model is valid over a wide range of physical conditions, from cold, molecular gas to hot, highly ionized gas. This enables us to self-consistently follow the chemistry (along with associated radiative cooling and heating processes) in all phases of the AGN wind.

We consider an idealized set-up, with a uniform ambient ISM, as this matches the set-up studied analytically by FGQ12, providing us with a useful framework in which to interpret our results. The simplified geometry also allows us to focus on the chemistry in this paper, which is the most novel aspect of our simulations. We explore a range of ambient ISM densities (1–10 cm $^{-3}$), AGN luminosities (10^{45} – 10^{46} erg s $^{-1}$), and metallicities (0.1–1 Z $_{\odot}$ ¹), to determine under what conditions molecular outflows are likely to form.

We also perform radiative transfer calculations on our simulations in post-processing to compute observable molecular lines of CO, OH, and IR-traced warm H $_2$. These allow us to rigorously compare

¹ Throughout this paper we use solar abundances taken from table 1 of Wiersma, Schaye & Smith (2009), with a solar metallicity Z $_{\odot} = 0.0129$.

our simulations to observations, and to compute the conversion factors between CO emission and H₂ mass.

The remainder of this paper is divided into the following sections. We describe the AGN wind model, the chemistry solver module and the simulation suite in Section 2. In Section 3, we explore the effects of the ambient medium density, AGN luminosity, and metallicity on the molecular content of AGN winds. We compare variations of the model in Section 4 to quantify how uncertainties in the model parameters affect our results. In Section 5 we compute emission and absorption lines of CO, OH, and IR-traced H₂ from the simulations in post-processing, which we compare to observations. We summarize our main results in Section 6. In Appendix A we present resolution tests, and we explore the uncertainties in the local shielding approximation used in the simulations in Appendix B.

2 SIMULATIONS

We ran a series of 3D hydro-chemical simulations of an idealized, isotropic AGN wind interacting with an initially uniform ambient ISM. We used the hydrodynamics+gravity code GIZMO (Hopkins 2015) with the Meshless Finite Mass (MFM) hydro solver option. MFM is a Lagrangian hydrodynamics method that combines many of the advantages of grid-based methods (such as the accurate capture of shocks and fluid instabilities) and Smoothed Particle Hydrodynamics (SPH) methods (such as exact conservation of mass, energy, and momentum, and better angular momentum conservation); see Hopkins (2015) for more details.

To follow the chemical evolution of molecules and ions in the simulations, we have coupled GIZMO to the CHIMES chemistry solver module (Richings, Schaye & Oppenheimer 2014a,b). The chemical and thermal processes that are included in CHIMES are summarized in Section 2.2 below.

We simulate a periodic box 1.6–5.0 kpc across. Each simulation is run for 1 Myr, which is approximately the time it would take an outflow of 1000 km s⁻¹ to reach 1 kpc. For comparison, González-Alfonso et al. (2017) modelled the molecular outflows observed in OH lines from a sample of 14 ULIRGs that host AGN, and they found flow times (r/v) of 0.1–2 Myr. The box size for each simulation was chosen to ensure that the outflow does not reach the edge of the box within 1 Myr. The gas is initially in thermal and chemical equilibrium in the presence of the redshift zero extragalactic UV background of Haardt & Madau (2001), i.e. before the radiation from the AGN is switched on.

To reduce the computational cost, we only simulate one octant of the box at high resolution, with the remainder of the box at a factor 8 lower mass resolution. We use a fiducial resolution of 30 M_⊙ per gas particle, with 32 kernel neighbours, in the high-resolution region. Some runs use a lower resolution of 240 M_⊙ per particle, and we include one high-resolution run with 10 M_⊙ per particle (see Section 2.3).

We include gravity from a central black hole using a single collisionless particle with a mass $M_{\text{BH}} = 10^8 M_{\odot}$. We also include a static potential from an isothermal sphere to represent the galaxy potential, with an enclosed mass at radius R of:

$$M_{\text{gal}}(< R) = \frac{2\sigma^2 R}{G}. \quad (1)$$

We use a velocity dispersion $\sigma = 200 \text{ km s}^{-1}$, which corresponds to a black hole mass of $10^8 M_{\odot}$ on the $M_{\text{BH}} - \sigma$ relation (e.g. Gültekin et al. 2009).

We also include self-gravity of the gas. For gas particles, we use an adaptive gravitational softening that is set to be equal to the

inter-particle spacing, which is defined as $h_{\text{inter}} = (m_{\text{gas}}/\rho)^{1/3}$ for a gas particle with mass m_{gas} and density ρ . We impose a minimum gravitational softening of 0.1 pc at the fiducial mass resolution. The gravitational softening of the black hole particle is fixed at 1 pc.

Note that, since we consider a uniform ambient medium, these simulations may represent an outflow from a buried quasar nucleus. Once the outflow breaks out of the galactic disc, the density of the ambient medium will decrease rapidly and the wind bubble may lose its internal pressure support when the hot gas vents out (e.g. Stern et al. 2016).

2.1 AGN wind model

We inject an isotropic AGN wind at the centre of the simulation box by spawning new gas particles within the central 1 pc with an outward radial velocity $v_{\text{in}} = 30\,000 \text{ km s}^{-1}$. Such velocities have been observed in broad absorption line quasars (e.g. Weymann, Carswell & Smith 1981; Gibson et al. 2009), and in ultra-fast outflows from quasars observed in X-rays (e.g. Feruglio et al. 2015; Nardini et al. 2015; Tombesi et al. 2015), and could be launched, for example, from the accretion disc around the black hole (Murray et al. 1995), although we do not explicitly model the launching mechanism in this work. Following FGQ12, we parametrize the momentum injection rate of the AGN wind as:

$$\dot{M}_{\text{in}} v_{\text{in}} = \tau_{\text{in}} \frac{L_{\text{AGN}}}{c}, \quad (2)$$

where \dot{M}_{in} is the mass injection rate of the wind and L_{AGN} is the AGN luminosity. τ_{in} is a parameter that is related to the mass loading of the wind (see equation 3 of FGQ12); we use $\tau_{\text{in}} = 1$.

For our choice of v_{in} and τ_{in} , and a given L_{AGN} , we can calculate \dot{M}_{in} from equation 2. For each hydrodynamic time-step we can then calculate the mass of gas to inject. We spawn wind particles in pairs oriented in opposite directions, so that their net momentum vector is zero. If the injected mass in a time-step exceeds twice the gas particle mass in the high-resolution region, we spawn pairs of gas particles with random orientations at random distances between 0 and 1 pc from the centre of the box. All wind particles have a fixed mass equal to that of the high-resolution particles, even if they are injected into a low-resolution region. Wind particles are initially cool (10^4 K), and will later be shock-heated by the reverse shock. Any injected mass that is insufficient to create a pair of gas particles is carried over to later time-steps. For a more detailed discussion of the different components of the wind bubble that results, we refer readers to FGQ12.

2.2 Chemistry and cooling

We use the CHIMES chemistry solver module (Richings et al. 2014a,b) to follow the evolution of 157 chemical species, including all ionization states of 11 elements² and 20 molecular species.³ CHIMES also calculates the radiative cooling and heating rates using the non-equilibrium chemical abundances. This chemical model covers a wide range of physical conditions relevant to AGN winds, from hot, highly ionized gas to the cold, CO-bright molecular gas that is the focus of this study. The chemical model gives us a system of 158 coupled differential equations (157 chemical rate equations

² H, He, C, N, O, Ne, Mg, Si, S, Ca, Fe.

³ H₂, H₂⁺, H₃⁺, OH, H₂O, C₂, O₂, HCO⁺, CH, CH₂, CH₃⁺, CO, CH⁺, CH₂⁺, OH⁺, H₂O⁺, H₃O⁺, CO⁺, HOC⁺, O₂⁺.

and the thermal energy equation). Each time-step CHIMES integrates these equations for every active gas particle using the CVOICE non-linear differential equation solver from the SUNDIALS⁴ suite. For this integration, we use relative and absolute tolerances of 10^{-4} and 10^{-10} , respectively, for most particles. However, for particles with temperatures $T > 10^6$ K or densities $n_{\text{H}} < 50 \text{ cm}^{-3}$, we found it necessary to use a relative tolerance of 10^{-6} . We impose a cooling floor at $T = 10$ K.

A full description of the chemical reactions and thermal processes included in CHIMES can be found in Richings et al. (2014a). For this work, we have updated the photoelectric heating rate from dust grains used in CHIMES to that given in equations 19 and 20 of Wolfire et al. (2003), which better models the effects of PAHs. Below we summarize the most important thermo-chemical processes that are relevant for this study.

H_2 formation. We include the formation of H_2 in the gas phase (via H^-) and on dust grains. For the formation rate of H_2 on dust grains, we use equation 18 of Cazaux & Tielens (2002), where we scale the dust abundance linearly with metallicity. This assumes that the dust-to-metals ratio is constant. However, it is unclear whether dust grains can survive the strong shocks and high gas temperatures found in AGN winds, and some studies have suggested that the dust grains will be rapidly destroyed in such conditions (e.g. Ferrara & Scannapieco 2016). As we will discuss further in Section 4, the formation of molecules in an AGN wind is sensitive to the assumed dust-to-metals ratio, and this is therefore a major uncertainty in our model.

Photoionization and photoheating. We include the photoionization of atoms and ions and the photodissociation of molecules, along with the associated heating, due to the AGN radiation field. To calculate the photochemical rates, we use the spectrum of an average quasar from Sazonov, Ostriker & Sunyaev (2004), scaled to a given bolometric AGN luminosity L_{AGN} . At a radius R from the black hole, the flux of hydrogen-ionizing photons (with energy $h\nu > 1$ Ryd) is then:

$$f_{\text{ion}} = 1.62 \times 10^{11} \text{ cm}^{-2} \text{ s}^{-1} \left(\frac{L_{\text{AGN}}}{10^{46} \text{ erg s}^{-1}} \right) \left(\frac{R}{\text{kpc}} \right)^{-2}. \quad (3)$$

To model shielding of gas from the AGN radiation, we use a local shielding approximation that assumes that a gas particle is shielded by material within a shielding length L_{sh} of the particle. The column density of species i , with volume density n_i , is then:

$$N_i = n_i L_{\text{sh}}. \quad (4)$$

This also assumes that n_i is uniform over the shielding length.

For the shielding length, we use a Sobolev-like approximation based on the local gradient of the density, ρ :

$$L_{\text{sh}} = \frac{1}{2} \left(h_{\text{inter}} + \frac{\rho}{|\nabla \rho|} \right), \quad (5)$$

where the inter-particle spacing h_{inter} for a gas particle of mass m_{gas} is defined such that $\rho = m_{\text{gas}}/h_{\text{inter}}^3$, and is related to the kernel smoothing length h_{sml} and the number of neighbours N_{ngb} as $h_{\text{inter}} = h_{\text{sml}}/N_{\text{ngb}}^{1/3}$. This approximation has previously been used in cosmological simulations (e.g. Gnedin, Tassis & Kravtsov 2009; Hopkins et al. 2017), and simulations of isolated disc galaxies that use the CHIMES chemical model (Richings & Schaye 2016), although not always with the h_{inter} term.

Safraneck-Shrader et al. (2017) recently compared several local approximations for radiative shielding, including the density gradient-based approach that we use, in galactic disc simulations in post-processing. They found that, of the methods that depend only on local quantities, using the Jeans length, with a temperature capped at 40 K, gave the most accurate results for H_2 and CO fractions, as compared to a full radiative transfer calculation. However, in Appendix B we compare these two local approximations in our simulations to the true column densities computed from a simple ray tracing algorithm in post-processing, and we find that the density gradient-based approach is the most accurate. This is likely because the AGN winds that we simulate are a very different environment to the turbulent galactic discs that Safraneck-Shrader et al. (2017) studied, so it is possible that the Jeans length assumption is not valid for cold clumps in an AGN wind. Nevertheless, there is still a large scatter of an order of magnitude between the density gradient-based approximation and the true ray tracing-based column densities. We explore the impact of these uncertainties on our results in Section 4.

Once we have calculated the column densities for a given particle from equations 4 and 5, we then attenuate the photochemical and photoheating rates as a function of the column densities of H I , H_2 , He I , He II , CO , and the dust extinction, A_{v} , as described in Richings et al. (2014b). We use $A_{\text{v}} = 4.0 \times 10^{-22} N_{\text{Htot}} Z/Z_{\odot}$, where N_{Htot} is the total hydrogen column density. This is an average of the dust models of Weingartner & Draine (2001) for the Milky Way and the Large and Small Magellanic Clouds (see appendix A of Krumholz, Leroy & McKee 2011). This also assumes a constant dust-to-metals ratio, which is a major uncertainty in this work, as discussed above.

We use a temperature-dependent H_2 self-shielding function that was fit to a series of photodissociation region (PDR) models created with version 13.01 of CLOUDY⁵ (last described by Ferland et al. 2013), as described in appendix B of Richings et al. (2014b). Doppler broadening of the Lyman-Werner lines, for example by turbulence, tends to weaken the self-shielding of H_2 (e.g. Draine & Bertoldi 1996; Ahn & Shapiro 2007). We assume a constant turbulent broadening parameter $b_{\text{turb}} = 7.1 \text{ km s}^{-1}$, which is typical for molecular clouds in the Milky Way. This is added to the thermal broadening in quadrature to give the total broadening parameter, which is used in the self-shielding function.

Compton cooling and heating from AGN radiation. The original CHIMES model only included Compton cooling from the cosmic microwave background. For this work, we have added non-relativistic Compton cooling and heating from the AGN radiation field, with a Compton temperature $T_{\text{C}} = 2 \times 10^7$ K (Sazonov et al. 2004). The cooling rate (where a negative value corresponds to heating) for an AGN luminosity L_{AGN} at a distance R is then:

$$\Lambda_{\text{Compton}} = \frac{\sigma_{\text{T}}}{m_{\text{e}} c^2} \left(\frac{L_{\text{AGN}}}{\pi R^2} \right) k_{\text{B}} (T - T_{\text{C}}) n_{\text{e}} \text{ erg s}^{-1} \text{ cm}^{-3}, \quad (6)$$

where σ_{T} , m_{e} , and k_{B} are the Thomson cross-section, electron mass, and Boltzmann constant, respectively, and T and n_{e} are the temperature and electron density of the gas.

In cold gas ($T < T_{\text{C}}$), this will be a source of heating. At the high post-shock temperatures corresponding to the wind velocity, $v_{\text{in}} = 30\,000 \text{ km s}^{-1}$, used in our simulations ($\sim 10^{10}$ K), the non-relativistic assumption can break down. However, FGQ12 argued that, in this regime, the temperatures of electrons and protons become thermally decoupled, which leads to two-temperature effects that cause the shocked electrons to effectively remain

⁴ <https://computation.llnl.gov/casc/sundials/main.html>

⁵ <http://nublado.org/>

non-relativistic and the cooling time of the protons to become limited by Coulomb interactions with the electrons. This generally increases the cooling time of the shocked wind by a large factor relative to relativistic Compton cooling. Since we do not model 2-temperature effects in our simulations, we would overestimate the cooling in the hot AGN wind if we used the relativistic Compton cooling rate. We therefore use the non-relativistic rate given above. Since we find that the shocked wind does not significantly cool in our simulations, our results are not sensitive to the exact details of how Compton cooling is modelled.

Cosmic rays. We include ionization, dissociation of molecules and heating due to cosmic rays. The strength of the cosmic rays is parametrized by the primary ionization rate of H I by cosmic rays, ζ_{HI} . The ionization and dissociation rates of other species are then scaled to the H I rate according to their relative rates in the UMIST data base⁶ (McElroy et al. 2013) where available, or using the equations in Lotz (1967), Silk (1970), and Langer (1978) otherwise. Secondary ionizations of H I and He I from cosmic rays are calculated using the tables of Furlanetto & Stoever (2010), assuming a primary electron energy of 35 eV.

The cosmic ray ionization rate of H I in the Milky Way, inferred from observations of H₃⁺, is $\zeta_{\text{HI}}^{\text{MW}} = 1.8 \times 10^{-16} \text{ s}^{-1}$ (Indriolo & McCall 2012). However, the quasar host galaxies where fast molecular outflows have been observed are typically ULIRGs, with higher star formation rates than the Milky Way. We therefore expect that the cosmic ray density will also be higher in these ULIRGs than in the Milky Way. Lacki, Thompson & Quataert (2010) present 1D models of cosmic ray diffusion in star-forming galaxies. Their fig. 15 shows the cosmic ray energy density, U_{CR} , as a function of gas surface density, Σ_{g} , in their models. The gas surface density in the Milky Way is $\Sigma_{\text{g}}^{\text{MW}} \approx 10 M_{\odot} \text{ pc}^{-2} = 0.002 \text{ g cm}^{-2}$, while ULIRGs have surface densities up to $\Sigma_{\text{g}}^{\text{ULIRG}} \approx 10^4 M_{\odot} \text{ pc}^{-2} = 2 \text{ g cm}^{-2}$ (see fig. 11a of Kennicutt & Evans 2011). If we assume that $\zeta_{\text{HI}} \propto U_{\text{CR}}$ then, using fig. 15 of Lacki et al. (2010) to get U_{CR} , and with $\zeta_{\text{HI}}^{\text{MW}}$ from Indriolo & McCall (2012), we find $\zeta_{\text{HI}}^{\text{ULIRG}} = 5.7 \times 10^{-15} \text{ s}^{-1}$. We use this value in most of our simulations, although we include one run with the Milky Way value to explore how uncertainties in this rate affect our results.

2.3 Simulation suite

Table 1 summarizes the physical parameters of our simulations. For our fiducial model, we considered an AGN with a bolometric luminosity $L_{\text{AGN}} = 10^{46} \text{ erg s}^{-1}$, embedded in an ambient ISM with a uniform initial density $n_{\text{H}} = 10 \text{ cm}^{-3}$ and solar metallicity. These conditions are typical for the ULIRG host galaxies in which fast molecular outflows have been observed. Note that, while mean densities in the central $\approx 500 \text{ pc}$ of ULIRGs are typically much higher than this ($\sim 10^3\text{--}10^4 \text{ cm}^{-3}$; e.g. Downes & Solomon 1998), measured outflow parameters suggest that, on average, the outflows must have expanded along underdense paths (see FGQ12). Our spherically symmetric initial conditions may be viewed as an approximation to AGN winds expanding in the nuclei of buried quasars, in which all directions from the central AGN are significantly covered by ambient gas. Observations offer some evidence that this set-up is appropriate for understanding the generation of quasar-driven molecular outflows. For example, González-Alfonso et al. (2017) note that the highest molecular outflow velocities are found in buried sources.

This fiducial model is labelled nH10_L46_Z1. We then ran a further three simulations where the ambient ISM density, AGN luminosity, and metallicity were each reduced in turn by a factor of 10. This set of four simulations, run at our fiducial resolution of $30 M_{\odot}$ per gas particle, allows us to explore how molecule formation is affected by the physical parameters of the system. These parameter variations are presented in Section 3.

To test the effects of numerical resolution on our results, we repeated these four simulations with a lower resolution of $240 M_{\odot}$ per gas particle. We also performed one high-resolution run, nH10_L45_Z1_hiRes, with a resolution of $10 M_{\odot}$ per gas particle. We used the low-luminosity model for this high-resolution run, rather than our fiducial model, as it was computationally cheaper. These resolution tests are presented in Appendix A.

Finally, we performed a further six runs of nH10_L46_Z1_lowRes in which we varied different aspects of our AGN and chemistry modelling, to explore uncertainties in the models. These runs all used the lower resolution of $240 M_{\odot}$ per particle and are described in more detail below.

lowCR. In our fiducial model we assume a primary ionization rate of H I due to cosmic rays of $\zeta_{\text{HI}} = 5.7 \times 10^{-15} \text{ s}^{-1}$, which we expect to be appropriate for ULIRGs. However, this rate is uncertain as it is based on scaling the observed Milky Way rate up to typical ULIRG conditions using the 1D cosmic ray diffusion models of Lacki et al. (2010). To quantify how uncertainties in this rate might affect molecule formation, we performed one run with the Milky Way rate, $1.8 \times 10^{-16} \text{ s}^{-1}$ (Indriolo & McCall 2012), which is a factor ≈ 30 lower than our fiducial value.

jeansLimit. Bate & Burkert (1997) and Truelove et al. (1997) demonstrated that hydrodynamic simulations that include self-gravity may experience artificial fragmentation if they do not resolve the Jeans scale. Our simulations do not include an explicit Jeans limiter to ensure that the Jeans scale is always resolved. At our fiducial resolution, the thermal Jeans scale becomes unresolved at densities $n_{\text{H}} \gtrsim 10^4 \text{ cm}^{-3}$ at 100 K. Beyond this density, the Jeans mass is smaller than the particle mass, $30 M_{\odot}$. Note that, in the MFM hydro solver, a gas particle is roughly equivalent to a gas cell in a grid-based code. Thus the effective spatial and mass resolution in MFM is the inter-particle spacing and the particle mass, respectively, rather than the kernel smoothing length and the kernel mass as in SPH codes (see section 4.2.1 of Hopkins et al. 2017 for a more detailed discussion of the resolution requirements of MFM).

To investigate whether artificial fragmentation may affect our results, we ran one simulation in which we included a density-dependent pressure floor to ensure that the thermal Jeans scale is always resolved. This method to prevent artificial fragmentation has previously been used in galaxy simulations (e.g. Robertson & Kravtsov 2008; Hopkins, Quataert & Murray 2011; Richings & Schaye 2016). The pressure floor, P_{floor} , was chosen such that the Jeans length is always at least a factor $N_{\text{J},1}$ times the smoothing length, h_{sm1} . Then, from equation 2.14 of Richings & Schaye (2016):

$$P_{\text{floor}} = \frac{G}{\pi\gamma} N_{\text{J},1}^2 h_{\text{sm1}}^2 \rho^2, \quad (7)$$

where ρ is the gas density and the ratio of specific heats is $\gamma = 5/3$. We used $N_{\text{J},1} = 3.2$, which is equivalent to resolving the Jeans mass by 4 times the kernel mass, as used by Richings & Schaye (2016).

noAGNflux. The radiation field of the AGN may have an important effect on the formation of molecular outflows, as it can destroy

⁶ <http://www.udfa.net/>

Table 1. Simulation runs and parameters: initial ambient ISM density n_{H} , bolometric AGN luminosity L_{AGN} , gas metallicity Z , gas particle mass m_{gas} in the high-resolution octant, minimum gravitational softening ϵ_{min} for gas particles, and the side length of the high-resolution octant l_{hiRes} . The model variations are discussed further in the text.

Name	n_{H} (cm $^{-3}$)	L_{AGN} (erg s $^{-1}$)	Z/Z_{\odot}	m_{gas} (M_{\odot})	ϵ_{min} (pc)	l_{hiRes} (kpc)
Parameter variations						
nH10_L46_Z1	10	10^{46}	1	30	0.1	1.2
nH1_L46_Z1	1	10^{46}	1	30	0.1	2.5
nH10_L45_Z1	10	10^{45}	1	30	0.1	0.8
nH10_L46_Z0.1	10	10^{46}	0.1	30	0.1	1.2
Resolution variations						
nH10_L46_Z1_lowRes	10	10^{46}	1	240	0.2	1.2
nH1_L46_Z1_lowRes	1	10^{46}	1	240	0.2	2.5
nH10_L45_Z1_lowRes	10	10^{45}	1	240	0.2	0.8
nH10_L45_Z1_hiRes	10	10^{45}	1	10	0.07	0.8
nH10_L46_Z0.1_lowRes	10	10^{46}	0.1	240	0.2	1.2
Model variations						
lowCR	10	10^{46}	1	240	0.2	1.2
jeansLimit	10	10^{46}	1	240	0.2	1.2
noAGNflux	10	10^{46}	1	240	0.2	1.2
lowShield	10	10^{46}	1	240	0.2	1.2
lowDust10	10	10^{46}	1	240	0.2	1.2
lowDust100	10	10^{46}	1	240	0.2	1.2

molecules via photodissociation, and it can heat the cold gas via photoheating and photoelectric heating from dust grains. It can also affect how the AGN wind evolves, as Compton cooling due to the AGN radiation can enable the hot wind bubble to cool under certain conditions, which will make the wind momentum-driven rather than energy-driven (e.g. King 2003; FGQ12; Costa et al. 2014). We therefore performed one run with the AGN radiation switched off, to investigate the impact that it has on the chemistry and the AGN wind structure.

lowShield. As described in Section 2.2, we use a local approximation for the shielding column density of each gas particle. However, in Appendix B we show that, while this approximation does on average follow the true column density, as computed from ray tracing in post-processing, there is a scatter of a factor of 10 between the two. We therefore ran one simulation in which the local column density was reduced by a factor of 10, to quantify how these uncertainties might affect our results.

lowDust10/100. Dust grains are important for molecule formation as they catalyse H_2 formation and shield molecules from photodissociating radiation. In our chemical model we assume that the dust abundance scales linearly with metallicity. However, it is unclear whether dust grains will survive in AGN winds, as some studies have suggested that dust grains will be rapidly destroyed by shocks and sputtering (e.g. Ferrara & Scannapieco 2016). We therefore ran two simulations with dust abundances reduced by a factor of 10 and 100, to explore how uncertainties in the dust abundance might affect our results. The results of these model variations are presented in Section 4.

3 PARAMETER VARIATIONS

In this section we focus on the four fiducial-resolution simulations in which we vary the ambient ISM density, AGN luminosity, and metallicity. Fig. 1 shows snapshots of the high-resolution octant at the end of each simulation, after 1 Myr. The top and bottom rows show 2D slices of the gas density and temperature, respectively, along the $x = y$ plane, i.e. a vertical plane diagonally

through the high-resolution octant. The spatial extent of these images ranges from 0.6 to 2.0 kpc, as indicated at the bottom of the figure.

We see artefacts near the boundaries between the high- and low-resolution regions (i.e. the left and bottom edges of each panel), which break the spherical symmetry of the outflow. These are particularly noticeable in the nH10_L46_Z1 run (left-hand column). They occur because the kernel used by the MFM hydro solver smoothes over gas particles with different masses near these boundaries, which can lead to inaccuracies in the hydro solver. However, we find that these artefacts are only near the boundaries with the low-resolution regions, and most of the volume is unaffected. For example, in Fig. 2 we show a 3D volume rendering of the gas density in the high-resolution octant of run nH10_L46_Z1. We see in this 3D view that the artefacts are only found along the boundaries with the low-resolution regions.

To ensure that these artefacts do not affect our results, we restrict our analysis throughout this paper to particles in the high-resolution region with a polar and azimuthal angle (in spherical polar coordinates) between 15° and 75° . This defines a wedge through the high-resolution octant that avoids regions near the boundaries with the low-resolution regions. We will refer to this as the high-resolution wedge for the remainder of this paper. For quantities such as molecular masses and outflow rates, we scale up our results for the wedge to what we would get in a full sphere by multiplying by the ratio of the solid angle of a sphere to the solid angle subtended by the wedge (a factor of 17).

In all four simulations we see a bubble of hot ($\sim 10^{10}$ K), low-density ($n_{\text{H}} \sim 10^{-2}$ cm $^{-3}$) gas. This bubble has maintained its high temperature throughout the simulation, so radiative cooling is inefficient here and the outflows are in the energy-driven regime. Except for in the nH10_L45_Z1 run, the central ~ 100 pc is cooler. This corresponds to the position of the reverse shock (which we determine from the radial velocity of wind particles as a function of radius; not shown), and particles within this radius have not yet been shock-heated. In the low-luminosity run, we expect that the inner, pre-shock wind should also remain cool, but this is not apparent in

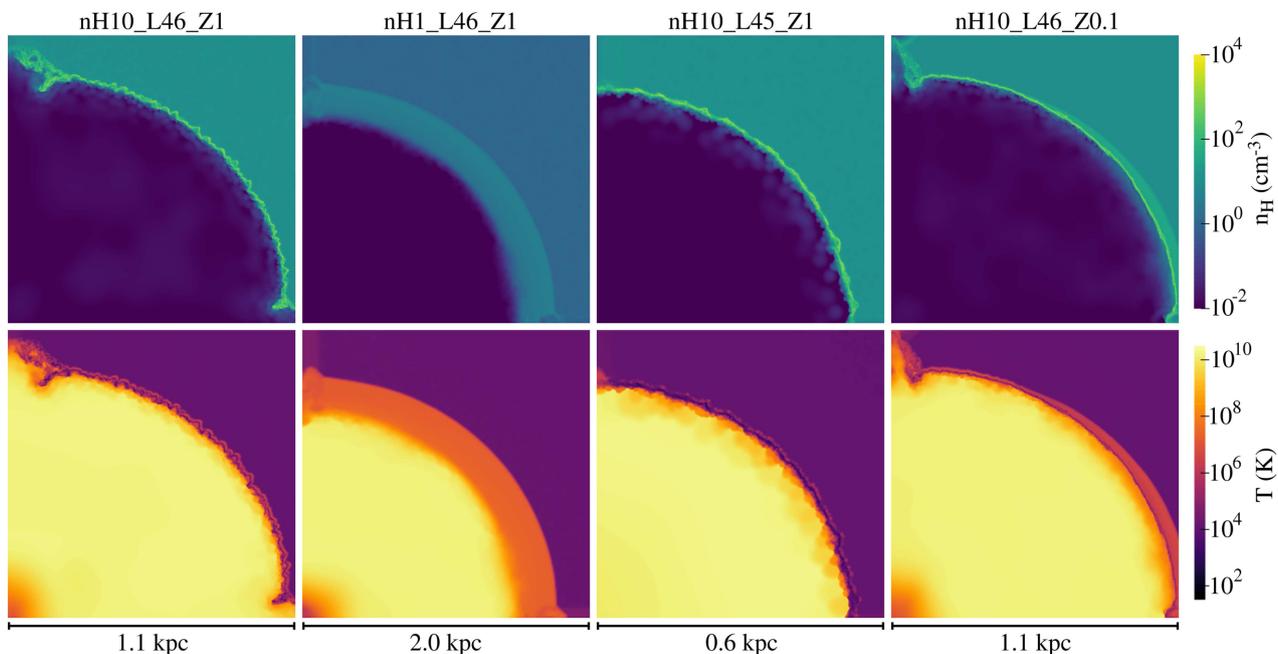


Figure 1. Snapshots from the four ‘parameter variations’ runs at the end of each simulation, after 1 Myr, showing 2D slices of the gas density (top row) and temperature (bottom row). These 2D slices are taken along the $x = y$ plane through the high-resolution octant. We then project the mean density and mass-weighted mean temperature in a slice 0.03 kpc thick on to this plane. The size of each panel is indicated at the bottom of the figure. The outflow is energy-driven in all cases, with a hot bubble that maintains a high temperature of $\sim 10^{10}$ K, surrounded by a thin shell of material swept up from the ambient ISM. This swept-up material is able to radiatively cool within 1 Myr in all runs except the low ambient density run, nH1_L46_Z1.

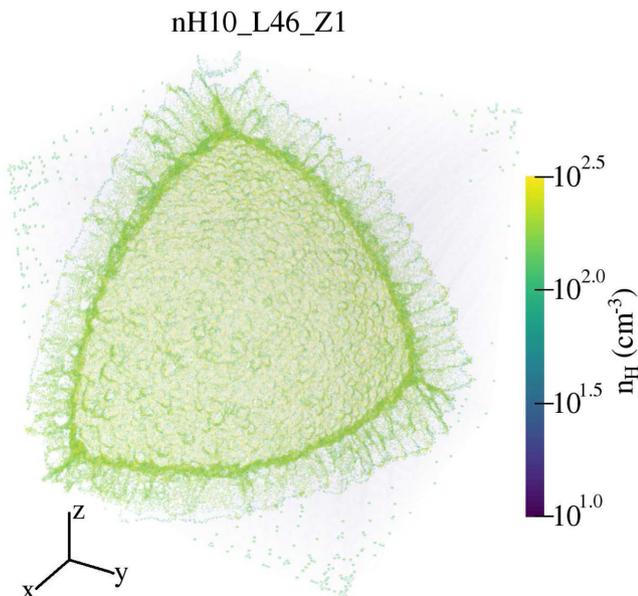


Figure 2. A 3D volume rendering of the gas density in the high-resolution octant of run nH10_L46_Z1. We see density inhomogeneities in the shell of swept-up gas, which are created when this material radiatively cools.

Fig. 1 because there the $10\times$ lower \dot{M}_{in} implies that there are very few resolution elements in the wind at any given time.

Close to the AGN, Compton cooling can potentially cool the hot bubble. This would lead to a momentum-driven wind (e.g. King 2003; Costa et al. 2014; Ferrara & Scannapieco 2016), provided that this occurs beyond the free expansion radius, i.e. the

radius at which the total mass injected by the nuclear wind is equal to the swept-up mass (FGQ12). Using the non-relativistic Compton cooling rate given in equation 6, the cooling time at a radius R from an AGN with luminosity L_{AGN} is:

$$t_{\text{cool}} = \frac{3 nm_e c^2}{2 n_e \sigma_T} \left(\frac{\pi R^2}{L_{\text{AGN}}} \right), \quad (8)$$

where n and n_e are the total (including electrons, H, He, and metals) and electron number densities, respectively. By equating t_{cool} to the flow time, $t_{\text{flow}} = R/v_s$, for the forward shock velocity v_s , we can calculate the radius R_C within which we expect Compton cooling to be efficient:

$$R_C = 9.3 \left(\frac{L_{\text{AGN}}}{10^{46} \text{ erg s}^{-1}} \right) \left(\frac{v_s}{300 \text{ km s}^{-1}} \right)^{-1} \text{ pc}, \quad (9)$$

where we assume $n/n_e \approx 2$ for a fully ionized plasma. For forward shock velocities $v_s \approx 300\text{--}1000 \text{ km s}^{-1}$, as we find in our simulations on scales ~ 1 kpc, we therefore expect Compton cooling to be inefficient beyond ~ 10 pc. At smaller radii, v_s can be higher than the fiducial 300 km s^{-1} in equation 9, so Compton cooling will often also be negligible at smaller radii. This is consistent with FGQ12’s result that AGN-driven wind bubbles should be generically energy-conserving, provided that they are inflated by nuclear winds with high initial velocities $v_{\text{in}} \gtrsim 10\,000 \text{ km s}^{-1}$.

The model of King (2003) suggests that Compton cooling should be efficient out to radii up to ~ 1 kpc, which would result in a momentum-driven wind. This discrepancy arises partly because King (2003) uses the relativistic Compton cooling time, which, for an initial wind velocity $v_{\text{in}} = 30\,000 \text{ km s}^{-1}$, is nearly two orders of magnitude shorter than for non-relativistic Compton cooling, as we use. However, FGQ12 argued that, in the hot shocked AGN wind, the electron and proton temperatures become thermally

decoupled, and two-temperature effects result in cooling times that are much longer than the relativistic Compton cooling time used in King (2003). Additionally, the model of King (2003) assumes an isothermal density profile ($n_{\text{H}} \propto r^{-2}$), whereas our simulations assume a uniform ambient medium.

While the inner wind shock typically does not cool, the outer shock with the ambient medium has a lower velocity ($\sim v_s$) and can efficiently cool at moderately high ambient densities (e.g. Zubovas & King 2014; Nims, Quataert & Faucher-Giguère 2015). At the outer edge of the hot bubble, there is a thin shell of gas swept up from the ambient ISM. In the three runs with an ambient density of 10 cm^{-3} (first, third, and fourth columns in Fig. 1), this shell has cooled to $\sim 10^4 \text{ K}$, and has been compressed, with densities up to $\sim 10^4 \text{ cm}^{-3}$ (we explore the density and temperature distributions in more detail in Section 3.2). However, in run nH1_L46_Z1 the swept-up ISM material is still at $\sim 10^7 \text{ K}$ (which corresponds to the post-shock temperature of the forward shock velocity, $v_s \sim 1000 \text{ km s}^{-1}$), with only a moderate enhancement in density. This is understandable as the cooling time is longer at lower densities, so in the low-density run the swept-up gas has had insufficient time to cool after 1 Myr.

In the nH1_L46_Z1 run, the shell of swept-up gas remains smooth. However, in the other three runs, radiative cooling in the swept-up material leads to compression and collapse of the gas, which creates density inhomogeneities, so the shell is no longer smooth. These inhomogeneities can be seen in both the 2D (Fig. 1) and 3D (Fig. 2) views. It is in this dense, radiatively cooling shell that we might expect molecules to form.

3.1 Molecular masses and outflow rates

Fig. 3 shows the total mass and molecular fractions of H_2 , CO, OH, and HCO^+ (top left, top right, bottom left, and bottom right, respectively) in gas that is outflowing with a radial velocity $v_{\text{out}} > 100 \text{ km s}^{-1}$. The masses are calculated in the high-resolution wedge only and then multiplied by a factor of 17 (the ratio of the solid angle of a sphere to the solid angle subtended by the wedge) to represent the mass we would find in the full spherical shell. The molecular fractions are defined as the fraction of the total hydrogen mass in H_2 , the fraction of the total carbon mass in CO or HCO^+ , or the fraction of the total oxygen mass in OH. These fractions thus go to unity when all of the given element is in the given molecular species.

In our fiducial run, nH10_L46_Z1, the total mass of H_2 in outflowing gas reaches $2.7 \times 10^8 M_{\odot}$ after 1 Myr (top left-hand panel of Fig. 2; red curve). For comparison, observations of ULIRGs that host luminous quasars have found H_2 outflows with masses up to a few times $10^8 M_{\odot}$, as measured from e.g. CO emission (Feruglio et al. 2010; Cicone et al. 2014) or OH emission and absorption (González-Alfonso et al. 2017). The molecular H_2 fraction of outflowing gas in this run is 0.24 after 1 Myr. Rupke & Veilleux (2013b) measure galactic winds in the neutral atomic phase (via Na I D) and the ionized phase (via O I , $\text{H}\alpha$, and N II) in six ULIRGs, of which three are powered by AGN (F08572+3915, Mrk 231, and Mrk 273). Molecular outflows have also been observed in these three systems from CO emission (Cicone et al. 2012, 2014). From these measurements, the molecular H_2 fractions of these observed outflows are 0.26–0.83.

When we decrease the AGN luminosity by a factor of 10 (nH10_L45_Z1; blue curves), the outflowing H_2 mass after 1 Myr decreases by a factor of 4. However, the H_2 fraction of outflowing gas is similar to the fiducial run, so the H_2 mass is lower because

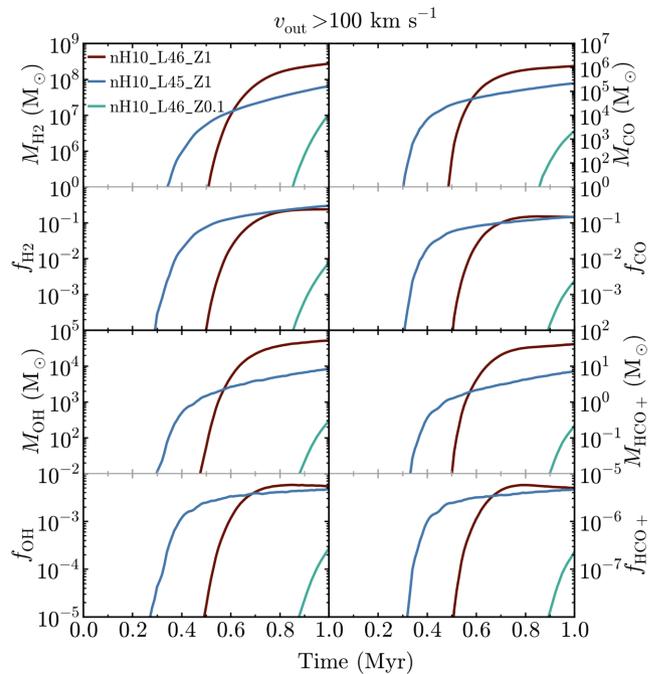


Figure 3. The molecular content of gas outflowing with a velocity $v_{\text{out}} > 100 \text{ km s}^{-1}$. Each pair of panels shows the mass (top) and molecular fraction (bottom) of a given species as a function of time, for H_2 (top left-hand pair of panels), CO (top right), OH (bottom left), and HCO^+ (bottom right). Molecular masses are calculated in the high-resolution wedge and then multiplied by 17 to represent the mass in the full spherical shell. Molecular fractions are defined as the fraction of the total mass of hydrogen (H_2), carbon (CO and HCO^+), or oxygen (OH) in the given molecular species. Our fiducial run (nH10_L46_Z1; red curves) has an H_2 mass of $2.7 \times 10^8 M_{\odot}$ in the outflow after 1 Myr, which decreases with decreasing AGN luminosity (blue curves) and metallicity (green curves). Our low-density run (nH1_L46_Z1) is not shown here, as it did not form any molecules.

the total mass swept up by the outflow is lower at lower AGN luminosities, which in our simulations correspond to less energetic winds.

The H_2 mass in the low-metallicity run (nH10_L46_Z0.1; green curves) is a factor 26 lower after 1 Myr compared to the fiducial run. The H_2 fraction is also correspondingly lower. This is partly due to the lower dust abundance, which is needed to catalyse H_2 formation and to shield H_2 from UV radiation. We will also see in Section 3.2 that there is less cold, dense gas in the nH10_L46_Z0.1 run, which is likely at least partly because H_2 is an important coolant below 10^4 K (in Section 5.3 we will show that warm H_2 emits strongly in infrared lines in our simulations). Metal line cooling is also reduced in the low-metallicity run.

Note that, while the H_2 fraction in outflowing gas has reached a steady state in the fiducial run after 1 Myr, in the low-luminosity and low-metallicity runs it is still increasing at the end of the simulation. In these two runs, the H_2 fraction that we measure therefore depends on the time that we choose to measure it at.

The low-density run (nH1_L46_Z1; not shown) did not form any molecules after 1 Myr. We saw in Fig. 1 that the swept-up material in this run has had insufficient time to cool by the end of the simulation, so there is no cold, dense gas, which is needed for molecules to form.

The masses and fractions of CO, OH, and HCO^+ in outflowing gas show similar trends between the four parameter variations runs. After 1 Myr, we find OH fractions relative to the total oxygen mass

Table 2. OH abundances relative to H₂ after 1 Myr in the three parameter variations runs that form molecules. For comparison, observations typically assume $\chi(\text{OH}) = 5 \times 10^{-6}$ when inferring H₂ masses from OH emission and absorption.

Simulation	$\chi(\text{OH}) = n_{\text{OH}}/n_{\text{H}_2}$
nH10_L46_Z1	2.3×10^{-5}
nH10_L45_Z1	1.6×10^{-5}
nH10_L46_Z0.1	3.3×10^{-6}

($f_{\text{OH}} = (16/17)M_{\text{OH}}/M_{\text{O,tot}}$) of 2.7×10^{-4} – 5.4×10^{-3} (excluding the low-density run, which does not form OH). Observational studies of OH absorption and emission typically fit radiative transfer models to the OH spectra to determine the OH column density, then convert this to the H₂ column density using an assumed OH fraction relative to H₂ ($\chi(\text{OH}) = n_{\text{OH}}/n_{\text{H}_2}$). This OH conversion factor is a ratio of OH to H₂ number densities, unlike the CO conversion factor (α_{CO} , which we explore further in Section 5.1), which is a ratio of H₂ mass to CO luminosity. Observational studies typically assume $\chi(\text{OH}) = 5 \times 10^{-6}$ (e.g. Sturm et al. 2011), which is based on observations of multiple OH lines in the Sagittarius B2 molecular cloud in the Galactic Centre (Goicoechea & Cernicharo 2002).

Using the OH and H₂ fractions from Fig. 3, along with the total oxygen abundance, we calculate $\chi(\text{OH})$ in outflowing gas ($v_{\text{out}} > 100 \text{ km s}^{-1}$) after 1 Myr in each simulation as follows:

$$\chi(\text{OH}) = \frac{n_{\text{OH}}}{n_{\text{H}_2}} = \frac{2}{17} \frac{M_{\text{OH}}}{M_{\text{H}_2}} = \frac{f_{\text{OH}} M_{\text{O,tot}}/16}{f_{\text{H}_2} M_{\text{H,tot}}/2}, \quad (10)$$

where $M_{\text{O,tot}}$ and $M_{\text{H,tot}}$ are the total masses of oxygen and hydrogen, respectively. The values of $\chi(\text{OH})$ from the simulations are shown in Table 2. Based on the fiducial run, which produces the strongest molecular outflows, observations of molecular outflows based on OH may overestimate H₂ masses by up to a factor ≈ 4 . However, we note that González-Alfonso et al. (2017) assume an OH abundance relative to all H nuclei (not just H₂) of 2.5×10^{-6} , which is in good agreement with our simulations.

The value of $\chi(\text{OH})$ that we calculate from the simulations depends on the time at which we calculate it, although 1 Myr corresponds to the typical flow times (r/v_{out}) that are seen in observed AGN-driven molecular outflows (e.g. González-Alfonso et al. 2017).

We can also use our simulations to calculate the mass outflow rate in molecular gas:

$$\begin{aligned} \frac{dM_{\text{H}_2}}{dt} &= 17 \sum_{\text{hiResWedge}} \left(\frac{m_{\text{H}_2}}{t_{\text{flow}}} \right) \\ &= 17 \sum_{\text{hiResWedge}} \left(\frac{m_{\text{H}_2} v_{\text{out}}}{r} \right), \end{aligned} \quad (11)$$

where m_{H_2} is the H₂ mass of a given particle, $t_{\text{flow}} = r/v_{\text{out}}$ is the flow time of the particle, v_{out} is its radial velocity, and r is its radius from the AGN. The summation is over particles in the high-resolution wedge. We include only particles outflowing with a velocity $v_{\text{out}} > 100 \text{ km s}^{-1}$, for consistency with Fig. 3 (although the outflow rates are unchanged if we include all outflowing particles with $v_{\text{out}} > 0 \text{ km s}^{-1}$, since in our set-up the outflowing gas is concentrated in a thin shell). The factor of 17 then gives us the outflow rate from the full spherical shell.

Fig. 4 shows the H₂ mass outflow rate versus time in the three parameter variations runs that form molecules. In our fiducial run (nH10_L46_Z1; red curve), the outflow rate reaches $140 M_{\odot} \text{ yr}^{-1}$.

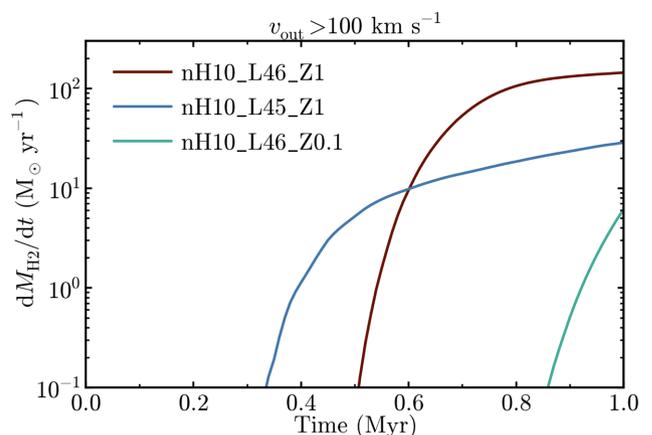


Figure 4. Mass outflow rates of H₂ in the fiducial (nH10_L46_Z1; red curve), low-luminosity (nH10_L45_Z1; blue curve), and low-metallicity (nH10_L46_Z0.1; green curve) runs, plotted versus time. These outflow rates are calculated in the high-resolution wedge and then multiplied by 17 to give the outflow rate in the full spherical shell. We include only particles outflowing with a velocity $v_{\text{out}} > 100 \text{ km s}^{-1}$. Our fiducial run reaches an H₂ outflow rate of $140 M_{\odot} \text{ yr}^{-1}$ after 1 Myr.

For comparison, observations of luminous quasar host galaxies find H₂ outflow rates up to ~ 100 – $1000 M_{\odot} \text{ yr}^{-1}$ (e.g. Ciccone et al. 2014; González-Alfonso et al. 2017). However, these observations infer H₂ masses from other molecules (e.g. CO or OH). As we saw above (and we will see in Section 5.1), the conversion factors between observed molecular lines and H₂ can introduce uncertainties of up to a factor ≈ 4 – 5 . Furthermore, observations typically estimate an outflow rate from the H₂ mass, size and velocity of the outflow as a whole, whereas we sum over individual particles in the simulations. In Section 5.1, we explore a more rigorous comparison between the observations and our simulations, using radiative transfer calculations of CO line emission performed on our simulations in post-processing to estimate H₂ outflow rates in the same way as the observations.

When the AGN luminosity is reduced by a factor of 10 (nH10_L45_Z1; magenta curve), the H₂ outflow rate is lower, reaching $29 M_{\odot} \text{ yr}^{-1}$ after 1 Myr. As we saw in Fig. 3, this is because less mass has been swept up from the ISM than in the fiducial run (the H₂ mass fraction in the low-luminosity run is roughly the same as in the reference run at 1 Myr).

The low-metallicity run (nH10_L46_Z0.1; green curve) produces an even lower H₂ outflow rate, reaching $6 M_{\odot} \text{ yr}^{-1}$ after 1 Myr. This reflects the lower H₂ mass in this run, primarily due to the lower dust abundance, as discussed above.

In Appendix A, we explore the effects of numerical resolution on the H₂ outflow rate in the simulations. We find that the molecular outflow rate is not well converged, as it increases by a factor of ≈ 4 as the mass resolution is increased from 240 to $10 M_{\odot}$ per gas particle (see Fig. A1). Therefore, our predicted molecular fractions can be interpreted as lower limits.

3.2 Physical properties of molecular gas

Fig. 5 shows the temperature-density distributions of gas particles in the high-resolution wedge after 1 Myr, in the four parameter variations runs (from left to right).

In the top row, the colour scale indicates the mass of gas in each temperature-density bin. In the fiducial run (nH10_L46_Z1; top left-hand panel), the injected wind particles have been shock-heated to

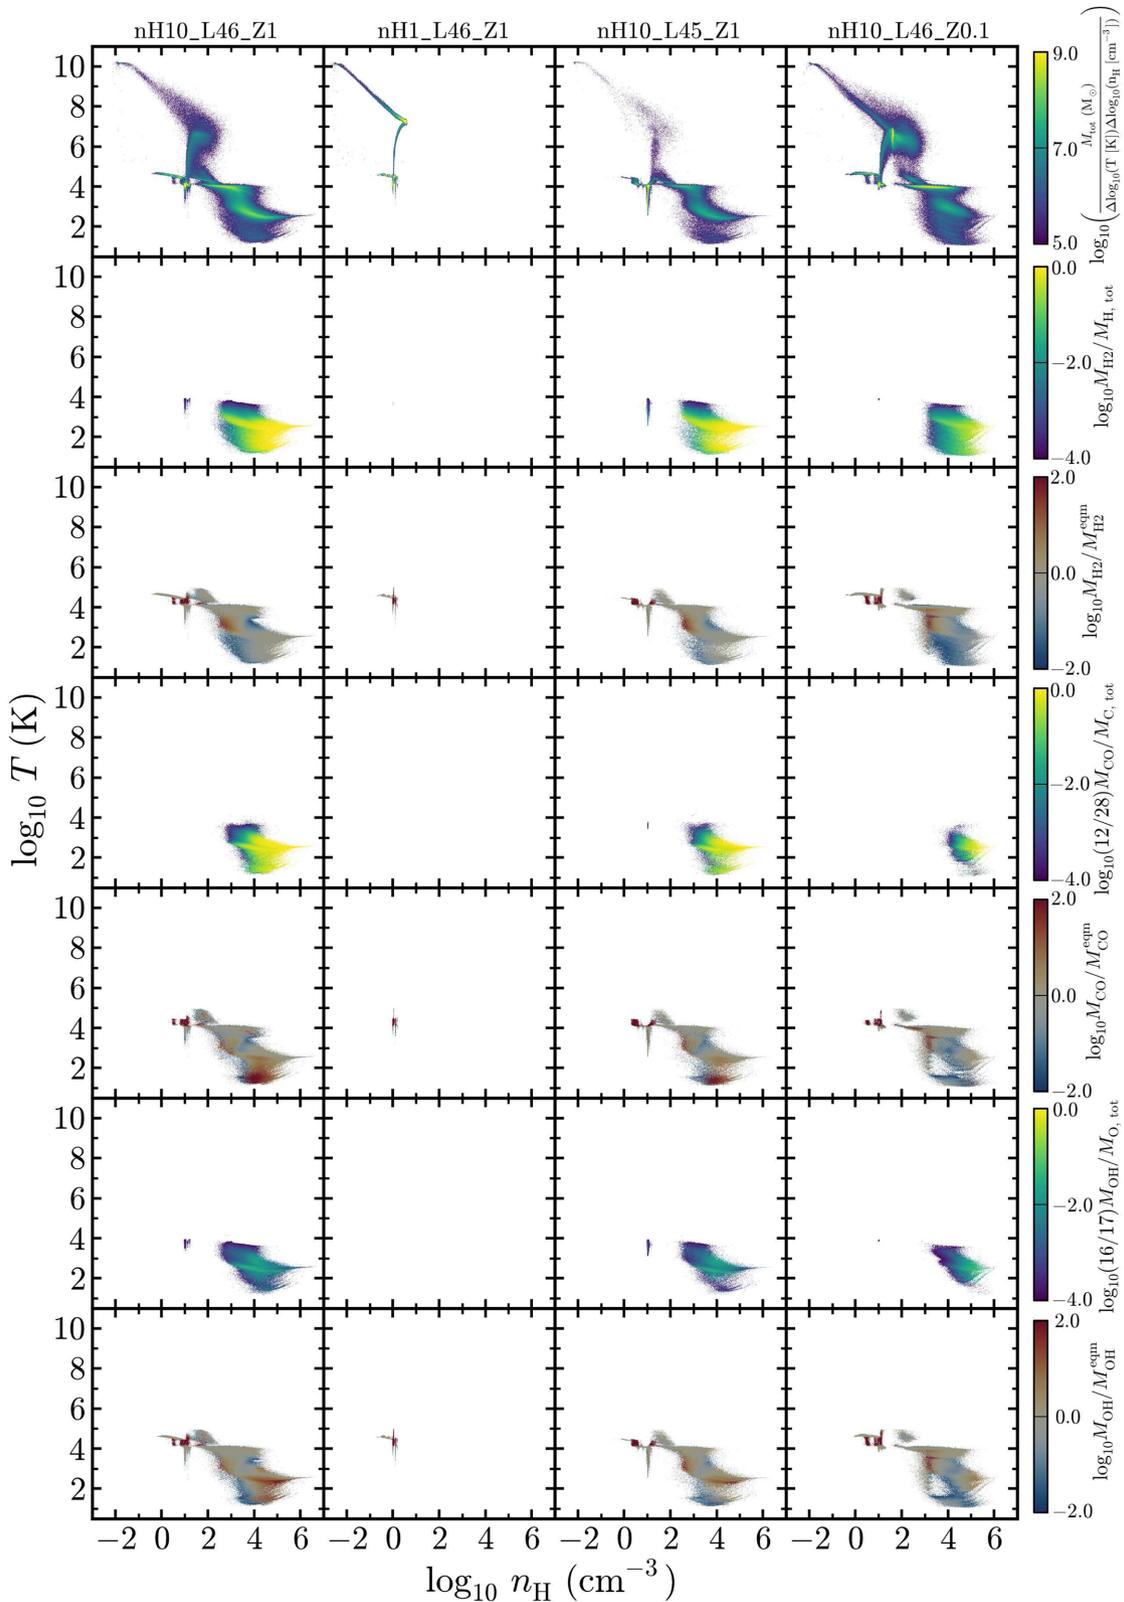


Figure 5. Temperature-density distributions of gas in the high-resolution wedge after 1 Myr in the four parameter variations runs (from left to right: nH10_L46_Z1, nH1_L46_Z1, nH10_L45_Z1, and nH10_L46_Z0.1). From top to bottom, the colour scale indicates (in each temperature-density bin): gas mass, H₂ fraction, ratio of H₂ masses using non-equilibrium and equilibrium abundances, CO fraction, ratio of non-equilibrium to equilibrium CO masses, OH fraction, and ratio of non-equilibrium to equilibrium OH masses. The factors 12/28 and 16/17 in the colour scales of the fourth and sixth rows, respectively, are the ratios of the atomic weights of C to CO, and of O to OH, which ensures that these fractions go to unity when all carbon is in CO, or when all oxygen is in OH.

$T \sim 10^{10}$ K, with a density $n_{\text{H}} \sim 10^{-2} \text{ cm}^{-3}$. The ambient ISM is initially at $n_{\text{H}} = 10 \text{ cm}^{-3}$, $T \sim 10^4$ K. The forward shock then heats this to $\sim 10^7$ K. The shock-heated ambient ISM is approximately in pressure equilibrium with the shocked wind material, albeit with a large scatter. Once the swept-up material is able to radiatively cool, its density increases. At intermediate densities ($\sim 10^2$ – 10^3 cm^{-3}), it cools to $\sim 10^4$ K, while at higher densities it cools below 10^3 K. At the highest densities reached in this simulation ($\sim 10^6 \text{ cm}^{-3}$), the thermal equilibrium temperature is ≈ 400 K.

In the low-density run (nH1_L46_Z1; second column), almost all of the swept-up material is still at the post-shock temperature and density of $\sim 10^7$ K and a few cm^{-3} , respectively, from passing through the forward shock. This material has not had sufficient time to cool by the end of the simulation after 1 Myr, so there is no cold, high-density gas.

The low-luminosity run (nH10_L45_Z1; third column) has a similar temperature-density structure as the fiducial run, although there is more swept-up material in the latter. We also saw in Fig. 1 that the outflow reaches a larger radius after 1 Myr in nH10_L46_Z1 than in nH10_L45_Z1, and the mass of swept-up gas at a given time is simply the mass of the ambient medium that was initially within the radius that the outflow has reached at that time.

The low-metallicity run (nH10_L46_Z0.1; fourth column) also has a similar temperature-density structure to the fiducial run, and with a similar mass of swept-up gas. However, less gas has cooled below 10^4 K in nH10_L46_Z0.1, which is likely at least partly due to the reduced molecular H_2 cooling (we will see in Section 5.3 that warm H_2 emits strongly in the infrared, and thus H_2 can be an important coolant). Metal line cooling is also lower in this run.

In the second, fourth, and sixth rows of Fig. 5, the colour scales indicate the H_2 , CO, and OH fractions, respectively, in each temperature-density bin. As above, these fractions have been normalized such that they are equal to 1 when all hydrogen is in H_2 , all carbon is in CO, or all oxygen is in OH. In the fiducial run, the H_2 fraction reaches close to unity at densities $n_{\text{H}} \gtrsim 10^3 \text{ cm}^{-3}$, even at temperatures up to $\sim 10^3$ K. In particular, comparing the first and second panels in the left-hand column, for the fiducial run, we see that there is a significant amount of warm H_2 (at a few hundred K), which is likely to exhibit bright emission from the near- and mid-infrared H_2 lines, as has been observed, for example, in shocked molecular gas in galaxy mergers and galaxy clusters (e.g. van der Werf et al. 1993; Sugai et al. 1997; Appleton et al. 2006, 2017; Egami et al. 2006; Zakamska 2010; Hill & Zakamska 2014). We will explore the IR-traced H_2 emission further in Section 5.3.

The CO fraction in the fiducial run approaches unity at a higher density ($n_{\text{H}} \gtrsim 10^4 \text{ cm}^{-3}$) than for the H_2 fraction. This is the same behaviour as in typical PDR models, as CO requires a higher density than H_2 to become shielded from dissociating radiation, and results in ‘CO-dark’ molecular gas at intermediate densities, which is rich in H_2 but not CO (e.g. Tielens & Hollenbach 1985; van Dishoeck & Black 1988; Wolfire, Hollenbach & McKee 2010; Smith et al. 2014). Similarly to H_2 , we also see significant amounts of CO at high temperatures, up to $\sim 10^3$ K. It would be interesting to look for warm CO in AGN-driven winds using high-J transitions of CO. Such high-J lines have been detected by *Herschel* in several ULIRGs, including the quasar Mrk 231 (e.g. van der Werf et al. 2010), although not specifically in the wind.

In regions of the temperature-density diagram where CO is close to unity, approximately half of the oxygen is in CO (for the metal abundance ratios that we assume here, which are taken from table 1 of Wiersma et al. 2009). The OH fraction in the fiducial run is

highest at densities $\gtrsim 10^4 \text{ cm}^{-3}$. However, the OH fraction never exceeds 0.3, and is typically much lower than this. In molecular gas with a CO fraction close to unity, most of the remaining oxygen is in either H_2O or O I, rather than in OH.

The low-density run shows no significant H_2 , CO, or OH, as the swept-up material has had insufficient time to cool, while in the low-luminosity run the H_2 , CO, and OH fractions show similar behaviour as in the fiducial run.

At low metallicity, the transition to high molecular fractions in H_2 , CO, and OH occurs at a higher density than in the fiducial model (a few times 10^4 cm^{-3} for H_2 , and $\sim 10^5 \text{ cm}^{-3}$ for CO and OH). This is partly due to the lower formation rate of H_2 on dust grains (which we assume scales linearly with metallicity), and partly due to reduced shielding by dust, self-shielding, and cross shielding of CO by H_2 . This trend of the atomic to molecular transition with metallicity has also been demonstrated in PDR models (e.g. McKee & Krumholz 2010; Krumholz 2013; Sternberg et al. 2014).

We see in Fig. 5 that our simulations (except the low-density run) produce dense ($n_{\text{H}} > 10^4 \text{ cm}^{-3}$) molecular gas in the outflow. Such dense gas has been observed, for example, from HCN in the wind of Mrk 231 (Aalto et al. 2015). While we do not track HCN in our chemical network, the bottom right-hand panels of Fig. 3 show that our simulations do form a significant mass of HCO^+ , which also traces dense gas.

Finally, the colour scales in the third, fifth, and seventh rows in Fig. 5 show, for H_2 , CO, and OH respectively, the ratio of the mass of the given species in each temperature-density bin calculated using the non-equilibrium abundances to the corresponding mass assuming chemical equilibrium, in the high-resolution wedge after 1 Myr. To compute the equilibrium abundances, we evolved the chemistry of all gas particles in the snapshot at 1 Myr at fixed temperature and fixed particle positions for a further 10 Gyr.

At $T \lesssim 300$ K and densities just below the transition to fully molecular, H_2 , CO, and OH are lower in non-equilibrium, by 1–2 orders of magnitude, i.e. the transition to fully molecular would occur at a slightly lower density if we assumed equilibrium abundances. Particles here have not had enough time to fully form molecules, because the formation time-scale of molecules increases with decreasing density. At higher temperatures of $\sim 10^3$ K at these densities, there is an enhancement in the non-equilibrium molecular fractions.

At $T \lesssim 100$ K and moderate densities ($10^3 \lesssim n_{\text{H}} \lesssim 10^4 \text{ cm}^{-3}$), the CO mass fraction is higher in non-equilibrium than in equilibrium, by ~ 2 orders of magnitude. We find that particles in this region are out of thermal equilibrium, and are evolving rapidly through the density-temperature space. This explains why the CO abundances of these particles can be so far out of chemical equilibrium, as the chemistry cannot keep up with the rapid evolution through temperature-density space.

The low-luminosity run shows similar trends in the non-equilibrium to equilibrium ratios as the fiducial run. However, the low-metallicity run shows a much stronger reduction of H_2 , CO, and OH abundances in non-equilibrium, compared to equilibrium, than we see in the fiducial run. This is unsurprising, as the formation rates of these molecules are lower at low metallicity, due to the lower dust abundance (for H_2) and the lower carbon and oxygen abundances (for CO and OH). It thus takes longer for them to reach equilibrium.

At solar metallicity, the total outflowing mass of H_2 is 24 per cent lower in non-equilibrium than in equilibrium, in both the fiducial and low-luminosity runs. The non-equilibrium effects are stronger

Table 3. Dominant chemical reactions that form CO, in order of decreasing total rate, calculated from simulation nH10_L46_Z1 after 1 Myr.

Reaction	Total CO formation rate ($M_{\odot} \text{ yr}^{-1}$)
$\text{HCO}^+ + \text{e}^- \rightarrow \text{CO} + \text{H}$	1.9×10^7
$\text{CO}^+ + \text{H} \rightarrow \text{CO} + \text{H}^+$	1.8×10^7
$\text{OH} + \text{C}^+ \rightarrow \text{CO} + \text{H}^+$	1.2×10^7
$\text{OH} + \text{C} \rightarrow \text{CO} + \text{H}$	8.0×10^6
$\text{O}_2 + \text{C} \rightarrow \text{CO} + \text{O}$	1.5×10^6

for CO and OH, where the fiducial and low-luminosity runs have 89 per cent and 65 per cent, respectively, more CO mass in non-equilibrium than equilibrium, and 96 per cent and 68 per cent, respectively, more OH mass. In the low-metallicity run, the non-equilibrium H_2 , CO, and OH masses are 78 per cent, 30 per cent, and 46 per cent lower, respectively, than in equilibrium.

3.3 Dominant CO formation channels

Following the non-equilibrium evolution of CO requires a complex chemical network, with several different formation channels that create CO. In the CHIMES model, there are 18 reactions that form CO, in addition to reactions that form intermediate species between atomic/ionized carbon and CO (see appendix B of Richings et al. 2014a for a full list of reactions included in the chemical network). To explore whether one formation channel dominates the CO chemistry, we calculate the total formation rate of CO from each of these 18 reactions in simulation nH10_L46_Z1 after 1 Myr. We include only particles with a CO fraction of at least 10 per cent, to exclude particles where both the formation and destruction rates are high that do not contribute significantly to the total CO mass.

We find that there is no single reaction that dominates the formation of CO. Table 3 lists the five most dominant reactions, in the order of decreasing rate. These reactions are also important in diffuse clouds in the Milky Way (see e.g. Fig. 4 in van Dishoeck & Black 1986). Nelson & Langer (1997) suggested a simplified CO model in which C^+ first forms CH_2^+ , which is converted rapidly to CH and CH_2 . These can then react with O to form CO, or they can be photodissociated. Nelson & Langer (1997) do not explicitly follow these intermediate species, instead treating the whole formation channel via CH and CH_2 together, according to the relevant rate coefficients (see their equations 18–21). However, we find that the reactions that form CO from CH or CH_2 contribute <0.1 per cent to the total formation rate in our reference simulation. Therefore, the simplified CO model of Nelson & Langer (1997) cannot be applied to the physical regimes that we study here.

4 MODEL VARIATIONS

In this section we consider the six model variations runs, which we compare to our standard model implementation. These runs all use the fiducial set-up, with an initial ambient ISM density of 10 cm^{-3} , an AGN luminosity of $10^{46} \text{ erg s}^{-1}$, and solar metallicity. We then vary different aspects of the modelling that are uncertain, to quantify how these uncertainties impact our results. These runs all used a resolution of $240 M_{\odot}$ per gas particle, which is a factor 8 lower mass resolution than was used in the main parameter variations runs in the previous section. We present a comparison of the different resolutions in Appendix A.

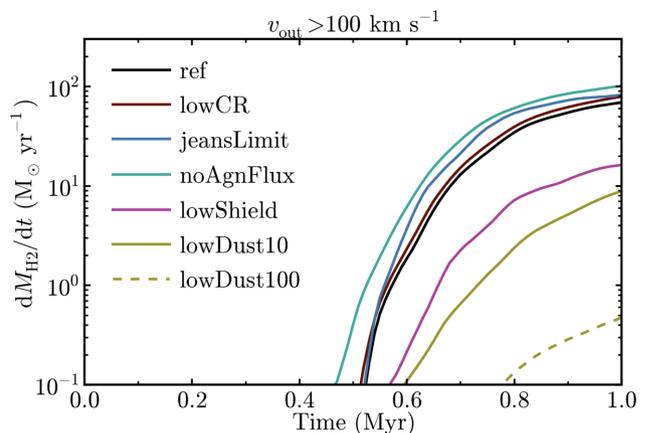


Figure 6. As Fig. 4, but for the six model variations runs (lowCR, red curve; jeansLimit, blue curve; noAGNflux, green curve; lowShield, magenta curve; lowDust10, yellow solid curve; lowDust100, yellow dashed curve), and the reference model (ref., i.e. nH10_L46_Z1_lowRes in Table 1; black curve). These runs are at a resolution of $240 M_{\odot}$ per gas particle. The largest effect on the H_2 outflow rate is seen when we reduce the shielding column density by a factor of 10 (lowShield), or we reduce the dust abundance by a factor of 10 (lowDust10) or 100 (lowDust100).

Fig. 6 shows the mass outflow rate of H_2 in gas outflowing with $v_{\text{out}} > 100 \text{ km s}^{-1}$ in the reference model (ref. i.e. nH10_L46_Z1_lowRes in Table 1; black curve) and in the six model variations runs. As in Fig. 4, the H_2 mass outflow rate is calculated in the high-resolution wedge and then multiplied by 17 to get the outflow rate in a full sphere, using equation 11.

The H_2 outflow rate in the reference model reaches $69 M_{\odot} \text{ yr}^{-1}$ after 1 Myr. This is a factor 2 lower than we found at higher resolution in run nH10_L46_Z1 (Fig. 4; see also Appendix A). In the lowCR run (red curve), we use a cosmic ray ionization rate of $\zeta_{\text{HI}} = 1.8 \times 10^{-16} \text{ s}^{-1}$, as measured in the Milky Way (Indriolo & McCall 2012) and a factor ≈ 30 lower than the fiducial value we use in the reference model for ULIRGs. This run produces a similar H_2 outflow rate as the reference model, reaching $79 M_{\odot} \text{ yr}^{-1}$ after 1 Myr. This suggests that uncertainties in the cosmic ray ionization rate do not have a strong impact on the H_2 outflow rate.

In the jeansLimit run (blue curve), we include a pressure floor such that the Jeans length is always resolved by at least 3.2 smoothing lengths. The reference model does not include this Jeans limiter, and so it might be susceptible to artificial fragmentation (e.g. Bate & Burkert 1997; Truelove et al. 1997). In Fig. 6, we see that the H_2 outflow rate in the jeansLimit run is also close to the reference model, reaching $82 M_{\odot} \text{ yr}^{-1}$ after 1 Myr. Therefore, if artificial fragmentation is occurring in the reference model, it does not significantly affect the H_2 outflow rate.

To test how the AGN radiation affects the formation and destruction of molecules in the AGN wind, we performed one run without the AGN radiation (noAGNflux; green curve). Unsurprisingly, the H_2 outflow rate is higher in this run, reaching $100 M_{\odot} \text{ yr}^{-1}$ after 1 Myr.

In the reference model, we calculate the shielding column density using a local approximation based on the density gradient (see Section 2.2). In Appendix B we show that, while this local approximation on average reproduces the true column density (computed via ray tracing in post-processing), there is a large scatter between the two of an order of magnitude. To test how these uncertainties might affect our results, we performed one run with column densities reduced by a factor of 10 (lowShield; magenta curve).

The H_2 outflow rate is noticeably lower in this run, reaching just $16 M_\odot \text{ yr}^{-1}$ after 1 Myr, a factor of 4.3 less than in the reference model. This highlights that uncertainties in the shielding column density can have a large impact on molecule formation, although the lowShield run is a worst-case scenario, as for most particles the local shielding approximation is closer to (or even less than) the true value (see Fig. B1).

In the reference model we assume that the dust abundance scales linearly with metallicity. However, it is not clear whether dust grains can survive the strong shocks and high gas temperatures found in AGN winds. To investigate how sensitive our results are to uncertainties in the dust abundance, we performed two runs at solar metallicity but with dust abundances reduced by a factor of 10 and 100 (lowDust10 and lowDust100, respectively; solid and dashed yellow curves in Fig. 6). We see that the dust abundance has the largest impact of all the model variations runs, reducing the H_2 outflow rate after 1 Myr by a factor of ≈ 8 and ≈ 150 when the dust abundance is lowered by a factor of 10 and 100, respectively. Therefore, to form a significant amount of H_2 in the AGN wind that can explain the observed fast molecular outflows, the swept-up material requires a dust-to-metals ratio that is close to the Milky Way value, or higher.

Ferrara & Scannapieco (2016) performed a series of hydrodynamic simulations of the swept-up ISM region of an AGN outflow. They followed the evolution of dust grains in their simulations, including destruction by sputtering in hot gas and accounting for the initial destruction of dust grains by the forward shock, both modelled based on the results of Dwek, Foster & Vancura (1996). Ferrara & Scannapieco (2016) found that dust grains are rapidly destroyed, within $\sim 10^4$ yr, in the swept-up ISM. We can estimate the sputtering time-scale in our fiducial model (nH10_L46_Z1) using equation 14 in Hirashita et al. (2015) (see also Nozawa, Kozasa & Habe 2006):

$$\tau_{\text{sput}} = 7.1 \times 10^5 \left(\frac{a}{1 \mu\text{m}} \right) \left(\frac{n_{\text{H}}}{1 \text{ cm}^{-3}} \right)^{-1} \text{ yr}, \quad (12)$$

where a is the grain radius. This equation is valid at gas temperatures $\gtrsim 10^6$ K, but at lower temperatures sputtering is inefficient. The post-shock density in the swept-up ISM is 4 times the ambient density, i.e. 40 cm^{-3} . So, for $a = 0.1 \mu\text{m}$ (as assumed by Ferrara & Scannapieco 2016), $\tau_{\text{sput}} = 1800 \text{ yr}$. In our fiducial model, it takes 0.5 Myr for the swept-up ISM to cool and start forming molecules (see Fig. 3). This suggests that dust grains in the swept-up ISM will be destroyed by sputtering before that gas can cool.

However, Ferrara & Scannapieco (2016) only included dust destruction, and they did not consider any dust formation mechanisms. For example, Hirashita & Nozawa (2017) highlighted the importance of dust growth via accretion of metals from the gas phase on to grains in their models of dust evolution in elliptical galaxies. We can estimate the accretion time-scale using equation 20 of Asano et al. (2013):

$$\tau_{\text{acc}} = 2.0 \times 10^7 \left(\frac{a}{0.1 \mu\text{m}} \right) \left(\frac{n_{\text{H}}}{100 \text{ cm}^{-3}} \right)^{-1} \times \left(\frac{T}{50 \text{ K}} \right)^{-1/2} \left(\frac{Z}{0.02} \right)^{-1} \text{ yr}, \quad (13)$$

where the metallicity $Z = Z_\odot = 0.0129$ in our fiducial model. From the top left-hand panel of Fig. 5, we see that the swept-up ISM covers a large range of densities and temperatures. As an example, consider $n_{\text{H}} = 10^4 \text{ cm}^{-3}$, $T = 1000 \text{ K}$. Then the accretion time-scale for $a = 0.1 \mu\text{m}$ grains is $\tau_{\text{acc}} = 0.07 \text{ Myr}$. Thus it is feasible

for dust grains to re-form after the swept-up ISM has cooled and before the end of the simulation at 1 Myr.

Dust growth from ISM accretion requires seed grains to be present. There are several possible sources of such seed dust in the post-shock cooling layers of AGN-driven galactic winds. For example, it is observed that star formation-driven galactic winds carry large dust masses (e.g. Hoopes et al. 2005; Roussel et al. 2010; Hutton et al. 2014; Meléndez et al. 2015; Hodges-Kluck, Cafmeyer & Bregman 2016). As the AGN wind interacts with star formation-driven outflows, some dust is likely to mix. Furthermore, while the main $\sim 1000 \text{ km s}^{-1}$ shock fronts in AGN-driven outflows can easily destroy dust, shocks driven in ambient overdensities will have lower velocities (Klein et al. 1994). Such overdensities can therefore protect dust from destruction by shocks. The dust protected by overdensities can subsequently mix with the cooling layers of AGN-driven outflows.

Additionally, Boschman et al. (2015) explored another formation channel for H_2 on polycyclic aromatic hydrocarbons (PAHs). They showed that this mechanism for H_2 formation can have an important impact on the structure of PDRs at high temperatures ($T > 200 \text{ K}$), which is particularly relevant for our AGN wind models, where we find a lot of dense gas at such temperatures. H_2 formation on PAHs is not included in our chemical model.

To conclusively determine whether sufficient dust can survive or re-form to produce the observed fast molecular outflows, we will need to model the evolution of dust grains, including destruction from shocks and sputtering and formation from accretion. The resulting dust abundances, and the resulting grain size distributions (as both τ_{sput} and τ_{acc} depend on the grain radius), along with a treatment for the PAH chemistry, will then need to be coupled to H_2 formation and shielding in the chemical model.

5 OBSERVABLE MOLECULAR LINES

Our chemical model includes several molecular species that are commonly seen in observations, such as CO and OH. We can therefore use our simulations to make predictions for observable emission and absorption lines from these molecules, which we can compare to observations. We model the molecular lines using radiative transfer calculations performed on snapshots from our simulations in post-processing, using the publicly available Monte Carlo radiative transfer code RADMC-3D,⁷ version 0.40 (Dullemond et al. 2012). RADMC-3D follows the line radiative transfer, along with thermal dust emission and absorption. We include anisotropic scattering by dust grains, and we include the radiation from the AGN, using the average quasar spectrum from Sazonov et al. (2004).

RADMC-3D includes several approximations to calculate the level populations of a given molecular species. The simplest approach is to assume Local Thermodynamic Equilibrium (LTE). However, Duarte-Cabral et al. (2015) showed that assuming LTE can overestimate CO emission in their hydrodynamic simulations of spiral galaxies. We therefore use an approximate non-LTE treatment using the Local Velocity Gradient (LVG) method (Sobolev 1957). See Shetty et al. (2011) for a full description of the LVG method implemented in RADMC-3D.

RADMC-3D operates on a grid, rather than on gas particles, so we first need to project the gas particles from the simulation snapshot on to a 3D grid. We construct an Adaptive Mesh Refinement (AMR) grid, with a base grid of 16^3 cells over the high-resolution octant.

⁷ <http://www.ita.uni-heidelberg.de/~dullemond/software/radmc-3d/>

We then refine any cells that contain more than eight particles by splitting the cell into eight smaller cells, and we continue refining the grid in this way until all cells contain no more than eight particles. This requires 8–10 levels of refinement, which results in a minimum cell size of 0.07–0.3 pc, depending on the simulation.

Once we have constructed the AMR grid, we then interpolate the particle densities, temperatures, and velocities on to the grid. We do this using a cubic spline kernel with a smoothing length that encloses 32 neighbouring particles, as used in the simulations. We include two species of dust grains, graphite and silicate, using opacities from Martin & Whittet (1990) (based on the grain size distribution of Mathis, Rumpl & Nordsieck 1977). We use dust-to-gas ratios of $2.4 \times 10^{-3} Z/Z_{\odot}$ and $4.0 \times 10^{-3} Z/Z_{\odot}$ for graphite and silicate, respectively, which we assume scale linearly with metallicity, Z . At solar metallicity, these are taken from CLOUDY version 13.01 (Ferland et al. 2013), using the ‘ISM’ grain abundances. In cells with a gas temperature $T > 10^5$ K, we set the dust abundance to zero. This temperature is chosen for consistency with the chemical model, as we switch off the molecular chemistry and dust processes above this temperature. At high gas temperatures, we expect dust grains to be rapidly destroyed by sputtering.

We also include a microturbulent Doppler broadening parameter $b_{\text{turb}} = 7.1 \text{ km s}^{-1}$ (in addition to thermal broadening), for consistency with the H_2 self-shielding function used in the simulations. However, we will see that the resolved gas motions produce lines that are much broader than this.

There is a possible inconsistency between the chemistry, which is calculated on the particles, and the radiative transfer, which is calculated on the grid cells. For example, suppose that there are two particles that contribute equally to a given cell. One of these particles is cold (100 K) with a CO fraction of 1, while the other is hot (10^7 K) with a CO fraction of 0. If we weight both particles equally, the cell will have a temperature $\approx 5 \times 10^6$ K and a CO fraction of 0.5, i.e. half the CO will be unrealistically hot. However, from the point of view of the chemistry solver, all of the CO was cold. To avoid such inconsistencies, when interpolating particle temperatures and velocities on to the grid, we weight each particle by its mass of the molecular species (CO in the above example), rather than the total particle mass.

Finally, we mirror the high-resolution octant in the line-of-sight direction, so that we include both the approaching (blueshifted) and receding (redshifted) sides of the outflow. We then run RADMC-3D to produce maps of CO (Section 5.1), OH (Section 5.2), and IR-traced H_2 (Section 5.3) emission and absorption. These maps contain 1024×1024 spatial pixels covering the high-resolution octant, and 2000 equally spaced velocity bins extending from -1000 to $+1000 \text{ km s}^{-1}$. For each line, we also create a map of the continuum emission only, which we then subtract from the full-line emission map.

5.1 CO emission

We calculate the emission from the three lowest rotational transitions of CO, $J = 1-0$, $2-1$, and $3-2$, with rest frame wavelengths of 2.6, 1.3 and 0.9 mm, respectively. These lines have been observed at high velocities (up to $\sim 1000 \text{ km s}^{-1}$) in several AGN host galaxies (e.g. Feruglio et al. 2010, 2015; Cicone et al. 2012, 2014). We use molecular data for CO from the LAMDA data base⁸ (Schöier et al. 2005). To calculate the non-LTE level populations of CO, we

include collisions with ortho- and para- H_2 , using collisional rate coefficients from Yang et al. (2010) and assuming an ortho-to-para ratio of 3:1. The collisional rate coefficients are only tabulated up to 3000 K, so we assume that they are constant above this temperature. We will primarily focus on the (1–0) line, but we will use the higher transitions to explore the excitation of CO.

The top row of Fig. 7 shows velocity-integrated maps of the continuum-subtracted CO (1–0) line emission in nH10_L46_Z1 (left), nH10_L45_Z1 (centre), and nH10_L46_Z0.1 (right), from the final snapshot after 1 Myr, while the bottom row shows the continuum-subtracted CO(1–0) spectra from these three runs. These spectra only show the flux from the high-resolution octant (which has also been mirrored in the line-of-sight direction, i.e. to create one quadrant), in other words, they have not been multiplied by 4 to get the spectrum from the full spherical outflow. The fluxes have been normalized to a distance of 184 Mpc, which corresponds to the distance of Mrk 231. We do not show the low-density run (nH1_L46_Z1), as it did not form any molecules. Note that the colour scale of the velocity-integrated map and the y-axis scale of the spectrum for the low-metallicity run (right-hand column) is lower than that for the other two runs.

In the velocity-integrated maps, we see that the CO emission is not smooth over the shell of swept-up ISM, but rather is clumpy. This clumpiness was also seen in the 2D and 3D views of the gas density in Figs 1 and 2, respectively, and is triggered by the radiative cooling in the swept-up material, as we see no clumpiness in the low-density run in Fig. 1, which does not cool. In the low-metallicity run (right-hand column of Fig. 7), the CO emission appears to be concentrated into much smaller clumps than in the other two runs. This is likely due to the lower dust abundance in this run, which results in weaker dust shielding. Since CO needs to be shielded from UV, it will therefore only trace the centres of dense clumps when the dust shielding is weaker.

All three spectra in the bottom row of Fig. 7 show a significant dip at the systemic velocity ($\Delta v = 0$). However, this does not appear to be physical, and is likely due to the artefacts along the boundary between high- and low-resolution regions that we saw in Fig. 1. Since the CO emission is in a spherical, expanding shell, emission at $\Delta v = 0$ will come from the part of the shell moving perpendicular to the line-of-sight direction, i.e. along the boundary with the low-resolution region. However, we saw in Fig. 1 that the spherical shell is disrupted near these boundaries. This is most evident in the top left-hand and top right-hand panels of Fig. 7, where we see that the CO emission does not extend all the way to the left-hand and bottom edges of the maps. By symmetry, CO emission will also be missing along the third high-resolution–low-resolution boundary (perpendicular to the line of sight). This would then explain the dip in the CO spectra. We will therefore need to be careful that these artefacts do not affect our analysis below.

In the two runs with an AGN luminosity of $10^{46} \text{ erg s}^{-1}$, the CO spectra extend to $\pm 500 \text{ km s}^{-1}$, while the low-luminosity run has a narrower CO spectrum, extending to $\pm 250 \text{ km s}^{-1}$. Run nH10_L46_Z1 has the brightest CO flux, followed by nH10_L45_Z1 and then nH10_L46_Z0.1. This follows the trend in outflowing CO mass between these runs that we saw in Fig. 3. Note that, unlike typical observations of AGN host galaxies, all of the CO emission in our simulations is from the outflow, and there is no pre-existing molecular disc in the host galaxy, which would appear as an additional narrow component in the CO spectrum (e.g. Feruglio et al. 2010; Cicone et al. 2014).

We can now use the CO spectra from our simulations to estimate the H_2 outflow rate in the same way as in the observations.

⁸ <http://home.strw.leidenuniv.nl/moldata/>

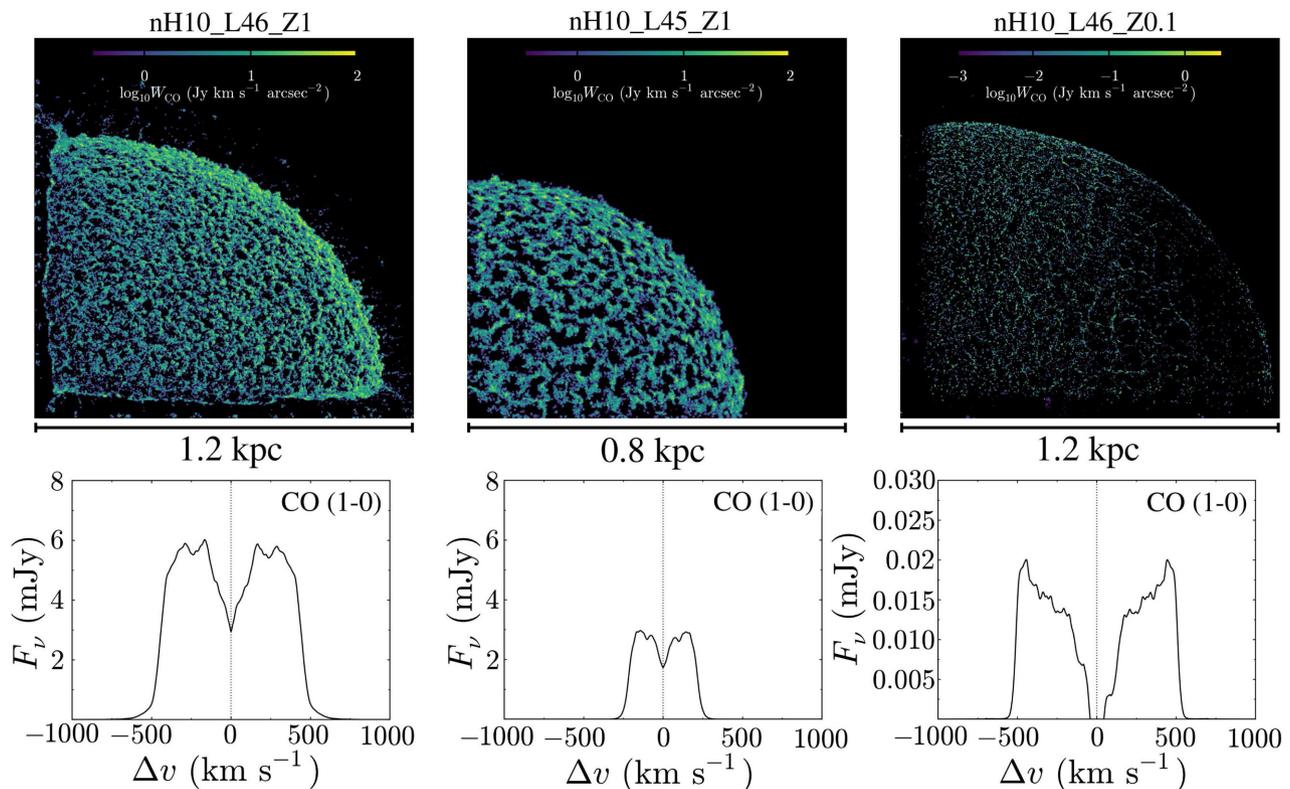


Figure 7. Velocity-integrated maps of CO (1–0) emission (top row) and CO (1–0) spectra (bottom row) from runs nH10_L46_Z1 (left), nH10_L45_Z1 (centre), and nH10_L46_Z0.1 (right) after 1 Myr. These were calculated in the high-resolution octant, which was mirrored in the line-of-sight direction to include both the receding and approaching sides of the outflow. The spectra have been normalized to a distance of 184 Mpc, which corresponds to the distance of Mrk 231. Note that the dips in the spectra at $\Delta v = 0$ are the result of artefacts along the boundary between the high- and low-resolution regions, and are not physical. These do not affect our comparisons with observed CO luminosities, for which we only integrate over the broad wings of the line, and also likely do not significantly affect our estimates of the CO to H_2 conversion factors (see the text).

Cicone et al. (2014) presented CO (1–0) observations of seven ULIRGs and quasar hosts, of which four showed detections of fast molecular outflows. They also compared these with a further 12 sources from the literature. Not all of these galaxies are luminous AGN, so we only compare our simulations to the Seyfert 1 and 2 galaxies in the extended sample of Cicone et al. (2014). This gives us a sample of 10 AGN host galaxies, of which 6 were observed by Cicone et al. (2014) and a further 4 were taken from the literature (Wiklund, Combes & Henkel 1995; Maiolino et al. 1997; Cicone et al. 2012; Feruglio et al. 2013a,b).

Cicone et al. (2014) measured the CO luminosity in the broad wings of the continuum-subtracted CO line. The velocity range that they integrated over depends on the source, but was typically from $\approx 300 \text{ km s}^{-1}$ to the maximum velocity that the wing could be detected at (see their table 2). To match the method used for the observations, we integrate our simulated CO spectra in each wing from 300 km s^{-1} to a maximum v_{max} , which we take to be the velocity at which the flux density falls to 1 per cent of its maximum value. This definition might not be equivalent to the v_{max} used in the observations, as we do not model the noise or detectability of CO in the simulated spectra, but it does not have a significant impact on the CO luminosity. For the low-luminosity run, we use a lower minimum velocity of 200 km s^{-1} , as $v_{\text{max}} < 300 \text{ km s}^{-1}$. By restricting the CO luminosity to the wings of the line, we also avoid the artificial dip at $\Delta v = 0$.

To obtain the H_2 mass from the CO luminosity, Cicone et al. (2014) assumed a conversion factor $\alpha_{\text{CO}(1-0)} = M_{H_2}/L_{\text{CO}(1-0)} =$

$0.8 M_{\odot} (\text{K km s}^{-1} \text{ pc}^2)^{-1}$. This value, which is ≈ 5 times lower than the Milky Way value, is commonly used for ULIRGs. Cicone et al. (2012) showed that the molecular gas in Mrk 231 has similar excitation in the host galaxy and the outflow, and so the outflow plausibly has the same conversion factor as its ULIRG host. To compare with observations, we use the same conversion factor, although we will investigate this assumption further below.

Cicone et al. (2014) then calculated the molecular mass outflow rate by assuming that the outflowing material is uniformly distributed over a radius R , measured from the CO maps. This gives an outflow rate of $\dot{M}_{H_2} = 3vM_{H_2}/R$, where the velocity v is the average velocity in the range used to integrate the wings of the CO line. In our simulations, the outflowing molecular gas is instead found in a shell, for which the outflow rate is $\dot{M}_{H_2} = vM_{H_2}/R$. We therefore divide the rates given in Cicone et al. (2014) by 3, for consistency with the simulations, although we note that this may underestimate the true outflow rate if the geometry of the observed systems is that of a uniform outflow rather than a shell.

The top panel of Fig. 8 shows the molecular outflow rates derived from our simulated CO spectra, from run nH10_L46_Z1 (red circle), nH10_L45_Z1 (blue circle), and nH10_L46_Z0.1 (green circle), plotted versus AGN luminosity. The outflow rates from the high-resolution quadrant have been multiplied by 4 to give the outflow rate in the full spherical outflow. The black symbols show the observed sources from Cicone et al. (2014) and references therein. The molecular outflow rates in the two simulations at solar metallicity are slightly lower than the observations, although they are still

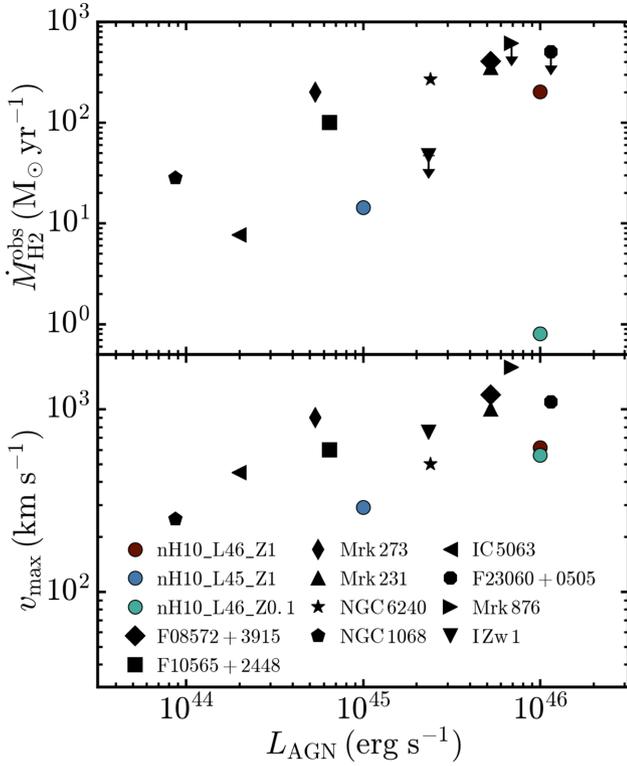


Figure 8. H_2 outflow rate (top panel) and maximum line-of-sight velocity (bottom panel), estimated from CO (1–0) spectra, plotted versus AGN luminosity. Coloured circles show the simulations: nH10_L46_Z1 (red), nH10_L45_Z1 (blue), and nH10_L46_Z0.1 (green). Black symbols show observed AGN host galaxies: *IRAS* F08572+3915, *IRAS* F10565+2448, *IRAS* F23060+0505, Mrk 273, Mrk 876, I Zw 1 (Cicone et al. 2014), IC 5063 (Wiklind et al. 1995), NGC 1068 (Maiolino et al. 1997), Mrk 231 (Cicone et al. 2012), and NGC 6240 (Feruglio et al. 2013a,b). The H_2 outflow rates estimated from the simulated CO spectra are slightly lower than the observations, although they are still reasonably close, while the simulated line of sight velocities are also lower than in the observations, by a factor ≈ 2 .

reasonably close, given the uncertainties in the geometry. The low-metallicity run is significantly below the observations, although this is unsurprising as these host galaxies are typically ULIRGs, with metallicities close to, or greater than, solar (e.g. Rupke, Veilleux & Baker 2008; Kilerci Eser, Goto & Doi 2014). The molecular outflow rate increases with increasing AGN luminosity, as found by Cicone et al. (2014).

In the bottom panel of Fig. 8, we show the maximum line-of-sight velocity of the CO wings versus AGN luminosity. In the simulations, this is defined as the velocity where the CO flux density falls to 1 per cent of its maximum value. These are generally lower than is seen in the observations, by a factor ≈ 2 . The definition of v_{\max} in the simulations is not exactly equivalent to the observational definition, as we do not model noise or the detectability of CO in the simulations, but we can see that the simulated spectra in Fig. 7 do not reach velocities up to $\sim 1000 \text{ km s}^{-1}$ as in the observations. This may be due to the simplified geometry and the assumption of a uniform ambient ISM in the simulations. For example, if there are inhomogeneities in the ISM, there will be a range of outflow velocities, and the maximum velocity of the outflow will be determined by the lowest density channels, through which the outflow can escape more quickly. The velocity of the outflow increases with increasing AGN luminosity in the simulations and the observations.

Table 4. Conversion factors from CO luminosity of the (1–0), (2–1), and (3–2) lines to H_2 mass after 1 Myr in the three parameter variations runs that form molecules. For comparison, observations typically assume $\alpha_{\text{CO}(1-0)} = 0.8 M_{\odot} (\text{K km s}^{-1} \text{pc}^2)^{-1}$ when inferring H_2 masses from CO (1–0) emission.

Simulation	$\alpha_{\text{CO}} = M_{H_2}/L_{\text{CO}}^a$		
	(1–0)	(2–1)	(3–2)
nH10_L46_Z1	0.13	0.08	0.06
nH10_L45_Z1	0.15	0.09	0.07
nH10_L46_Z0.1	1.77	0.82	0.80

Note. ^aUnits of α_{CO} are $M_{\odot} (\text{K km s}^{-1} \text{pc}^2)^{-1}$.

For the comparisons in Fig. 8, we assumed a conversion factor between CO and H_2 of $\alpha_{\text{CO}(1-0)} = 0.8 M_{\odot} (\text{K km s}^{-1} \text{pc}^2)^{-1}$, as used in the observations. However, we can also use our simulations to calculate this conversion factor, as well as conversion factors for the (2–1) and (3–2) lines. We first calculate the total CO luminosity in the high-resolution quadrant by integrating over the full CO line (not just the wings). We then calculate the total H_2 mass in the high-resolution quadrant (i.e. the high-resolution octant and mirrored in the line-of-sight direction). Since we integrate over the full CO line, the total CO luminosity will be affected by the artificial dip at $\Delta v = 0$. We therefore use the full high-resolution region to calculate the H_2 mass, rather than using only the high-resolution wedge and then scaling up. This way, the H_2 mass will also be affected by the artificial reduction in molecular gas along the boundaries with the low-resolution region. We find that, in the fiducial run, the masses of H_2 and CO in the high-resolution octant are 25 per cent lower than using the high-resolution wedge and scaling up (due to these artefacts), but the ratio of H_2 to CO mass changes by only 2 per cent. We therefore expect that the ratio of H_2 mass to CO luminosity (i.e. α_{CO}) calculated in this way should be unaffected by these artefacts.

Table 4 shows the α_{CO} conversion factors in our three simulations after 1 Myr, for the lines (1–0), (2–1), and (3–2). The value of α_{CO} depends on the time at which we calculate it, but we quote values at 1 Myr here as this corresponds to typical flow times, r/v_{out} , of observed AGN-driven molecular outflows (e.g. González-Alfonso et al. 2017). In the fiducial and low-luminosity runs, at solar metallicity, we find $\alpha_{\text{CO}(1-0)} = 0.13$ and $0.15 M_{\odot} (\text{K km s}^{-1} \text{pc}^2)^{-1}$, respectively, almost independent of the AGN luminosity. This is a factor of 5–6 lower than is typically assumed in observations (e.g. Cicone et al. 2014), which would result in a factor of 5–6 lower H_2 outflow rates. For example, the H_2 outflow rates from Cicone et al. (2014) would then all be below $120 M_{\odot} \text{yr}^{-1}$.

At $0.1 Z_{\odot}$, $\alpha_{\text{CO}(1-0)}$ is more than a factor of 10 higher than in the other two runs. This is unsurprising, as there is less carbon and oxygen relative to hydrogen, and is also reflected in the lower CO abundances at low metallicity (see Fig. 3). However, we caution that the α_{CO} factors in the low-metallicity run are poorly converged with numerical resolution (see Appendix A). The metallicity dependence of the CO to H_2 conversion factor has also been noted in previous studies of interstellar gas across a range of galaxy properties (e.g. Israel 1997; Leroy et al. 2011; Feldmann, Gnedin & Kravtsov 2012; Genzel et al. 2012; Bolatto, Wolfire & Leroy 2013a). However, as the quasar host galaxies with observed fast molecular outflows are typically ULIRGs, with metallicities close to, or greater than, solar (e.g. Rupke et al. 2008; Kilerci Eser et al. 2014), our solar metallicity estimate for $\alpha_{\text{CO}(1-0)}$ is more relevant for these systems.

The conversion factors for the (2–1) and (3–2) CO lines are a factor ≈ 2 lower than for the (1–0) line in all three simulations. For comparison, Cicone et al. (2014) assume that the CO lines are thermalized and optically thick, such that the CO line luminosities, and hence the α_{CO} conversion factors, of different rotational transitions are equal.

In typical molecular clouds, the CO 1–0 line is optically thick. However, this line can become optically thin, for example, in a highly turbulent medium. In the optically thin regime, the α_{CO} conversion factor reaches a lower limit. Bolatto et al. (2013a) showed that, for a CO to H₂ abundance $Z_{\text{CO}} = n_{\text{CO}}/n_{\text{H}_2} = 10^{-4}$ and a CO excitation temperature 30 K, the optically thin limit is $\alpha_{\text{CO}} = 0.34 M_{\odot} (\text{K km s}^{-1} \text{pc}^2)^{-1}$. While this appears to be inconsistent with the values that we find from our simulations, which are less than this at solar metallicity, we note that, for the solar abundances used in our simulations, $Z_{\text{CO}} = 5 \times 10^{-4}$ when all carbon is in CO and all hydrogen is in H₂. This reduces the optically thin limit of α_{CO} by a factor of 5, and so the values from the simulations are no longer inconsistent with this limit. Moreover, in Fig. 5 we saw that CO is found across a wide range of temperatures (~ 10 – 10^3 K) in our simulations, while the above optically thin limit assumes a single excitation temperature of 30 K and LTE.

Dasyra et al. (2016) measured CO (4–3)/(2–1) flux ratios of 5–11 in the jet-driven winds in the galaxy IC 5063, which indicates that the CO gas in this outflow is highly excited and optically thin. Based on these flux ratios, they inferred optically thin α_{CO} conversion factors of 0.27 – $0.44 M_{\odot} (\text{K km s}^{-1} \text{pc}^2)^{-1}$, which are reasonably close (within a factor ≈ 2 – 3) to our simulations at solar metallicity.

Weiß et al. (2001) demonstrated that the CO to H₂ conversion factor measured in different regions of the starburst galaxy M82 depends on the excitation of the molecular gas, with the conversion factor scaling with kinetic temperature, T_{kin} , and H₂ density, n_{H_2} , as $T_{\text{kin}}^{-1} n_{\text{H}_2}^{1/2}$. This scaling is also predicted by theoretical models of virialized clouds (e.g. Maloney & Black 1988), although it is not clear whether the molecular clouds in an AGN wind will also be virialized. If they are not virialized, this will affect their line widths, which will affect the CO (1–0) luminosity (since this is typically an optically thick line), and hence the α_{CO} factor, for a given T_{kin} . Cicone et al. (2012) showed that the (2–1)/(1–0) CO line ratio in the core and the broad components in Mrk 231 are similar, and thus they suggested that the same conversion factor should be used for the outflow and the host galaxy. It is therefore instructive to compare the excitation of molecular gas in the simulations to these observations.

In Fig. 9 we show the CO line ratios for the (2–1) and (3–2) transitions, relative to the (1–0) line, from our three simulations after 1 Myr (solid, coloured curves). We compare these to observed line ratios (black and grey curves) from Mrk 231 (Cicone et al. 2012; Feruglio et al. 2015), the starburst galaxy M82 (Weiß et al. 2005), the S0 galaxy NGC1266 with an AGN-driven molecular outflow (Glenn et al. 2015), and the inner disc of the Milky Way, with galactic longitude $2.5^\circ < |l| < 32.5^\circ$ (Fixsen et al. 1999). These line ratios are calculated with fluxes in units of Jy km s^{-1} , for consistency with Fig. 5 of Cicone et al. (2012). Note that the value of the line ratios will depend on the units used, because converting between commonly used units (e.g. Jy km s^{-1} , W m^{-2} , or K km s^{-1}) will introduce factors of the rest wavelength, which will be different for different lines.

The molecular gas in our simulations shows significantly higher excitation than the core or broad components of Mrk 231 measured by Cicone et al. (2012) and Feruglio et al. (2015). It is possible that

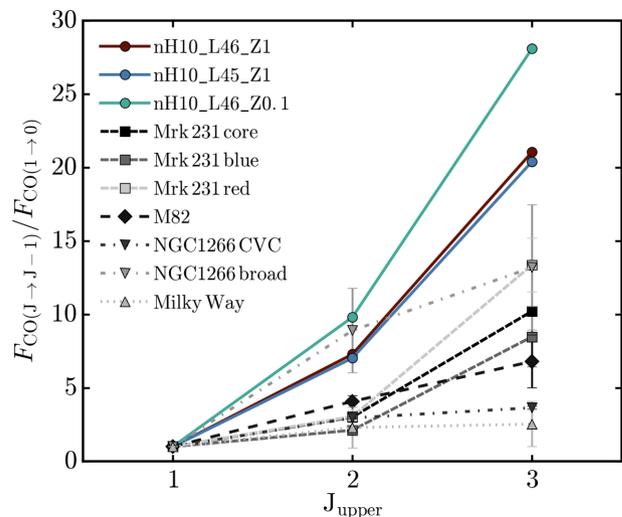


Figure 9. CO line ratios, relative to the (1–0) line, in the simulations nH10_L46_Z1 (red curve), nH10_L45_Z1 (blue curve), and nH10_L46_Z0.1 (green curve). Black and grey curves show observations of Mrk 231 (from the core and the red- and blueshifted components; Cicone et al. 2012; Feruglio et al. 2015), the starburst galaxy M82 (Weiß, Walter & Scoville 2005), the S0 galaxy NGC 1266 with an AGN-driven outflow (from the central velocity component and the broad component; Glenn et al. 2015), and the inner disc region of the Milky Way, with galactic longitude $2.5^\circ < |l| < 32.5^\circ$ (Fixsen, Bennett & Mather 1999). The line ratios use fluxes in units of Jy km s^{-1} . The simulations generally show higher CO excitation than the observed systems, although the broad component of NGC 1266 has a similar (2–1)/(1–0) ratio.

the unusually high excitation in our simulations is at least partly due to the lack of inhomogeneities in the ambient ISM, which would lead to a range of shock velocities throughout the shell of swept-up gas that would affect the shock heating, and hence excitation, of the molecular gas in the outflow. The higher excitation in the simulations suggests that the values of $\alpha_{\text{CO}(1-0)}$ that we calculate in Table 4 may be lower than the true values in Mrk 231. We must however note that the comparison with Mrk 231 is for a single quasar and it is not clear whether the CO excitation in its wind is representative of AGN-driven molecular outflows. Furthermore, the broad component in NGC 1266, with a width of 274 – 597 km s^{-1} (depending on the line), has a comparable (2–1)/(1–0) ratio to the simulations, although the (3–2)/(1–0) ratio is lower. This broad component has been attributed to an outflow powered by a buried AGN, as the star formation in this system is too weak to drive the observed outflow (Alatalo et al. 2011; Glenn et al. 2015). Note that NGC 1266 (a buried AGN in an S0 galaxy) is a very different system to Mrk 231 (a luminous quasar in a ULIRG).

5.2 OH absorption

Fast molecular outflows have been detected in several OH lines (see fig. 1 of González-Alfonso et al. 2014 for an energy-level diagram of OH). We use RADMC-3D to calculate the spectra of two commonly observed transitions to the ground state: the $^2\Pi_{3/2} J = 5/2 - 3/2$ doublet at $119 \mu\text{m}$, and the $^2\Pi_{1/2} - ^2\Pi_{3/2} J = 1/2 - 3/2$ doublet at $79 \mu\text{m}$. Each of these transitions consists of a doublet, separated by 520 km s^{-1} and 240 km s^{-1} , respectively. The molecular data for OH are taken from the LAMDA data base. We compute the non-LTE level populations of OH including collisions from ortho- and para-H₂, using collisional rate coefficients from Offer, van Hemert & van Dishoeck (1994), assuming an ortho-to-para ratio of 3:1. As these

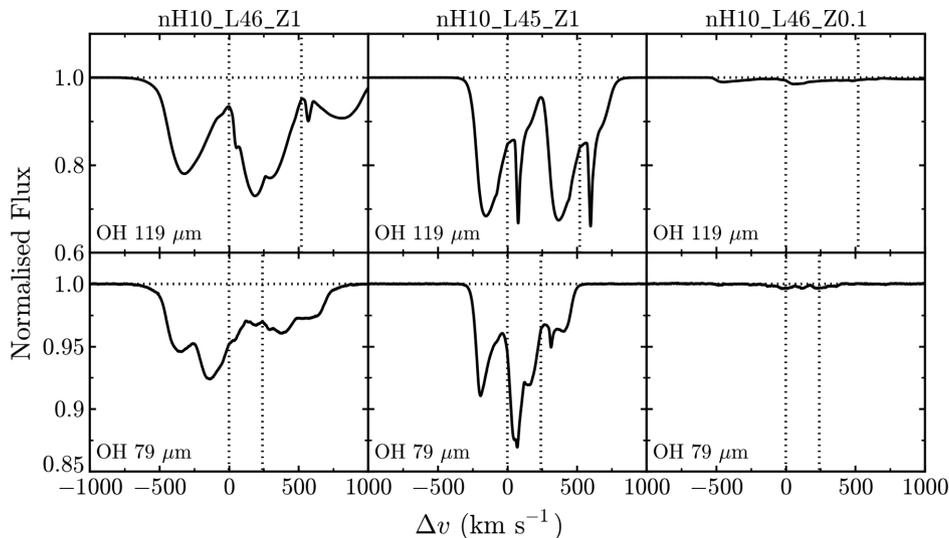


Figure 10. Continuum-normalized spectra of the OH doublets at 119 μm (top row) and 79 μm (bottom row) from the simulations nH10_L46_Z1 (left), nH10_L45_Z1 (centre), and nH10_L46_Z0.1 (right), after 1 Myr. The velocity scale is relative to the rest wavelength of the blueward component of the doublet. The positions of the two doublet components are indicated by vertical dotted lines.

rate coefficients are only tabulated up to 300 K in the LAMDA data base, we assume that they are constant above this temperature. We caution that, for OH, excitation by the far-infrared radiation field can be important. However, the implementation of the non-LTE LVG method in RADMC-3D does not currently account for the radiation from the dust continuum or from the AGN when computing level populations. Nevertheless, as we discuss below, the OH 119 μm line is typically optically thick, in which case the absorption strength of the 119 μm line should depend primarily on the covering fraction of OH in the outflow, rather than the details of the OH level populations.

Fig. 10 shows the continuum-subtracted spectra of the OH doublets at 119 μm (top row) and 79 μm (bottom row) in nH10_L46_Z1 (left), nH10_L45_Z1 (centre), and nH10_L46_Z0.1 (right), after 1 Myr. The velocity scale is relative to the rest wavelength of the blueward component of each doublet. The rest positions of the two components are shown by the vertical dotted lines. We calculate the normalized spectra from spatial pixels within the radius of the outflow, as measured from the CO emission maps in Fig. 7.

The spectra from nH10_L46_Z1 and nH10_L45_Z1 show several distinct absorption features. These are most clearly illustrated in the OH 119 μm spectrum of nH10_L45_Z1 (top centre panel of Fig. 10). First, there is broad blueshifted absorption seen in both components of the doublet. This arises from outflowing material coming towards the observer. Secondly, there is a narrow, redshifted absorption feature, which is due to OH in the ambient ISM. This is redshifted because the ambient ISM is inflowing (due to the gravitational potential of the black hole and the host galaxy), so the ambient ISM that lies in front of the outflow is moving away from the observer. Note that the gas inflows in the simulations are different from a realistic galaxy due to our idealized set-up (e.g., we do not include rotation of the gas). Finally, there is a broad redshifted absorption feature (seen more clearly in the redward doublet component in the top left-hand panel) due to outflowing material on the far side of the outflow. This is weaker than the broad blueshifted absorption because the receding side of the outflow lies behind most of the continuum-emitting dust, as viewed by the observer.

In the 79 μm spectra the two components of the doublet are closer together, so different absorption features overlap more. The absorption is weaker at 79 μm than at 119 μm , as seen in observations (e.g. González-Alfonso et al. 2017). The low-metallicity run, nH10_L46_Z0.1, shows only negligible OH absorption, as we will discuss further below.

It is likely that the artefacts along the boundaries between the high- and low-resolution regions will affect the OH spectra, similarly to the anomalous dips at $\Delta v = 0$ in the CO spectra in Fig. 7. In the OH spectra, this will result in weaker absorption at $\Delta v = 0$. However, the different absorption features overlap, which makes it difficult to isolate the effects of these artefacts on the OH absorption.

To compare the strength of the OH absorption in our simulations to observations, we measure the strength of the peak blueshifted absorption feature from the OH 119 μm transition, $1 - S_{\text{p,abs}}^{\text{OH}119} / C_{119}$. We take this to be the minimum in the continuum-normalized spectrum at velocities $< -80 \text{ km s}^{-1}$, which will correspond to the broad outflow component, where $S_{\text{p,abs}}^{\text{OH}119}$ is the flux density of the spectrum at the minimum and C_{119} is the flux density of the continuum. By looking at velocities $< -80 \text{ km s}^{-1}$, we also limit the impact of the artefacts along the boundaries between the high- and low-resolution region, which primarily affect velocities close to $\Delta v = 0$, as in the CO spectra (see Fig. 7). This OH absorption strength is shown in Fig. 11, plotted against the velocity of the absorption peak. The simulations are shown by the coloured circles, while the black symbols show the observed sample of local ULIRGs from González-Alfonso et al. (2017) (see the top panel of their fig. 3). In the observations, the absorption peak was measured by fitting Gaussians to the absorption features.

The two simulations at solar metallicity (nH10_L46_Z1 and nH10_L45_Z1) show OH 119 μm absorption strengths and velocities that are comparable to those observed in the sample of local ULIRGs. However, the absorption strength in the low-metallicity run (nH10_L46_Z0.1) is lower than is observed. González-Alfonso et al. (2017) highlighted that, since the OH 119 μm doublet is optically thick, the absorption strength measures the covering fraction of OH in the outflow. In Fig. 7, we saw that the covering fraction of dense, molecular clumps traced by CO emission is much

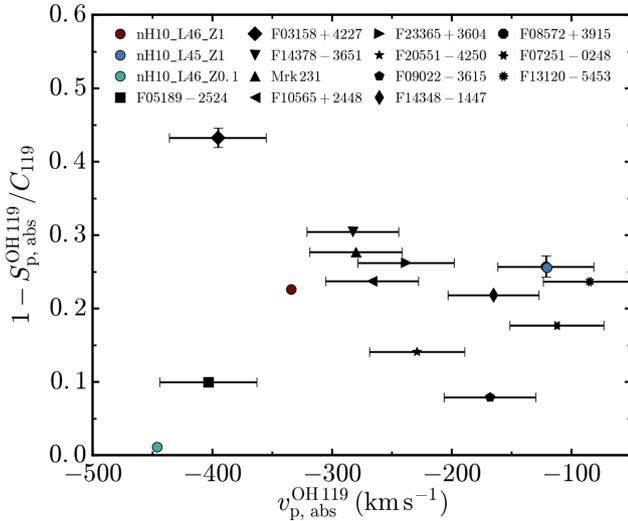


Figure 11. Peak blueshifted OH 119 μm absorption strength ($1 - S_{\text{p,abs}}^{\text{OH}119} / C_{119}$) versus peak absorption velocity ($v_{\text{p,abs}}^{\text{OH}119}$) from the simulations nH10_L46_Z1 (red circle), nH10_L45_Z1 (blue circle), and nH10_L46_Z0.1 (green circle), and from observations of a sample of local ULIRGs (González-Alfonso et al. 2017; black symbols). In the simulations, the peak absorption is defined as the minimum of the continuum-normalized spectrum at velocities $< -80 \text{ km s}^{-1}$, while the observations measure the peak absorption strength from Gaussian fits, again limited to velocities $< -80 \text{ km s}^{-1}$. The two simulations at solar metallicity are in good agreement with the observations, while the low-metallicity run has much weaker OH absorption, likely due to the lower OH covering factor.

lower in the low-metallicity run than in the other two simulations. Therefore, the weak OH absorption in this run is likely due to the low-covering fraction of OH. It is unsurprising that this run does not match the observations, as the observed systems are ULIRGs, with higher metallicities. We also caution that the peak absorption velocity in the low-metallicity run is poorly converged with numerical resolution (see Appendix A).

Several OH 119 μm spectra in the sample of González-Alfonso et al. (2017) show P-Cygni profiles, with redshifted emission in addition to blueshifted absorption. However, we do not find any OH emission in the simulations. It is possible that this is due to the simplified geometry of our simulations. For example, if most of the continuum emission comes from dusty material close to the AGN, inside the outflowing shell, then the side of the shell moving towards the observer will be seen in absorption, whereas the receding side will lie behind the continuum source and so will only be seen in emission, producing a P-Cygni profile. However, our simulations do not include a dusty torus or accretion disc around the AGN, and there is little dust inside the outflowing shell, as this gas is mostly hot ($\sim 10^{10} \text{ K}$), with zero dust abundance. Our simulations therefore represent AGN for which the continuum emission is primarily from the host galaxy itself, outside the outflowing shell, for which both sides of the shell are seen in absorption. See, for example, Fig. 12 of González-Alfonso et al. (2017) for schematic diagrams of the geometries that they assume in the radiative transfer models that they fit to their observations of OH lines in ULIRG outflows.

5.3 IR-traced warm H_2 emission

Cold H_2 at temperatures of a few tens of K, as found in typical giant molecular clouds, is difficult to observe directly, as the lowest rotational transition of the H_2 molecule has an energy of $E/k_B =$

510 K. However, warm H_2 , at temperatures of a few hundred K or higher, exhibits emission lines in the near- and mid-infrared from ro-vibrational transitions. These lines have been observed in a variety of extragalactic systems, including colliding galaxies (Rieke et al. 1985; van der Werf et al. 1993; Sugai et al. 1997; Appleton et al. 2006, 2017), starburst galaxies (Rosenberg, van der Werf & Israel 2013), ULIRGs (Zakamska 2010; Hill & Zakamska 2014), local normal star-forming galaxies from the SINGS sample (Roussel et al. 2007), radio galaxies (Ogle et al. 2007), and AGN-driven outflows (Dasyra & Combes 2011; Rupke & Veilleux 2013a).

We use RADMC-3D to calculate the line emission in our simulations from the four lowest rotational transitions of the ground vibrational state of the H_2 molecule, i.e. the $0 - 0 \text{ S}(0)$ to $\text{S}(3)$ transitions, with rest wavelengths of 28.2–9.7 μm . We use transition probabilities from Wolniewicz, Simbotin & Dalgarno (1998), and we include collisional excitation of H_2 by H I (Wrathmall, GUSDorf & Flower 2007), H II (Gerlich 1990), para- H_2 (Flower & Roueff 1998), ortho- H_2 (Flower & Roueff 1999), He I (Flower, Roueff & Zeippen 1998), and electrons (Draine, Roberge & Dalgarno 1983).

For the CO and OH lines, in Sections 5.1 and 5.2, respectively, we included thermal emission, scattering, and absorption from dust grains, and then subtracted the continuum emission from the line spectra (CO), or divided the total spectra by the continuum (OH). However, at the wavelengths of the warm H_2 lines, the dust continuum was stronger, and the resulting noise in the continuum (since RADMC-3D uses a Monte Carlo method) became large compared to the lines themselves. It was too computationally expensive to increase the number of photon packets to sufficiently reduce this noise, so instead we simply excluded thermal emission and scattering from dust grains for the H_2 lines, and only included dust absorption. The lack of thermal dust emission will not affect our results, as we subtract the continuum. We found that dust scattering had very little effect on the CO spectra (not shown), although the H_2 lines are at a shorter wavelength.

Fig. 12 shows the $\text{S}(0)$ spectra from nH10_L46_Z1 (left), nH10_L45_Z1 (centre), and nH10_L46_Z0.1 (right) after 1 Myr. The shape of this line is similar to that of the CO (1–0) line in these three simulations (see Fig. 7). The dip at $\Delta v = 0 \text{ km s}^{-1}$ is again due to artefacts along the boundary between the high- and low-resolution regions, as discussed in Section 5.1.

Using the line fluxes of each rotational transition, we can calculate the column densities of the different energy levels, and hence deduce an excitation temperature of the IR-traced H_2 gas in a manner similar to what is done in observations. For a given transition, the column density N_u of the upper energy level is:

$$N_u = \frac{4\pi\lambda}{hc} \frac{I}{A} \text{ cm}^{-2}, \quad (14)$$

where λ is the rest wavelength of the line, A is the Einstein coefficient, and I is the intensity of the line in units of $\text{erg s}^{-1} \text{ cm}^{-2} \text{ sr}^{-1}$.

We can then compare the level populations to a thermal distribution with an excitation temperature T_{exc} :

$$N_u/g_u = \frac{N_{\text{tot}} \exp(-E_u/k_B T_{\text{exc}})}{Z(T_{\text{exc}})}, \quad (15)$$

where N_{tot} is the total H_2 column density, E_u is the energy of the given level, and $Z(T)$ is the partition function, which we take from Herbst et al. (1996):

$$Z(T) = \frac{0.0247(T/\text{K})}{1 - \exp(-6000 \text{ K}/T)}. \quad (16)$$

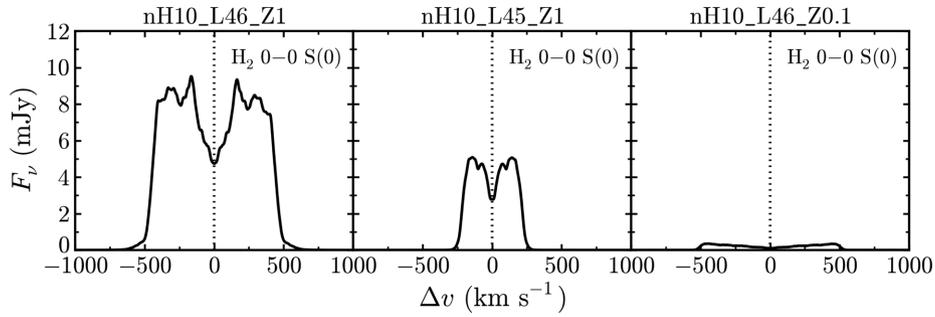


Figure 12. H_2 0 – 0 S(0) spectra from nH10_L46_Z1 (left), nH10_L45_Z1 (centre), and nH10_L46_Z0.1 (right) after 1 Myr, normalized to a distance of 184 Mpc. The dip at $\Delta v = 0 \text{ km s}^{-1}$ is due to artefacts along the boundary between the high- and low-resolution regions. These artefacts will reduce the H_2 column densities calculated from the full line flux, but we find that the H_2 excitation temperatures calculated from line ratios are not strongly affected.

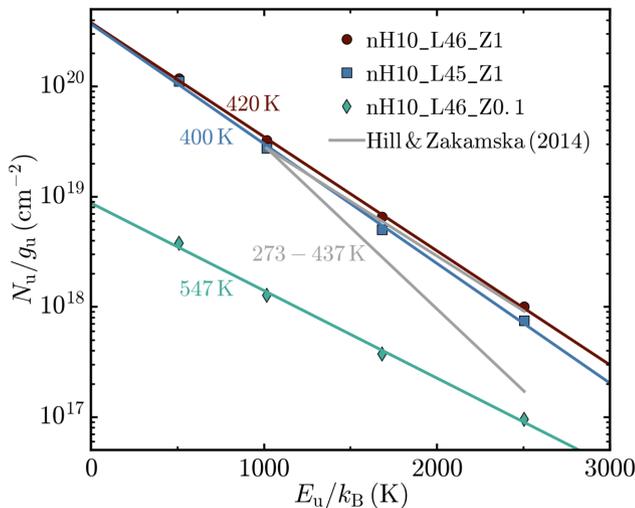


Figure 13. Level populations of H_2 derived from the line fluxes of the four lowest rotational transitions in the ground vibrational state, where N_u and g_u are the column density and statistical weight, respectively, of the upper energy level of each transition, plotted against the upper-level energy. We show simulations nH10_L46_Z1 (red circles), nH10_L45_Z1 (blue squares), and nH10_L46_Z0.1 (green diamonds). The coloured lines show single-temperature thermal distributions fit to each simulation, with excitation temperatures as indicated. The slopes of the two grey lines show the range of excitation temperatures found in the 36 ULIRGs classified as AGN in the sample of Hill & Zakamska (2014) (the normalization of the grey lines is arbitrary). The excitation temperatures in the two simulations at solar metallicity are in good agreement with the observations of Hill & Zakamska (2014), although the low-metallicity run shows a higher excitation temperature.

In equation 15, g_u is the statistical weight of the upper energy level. For rotational state J , $g_u = (2I + 1)(2J + 1)$, where the spin number I is 0 for even J and 1 for odd J .

Fig. 13 shows the level populations, N_u/g_u , calculated from the four rotational transitions in the simulations (coloured symbols), plotted against the energy of the upper level. To calculate N_u from the simulations, we first calculate the intensity I of each line by integrating the line over all velocity channels. This intensity is an average intensity over the area of the emitting region. We therefore only include spatial pixels within the projected radius of the outflow, as determined from the CO emission maps in Fig. 7. We then use the intensity to calculate the column density N_u , averaged over the projected area of the outflow, using equation 14.

The coloured lines in Fig. 13 show single-temperature thermal distributions fit to the level populations of each simulation. Both the

Table 5. H_2 line luminosities from the simulations after 1 Myr. Note that these luminosities are from the high-resolution quadrant of the radiative transfer grid multiplied by 4, to represent the full sphere.

Simulation	Line luminosity (erg s^{-1})			
	S(0)	S(1)	S(2)	S(3)
nH10_L46_Z1	4.0×10^{40}	1.2×10^{42}	8.5×10^{41}	2.2×10^{42}
nH10_L45_Z1	1.1×10^{40}	3.1×10^{41}	1.9×10^{41}	4.8×10^{41}
nH10_L46_Z0.1	1.4×10^{39}	5.3×10^{40}	5.3×10^{40}	2.3×10^{41}

excitation temperature, T_{exc} , and the total H_2 column density of the thermal distribution, N_{tot} , were fit to the values of N_u/g_u from each simulation. We find excitation temperatures of 400–547 K. Hill & Zakamska (2014) used archival *Spitzer* data to measure the S(1) and S(3) lines in 115 ULIRGs from the *IRAS* 1 Jy sample, including 8 type 1 and 28 type 2 AGN. They derived excitation temperatures of 273–437 K, which is in good agreement with our two simulations at solar metallicity, although the low-metallicity run has a higher excitation temperature. The grey lines in Fig. 13 indicate the slopes of thermal distributions with this range of temperatures. The normalization of the grey lines is arbitrary; they are intended to illustrate the slopes of the observed distributions. Although our solar metallicity wind simulations are consistent with the observations of Hill & Zakamska (2014), we note that the observations include all H_2 emission from the galaxies in their sample, which may include a significant non-wind component.

The column densities, N_u , will be affected by the artificial dips in the H_2 lines at $\Delta v = 0$ seen in Fig. 12, which reduce the total line flux, and hence N_u . However, the excitation temperatures are based on line ratios. We find that if we calculate the excitation temperatures in the same way but using only the flux at velocities $|\Delta v| > 300 \text{ km s}^{-1}$ (or $|\Delta v| > 200 \text{ km s}^{-1}$ for nH10_L45_Z1), which avoids the artificial dips, then the excitation temperatures change by at most 2 per cent, compared to using the full lines. Thus the excitation temperatures are not strongly affected by these artificial dips in the spectra.

Using the total column density, N_{tot} , from the best-fitting thermal distributions, we calculate the total mass of IR-traced H_2 that would be inferred from these mid-infrared emission lines. We find that the mass of IR-traced H_2 is approximately equal to the total H_2 mass in the simulations, to within 5 per cent. Thus nearly all of the molecular hydrogen in the outflow is warm and can be traced by the mid-infrared emission lines. Molecular hydrogen in the host galaxy – not modelled here – can have a substantial colder component.

In Table 5 we show the line luminosities for the four lowest rotational H_2 lines in the ground vibrational state from each

simulation. These were calculated in the high-resolution quadrant of the radiative transfer grid and multiplied by 4 to represent the luminosity from the full spherical outflow. For comparison, the 36 AGN host galaxies included in Hill & Zakamska (2014) show S(1) luminosities of 1.7×10^{41} – 1.2×10^{42} erg s⁻¹ and S(3) luminosities of 7.0×10^{40} – 5.8×10^{41} erg s⁻¹. This is comparable to our simulations, at least at solar metallicity, although the S(3) luminosity in the nH10_L46_Z1 run is a little high. We again note that the Hill & Zakamska (2014) sample may include emission from the host galaxy, and not just the AGN outflow.

6 CONCLUSIONS

We have run a suite of hydro-chemical simulations of an isotropic AGN wind interacting with a uniform medium to explore whether *in situ* molecule formation within an AGN-driven wind can explain the fast molecular outflows that have been observed in quasars. This idealized set-up matches the set-up studied analytically by FGQ12. In this work, we focus on the chemistry of the gas to understand under what conditions molecules can form in the wind.

We followed the non-equilibrium chemistry of the gas using a chemical network of 157 species, including 20 molecular species such as H₂, CO, OH, and HCO⁺. The simulations were run for 1 Myr, which corresponds to the typical flow times of observed molecular outflows (e.g. González-Alfonso et al. 2017), at a fiducial resolution of 30 M_⊙ per gas particle in the high-resolution octant for our main runs. We considered a range of initial ambient medium densities (1–10 cm⁻³), AGN luminosities (10⁴⁵–10⁴⁶ erg s⁻¹), and metallicities (0.1–1 Z_⊙). We also varied certain aspects of the model (the cosmic ray ionization rate, AGN UV flux, shielding column density, dust abundance, and the presence of a Jeans limiter), to test how uncertainties in these affect our results.

We then used the publicly available radiative transfer code RADMC-3D (Dullemond et al. 2012) to compute CO, OH, and IR-traced warm H₂ emission and absorption from our simulation snapshots in post-processing. This allowed us to compare the molecular outflows in our simulations to observations, and to make predictions for the conversion factors between CO emission and H₂ mass.

We summarize our main results below.

(i) The hot wind bubble remains at high temperatures ($\gtrsim 10^{10}$ K) throughout the simulations, and so the outflow is approximately energy conserving. However, in all simulations except the low-metallicity run (nH1_L46_Z1), the gas swept up from the ambient ISM cools to below 10⁴ K within 1 Myr (Fig. 1).

(ii) The fiducial simulation (nH10_L46_Z1) produces 2.7×10^8 M_⊙ of H₂ in outflowing gas after 1 Myr, with a molecular fraction of 0.24 and an H₂ outflow rate of 140 M_⊙ yr⁻¹ (Figs 3 and 4). The molecular mass and outflow rate increase with increasing density, AGN luminosity, and metallicity.

(iii) The molecular gas is found at high densities ($n_{\text{H}} \gtrsim 10^3$ cm⁻³) with a wide range of temperatures, from tens of K up to $\sim 10^3$ K. (Fig. 5). A significant fraction of the wind mass reaches densities $n_{\text{H}} > 10^4$ cm⁻³ that may be observable with dense gas tracers such as HCN and HCO⁺.

(iv) The H₂ outflow rate depends only weakly on the cosmic ray ionization rate, while it decreases strongly as we decrease the shielding column density or the dust-to-metals ratio (Fig. 6).

(v) We used the CO line emission computed from our simulations in post-processing to calculate the H₂ outflow rates in the same way as in the observations. We find that the simulations at solar metallicity are slightly lower than, but still reasonably close to,

the observed outflow rates from Ciccone et al. (2014), while the low-metallicity run (nH10_L46_Z0.1) is more than two orders of magnitude below the observations (Fig. 8). The maximum line-of-sight velocities inferred from the CO (1–0) line in the simulations are a factor ≈ 2 lower than in the observations.

(vi) In our reference simulation (nH10_L46_Z1), the CO conversion factor for the (1–0) line is $\alpha_{\text{CO}(1-0)} = 0.13 \text{ M}_{\odot} (\text{K km s}^{-1} \text{ pc}^2)^{-1}$ (Table 4), a factor of 6 lower than is commonly assumed in observations. For the (2–1) and (3–2) lines, we find conversion factors of 0.08 and 0.06 M_⊙ (K km s⁻¹ pc²)⁻¹, respectively.

(vii) The simulations at solar metallicity show significant absorption from the OH doublets at 119 and 79 μm, while the low-metallicity run shows only weak absorption (Fig. 10). The strength of the 119 μm absorption in the solar metallicity runs is in good agreement with the observations of González-Alfonso et al. (2017) (Fig. 11).

(viii) We find strong mid-infrared emission from the four lowest rotational transitions of the ground vibrational state of H₂. The mass of H₂ inferred from these emission lines is within a few percentages of the total H₂ mass, i.e. nearly all of the molecular hydrogen in the outflow is traced by IR emission. The level populations derived from these lines indicate an excitation temperature of 400–547 K (Fig. 13), in good agreement with the observations of Hill & Zakamska (2014), at least at solar metallicity.

We have thus demonstrated that it is possible to form fast molecular outflows *in situ* within an AGN-driven wind, provided that the ambient ISM density and metallicity are sufficiently high (at least 10 cm⁻³ and solar metallicity, respectively). The molecular emission and absorption lines in our simulations are generally in good agreement with those seen in observations of AGN, although the maximum line-of-sight velocities from the CO (1–0) line are a factor ≈ 2 lower than is observed at a given AGN luminosity, and the molecular gas in the simulations may be more highly excited than in observations. However, there are several caveats that we will need to address before we can fully understand the origin of fast molecular outflows in quasars.

The biggest uncertainty is the role that dust grains play in AGN winds. Our simulations assume a constant dust-to-metals ratio equal to that in the Milky Way. If we reduce the dust-to-metals ratio by a factor of 10 or 100 in our reference model, the mass of outflowing H₂ after 1 Myr decreases by a factor of ≈ 8 or ≈ 150 , respectively (Fig. 6). However, the simulations run by Ferrara & Scannapieco (2016) found that dust grains in an AGN wind will be rapidly destroyed by shocks and sputtering, within 10⁴ yr. As discussed in Section 4, it may be possible for dust formation mechanisms, such as the accretion of metals from the gas phase on to grains, to reform the grains after the swept-up gas has cooled. To resolve these issues, we will need to develop a more rigorous treatment for the dust grains, including formation and destruction mechanisms and the resulting evolution of the grain size distribution. This will need to be coupled to the chemistry solver, along with alternative H₂ formation channels via PAH chemistry (e.g. Boschman et al. 2015).

Secondly, our simulations do not include thermal conduction. Brüggén & Scannapieco (2016) explored the effects of thermal conduction in simulations of cold clouds interacting with a hot, fast gas flow. They found that thermal conduction can evaporate the cold clouds, but it can also compress the clouds, allowing them to survive longer. Ferrara & Scannapieco (2016) found that, while conduction can result in a more uniform medium in their simulations of an AGN wind, it does not have a significant impact on the final conditions

of the gas, although their simulations did not include the hot wind bubble.

Furthermore, thermal conduction can cause some of the cold gas in the shell of swept-up material to expand into, and subsequently mix with, the hot wind bubble. Castor, McCray & Weaver (1975) showed that, in stellar wind bubbles, the mass of the hot, shocked wind region is dominated by gas that came from the swept-up shell in such a conductive flow. The AGN-driven winds that we study here are analogous to the stellar wind bubbles of Castor et al. (1975), although the latter are at much lower velocities and energies. In AGN-driven galactic winds, magnetic fields may act to limit the effects of thermal conduction.

Thirdly, we show in Appendix A that the H_2 outflow rate is not well converged with resolution in our simulations. In the low-luminosity run (nH10_L45_Z1), the H_2 outflow rate increases by a factor of ≈ 4 as the mass resolution is increased from 240 to $10 M_\odot$ per gas particle. It is encouraging that the H_2 outflow rate increases with increasing resolution, as this strengthens our main conclusion that molecules can form in an AGN wind.

Finally, our simulations follow an idealized scenario of an isotropic AGN wind interacting with a uniform ambient medium. This allows us to focus on the details of the molecular chemistry, and makes it easier to understand how the chemistry is affected by the physical conditions such as density and AGN luminosity. However, we will need to explore the effects of more realistic geometries, such as a radially declining density profile, inhomogeneities in the ambient medium, and the presence of a galactic disc. It will therefore be interesting to apply our methods for injecting an AGN wind and following the non-equilibrium chemistry to more realistic galaxy simulations.

In this work we have focused on AGN-driven winds. However, molecular outflows have also been observed in starburst galaxies (e.g. Nakai et al. 1987; Veilleux, Rupke & Swaters 2009; Bolatto et al. 2013b; Hill & Zakamska 2014; Beirão et al. 2015). It will thus be interesting to extend this study to the case of star formation-driven galactic winds. Indeed, we saw that some of our results, such as the H_2 outflow rate and velocity (Fig. 8), scale with the AGN luminosity. Therefore, it might be possible to generalize our results to star formation-driven outflows by scaling down to the lower luminosities of starburst galaxies.

Our simulations predict strong infrared emission from warm H_2 , tracing nearly all of the mass of outflowing molecular gas. These H_2 lines may be observable by the upcoming *James Webb Space Telescope* (*JWST*). Future *JWST* observations could therefore provide valuable insight into the nature of these fast molecular outflows and allow us to further constrain our models.

ACKNOWLEDGEMENTS

We thank the anonymous referee for their detailed report, which improved the quality of this manuscript. We also thank Chiara Feruglio and Eduardo González-Alfonso for their detailed comments, as well as Nadia Zakamska, Mike Shull, Desika Narayanan and Fabrizio Fiore for useful discussions. We are grateful to Phil Hopkins and Paul Torrey for help with the *GIZMO* code. AJR is supported by the Lindheimer fellowship at Northwestern University. CAFG was supported by NSF through grants AST-1412836, AST-1517491, and CAREER award AST-1652522, by NASA through grant NNX15AB22G, and by CXO through grant TM7-18007X. The simulations used in this work were run on the Stampede supercomputer at the Texas Advanced Computing Center (TACC) through allocations TG-AST160035 and TG-AST160059 granted

by the Extreme Science and Engineering Discovery Environment (XSEDE), which is supported by NSF grant number ACI-154562; the Pleiades supercomputer through allocation s1480, provided through the NASA Advanced Supercomputing (NAS) Division at Ames Research Center; and the Quest computing cluster at Northwestern University, which is jointly supported by the Office of the Provost, the Office for Research, and Northwestern University Information Technology.

REFERENCES

- Aalto S., Garcia-Burillo S., Muller S., Winters J. M., van der Werf P., Henkel C., Costagliola F., Neri R., 2012, *A&A*, 537, 44
Aalto S. et al., 2015, *A&A*, 574, 85
Ahn K., Shapiro P. R., 2007, *MNRAS*, 375, 881
Alatalo K. et al., 2011, *ApJ*, 735, 88
Anglés-Alcázar D., Özel F., Davé R., 2013, *ApJ*, 770, 5
Anglés-Alcázar D., Özel F., Davé R., Katz N., Kollmeier J. A., Oppenheimer B. D., 2015, *ApJ*, 800, 127
Anglés-Alcázar D., Davé R., Faucher-Giguère C.-A., Özel F., Hopkins P. F., 2017, *MNRAS*, 464, 2840
Appleton P. N. et al., 2006, *ApJ*, 639, L51
Appleton P. N. et al., 2017, *ApJ*, 836, 76
Asano R. S., Takeuchi T. T., Hirashita H., Inoue A. K., 2013, *Earth Planets Space*, 65, 213
Bate M. R., Burkert A., 1997, *MNRAS*, 288, 1060
Beirão P. et al., 2015, *MNRAS*, 451, 2640
Bolatto A. D., Wolfire M., Leroy A. K., 2013a, *ARA&A*, 51, 207
Bolatto A. D. et al., 2013b, *Nature*, 499, 450
Booth C. M., Schaye J., 2009, *MNRAS*, 398, 53
Boschman L., Cazaux S., Spaans M., Hoekstra R., Schlathöler T., 2015, *A&A*, 579, 72
Bower R. G., Benson A. J., Crain R. A., 2012, *MNRAS*, 422, 2816
Bower R. G., Schaye J., Frenk C. S., Theuns T., Schaller M., Crain R. A., McAlpine S., 2017, *MNRAS*, 465, 32
Brüggen M., Scannapieco E., 2016, *ApJ*, 822, 31
Castor J., McCray R., Weaver R., 1975, *ApJ*, 200, L107
Cazaux S., Tielens A. G. G. M., 2002, *ApJ*, 575, L29
Cicone C., Feruglio C., Maiolino R., Fiore F., Piconcelli E., Menci N., Aussel H., Sturm E., 2012, *A&A*, 543, 99
Cicone C. et al., 2014, *A&A*, 562, 21
Cicone C. et al., 2015, *A&A*, 574, 14
Costa T., Sijacki D., Haehnelt M. G., 2014, *MNRAS*, 444, 2355
Crain R. A. et al., 2015, *MNRAS*, 450, 1937
Croton D. J. et al., 2006, *MNRAS*, 365, 11
Dasyra K. M., Combes F., 2011, *A&A*, 533, L10
Dasyra K. M., Combes F., Oosterloo T., Oonk J. B. R., Morganti R., Salomé P., Vlahakis N., 2016, *A&A*, L7
Di Matteo T., Springel V., Hernquist L., 2005, *Nature*, 433, 604
Downes D., Solomon P. M., 1998, *ApJ*, 507, 615
Draine B. T., Bertoldi F., 1996, *ApJ*, 468, 269
Draine B. T., Roberge W. G., Dalgarno A., 1983, *ApJ*, 264, 485
Duarte-Cabral A., Acreman D. M., Dobbs C. L., Mottram J. C., Gibson J., Brunt C. M., Douglas K. A., 2015, *MNRAS*, 447, 2144
Dullemond C. P., Juhasz A., Pohl A., Sereshti F., Shetty R., Peters T., Commercon B., Flock M., 2012, *Astrophysics Source Code Library*, record ascl:1202.015
Dwek E., Foster S. M., Vancura O., 1996, *ApJ*, 457, 244
Egami E., Rieke G. H., Fadda D., Hines D. C., 2006, *ApJ*, 652, L21
Faucher-Giguère C.-A., Quataert E., 2012, *MNRAS*, 425, 605 (FGQ12)
Faucher-Giguère C.-A., Quataert E., Murray N., 2012, *MNRAS*, 420, 1347
Feldmann R., Gnedin N. Y., Kravtsov A. V., 2012, *ApJ*, 747, 124
Feldmann R., Quataert E., Hopkins P. F., Faucher-Giguère C.-A., Kereš D., 2017, *MNRAS*, 470, 1050
Ferland G. J. et al., 2013, *Rev. Mex. Astron. Astrofis.*, 49, 137
Ferrara A., Scannapieco E., 2016, *ApJ*, 833, 46

- Feruglio C., Maiolino R., Piconcelli E., Menci N., Aussel H., Lamastra A., Fiore F., 2010, *A&A*, 518, L155
- Feruglio C. et al., 2013a, *A&A*, 549, 51
- Feruglio C., Fiore F., Piconcelli E., Ciccone C., Maiolino R., Davies R., Sturm E., 2013b, *A&A*, 558, 87
- Feruglio C. et al., 2015, *A&A*, 583, 99
- Field G. B., 1965, *ApJ*, 142, 531
- Fiore F. et al., 2017, *A&A*, 601, 143
- Fischer J. et al., 2010, *A&A*, 518, L41
- Fixsen D. J., Bennett C. L., Mather J. C., 1999, *ApJ*, 526, 207
- Flower D. R., Roueff E., 1998, *J. Phys. B*, 31, 2935
- Flower D. R., Roueff E., 1999, *J. Phys. B*, 32, 3399
- Flower D. R., Roueff E., Zeppen C. J., 1998, *J. Phys. B*, 31, 1105
- Furlanetto S. R., Stoever S. J., 2010, *MNRAS*, 404, 1869
- Gabor J. M., Davé R., 2012, *MNRAS*, 427, 1816
- Gaspari M., Sadowski A., 2017, *ApJ*, 837, 149
- Genzel R. et al., 2012, *ApJ*, 746, 69
- Gerlich D., 1990, *J. Chem. Phys.*, 92, 2377
- Gibson R. R. et al., 2009, *ApJ*, 692, 758
- Glenn J., Rangwala N., Maloney P. R., Kamenetzky J. R., 2015, *ApJ*, 800, 105
- Gnedin N. Y., Tassis K., Kravtsov A. V., 2009, *ApJ*, 697, 55
- Goicoechea J. R., Cernicharo J., 2002, *ApJ*, 576, L77
- González-Alfonso E. et al., 2014, *A&A*, 561, 27
- González-Alfonso E. et al., 2017, *ApJ*, 836, 11
- Greene J. E., Zakamska N. L., Smith P. S., 2012, *ApJ*, 746, 86
- Gültekin K. et al., 2009, *ApJ*, 698, 198
- Haardt F., Madau P., 2001, in Neumann D. M., Tran J. T. V., eds, *XXIst Moriond Astrophys. Meeting, Clusters of Galaxies and the High Redshift Universe Observed in X-rays Editions Frontières, Paris*, 64
- Harrison C. M. et al., 2012, *MNRAS*, 426, 1073
- Harrison C. M., Alexander D. M., Mullaney J. R., Swinbank A. M., 2014, *MNRAS*, 441, 3306
- Henriques B. M. B. et al., 2015, *MNRAS*, 451, 2663
- Herbst T. M., Beckwith S. V. W., Glindemann A., Tacconi-Garman L. E., Kroker H., Krabbe A., 1996, *AJ*, 111, 2403
- Hill M. J., Zakamska N. L., 2014, *MNRAS*, 439, 2701
- Hirashita H., Nozawa T., 2017, preprint ([arXiv:1701.07200](https://arxiv.org/abs/1701.07200))
- Hirashita H., Nozawa T., Villaume A., Srinivasan S., 2015, *MNRAS*, 454, 1620
- Hodges-Kluck E., Cafmeyer J., Bregman J. N., 2016, *ApJ*, 833, 58
- Hoopes C. G. et al., 2005, *ApJ*, 619, L99
- Hopkins P. F., 2015, *MNRAS*, 450, 53
- Hopkins P. F., Cox T. J., Kereš D., Hernquist L., 2008, *ApJS*, 175, 390
- Hopkins P. F., Quataert E., Murray N., 2011, *MNRAS*, 417, 950
- Hopkins P. F. et al., 2017, preprint ([arXiv:1702.06148](https://arxiv.org/abs/1702.06148))
- Hutton S., Ferreras I., Wu K., Kuin P., Breeveld A., Yershov V., Cropper M., Page M., 2014, *MNRAS*, 440, 150
- Indriolo N., McCall B. J., 2012, *ApJ*, 745, 91
- Israel F. P., 1997, *A&A*, 328, 471
- Kennicutt R. C., Jr, Evans N. J., II, 2012, *ARA&A*, 50, 531
- Kilerci Eser E., Goto T., Doi Y., 2014, *ApJ*, 797, 54
- King A., 2003, *ApJ*, 596, L27
- Klein R. I., McKee C. F., Colella P., 1994, *ApJ*, 420, 213
- Krumholz M. R., 2013, *MNRAS*, 436, 2747
- Krumholz M. R., Leroy A. K., McKee C. F., 2011, *ApJ*, 731, 25
- Lacki B. C., Thompson T. A., Quataert E., 2010, *ApJ*, 717, 1
- Langer W. D., 1978, *ApJ*, 225, 860
- Leroy A. K. et al., 2011, *ApJ*, 737, 12
- Liu G., Zakamska N. L., Greene J. E., Nesvadba N. P. H., Liu X., 2013, *MNRAS*, 436, 2576
- Lotz W., 1967, *ApJS*, 14, 207
- Maiolino R., Ruiz M., Rieke G. H., Papadopoulos P., 1997, *ApJ*, 485, 552
- Maiolino R. et al., 2012, *MNRAS*, 425, 66
- Maloney P., Black J. H., 1988, *ApJ*, 325, 389
- Martin P. G., Whittet D. C. B., 1990, *ApJ*, 357, 113
- Martin C. L., Dijkstra M., Henry A., Soto K. T., Danforth C. W., Wong J., 2015, *ApJ*, 803, 6
- Mathis J. S., Rimpl W., Nordsieck K. H., 1977, *ApJ*, 217, 425
- McElroy D., Walsh C., Markwick A. J., Cordiner M. A., Smith K., Millar T. J., 2013, *A&A*, 550, 36
- McKee C. F., Krumholz M. R., 2010, *ApJ*, 709, 308
- Meléndez M. et al., 2015, *ApJ*, 804, 46
- Moe M., Arav N., Bautista M. A., Korista K. T., 2009, *ApJ*, 706, 525
- Murray N., Chiang J., Grossman S. A., Voit G. M., 1995, *ApJ*, 451, 498
- Murray N., Quataert E., Thompson T. A., 2005, *ApJ*, 618, 569
- Nakai N., Hayashi M., Handa T., Sofue Y., Hasegawa T., Sasaki M., 1987, *PASJ*, 39, 685
- Narayanan D. et al., 2006, *ApJ*, 642, L107
- Narayanan D. et al., 2008, *ApJS*, 176, 331
- Nardini E. et al., 2015, *Science*, 347, 860
- Nelson R. P., Langer W. D., 1997, *ApJ*, 482, 796
- Neufeld D. A., Dalgarno A., 1989, *ApJ*, 340, 869
- Nims J., Quataert E., Faucher-Giguère C.-A., 2015, *MNRAS*, 447, 3612
- Nozawa T., Kozasa T., Habe A., 2006, *ApJ*, 648, 435
- Offer A. R., van Hemert M. C., van Dishoeck E. F., 1994, *J. Chem. Phys.*, 100, 362
- Ogle P., Antonucci R., Appleton P. N., Whysong D., 2007, *ApJ*, 668, 699
- Richings A. J., Schaye J., 2016, *MNRAS*, 458, 270
- Richings A. J., Schaye J., Oppenheimer B. D., 2014a, *MNRAS*, 440, 3349
- Richings A. J., Schaye J., Oppenheimer B. D., 2014b, *MNRAS*, 442, 2780
- Rieke G. H., Cutri R. M., Black J. H., Kailey W. F., McAlary C. W., Lebofsky M. J., Elston R., 1985, 290, 116
- Robertson B. E., Kravtsov A. V., 2008, *ApJ*, 680, 1083
- Rosenberg M. J. F., van der Werf P. P., Israel F. P., 2013, *A&A*, 550, A12
- Roussel H. et al., 2007, *ApJ*, 669, 959
- Roussel H. et al., 2010, *A&A*, 518, L66
- Rupke D. S. N., Veilleux S., 2011, *ApJ*, 729, 27
- Rupke D. S. N., Veilleux S., 2013a, *ApJ*, 775, L15
- Rupke D. S. N., Veilleux S., 2013b, *ApJ*, 768, 75
- Rupke D. S. N., Veilleux S., Baker A. J., 2008, *ApJ*, 674, 172
- Safrank-Shrader C., Krumholz M. R., Kim C.-G., Ostriker E. C., Klein R. I., Li S., McKee C. F., Stone J. M., 2017, *MNRAS*, 465, 885
- Sazonov S. Y., Ostriker J. P., Sunyaev R. A., 2004, *MNRAS*, 347, 144
- Scannapieco E., 2017, *ApJ*, 837, 28
- Scannapieco E., Brügggen M., 2015, *ApJ*, 805, 158
- Schneider E. E., Robertson B. E., 2017, *ApJ*, 834, 144
- Schöier F. L., van der Tak F. F. S., van Dishoeck E. F., Black J. H., 2005, *A&A*, 432, 369
- Shetty R., Glover S. C., Dullemond C. P., Klessen R. S., 2011, *MNRAS*, 412, 1686
- Silich S., Tenorio-Tagle G., Muñoz-Tuñón C., 2003, *ApJ*, 590, 791
- Silich S., Tenorio-Tagle G., Rodríguez-González A., 2004, *ApJ*, 610, 226
- Silk J., 1970, *ApL*, 5, 283
- Silk J., Rees M. J., 1998, *A&A*, 331, L1
- Smith R. J., Glover S. C. O., Clark P. C., Klessen R. S., Springel V., 2014, *MNRAS*, 441, 1628
- Sobolev V. V., 1957, *SvA*, 1, 678
- Spoon H. W. W. et al., 2013, *ApJ*, 775, 127
- Springel V., Di Matteo T., Hernquist L., 2005, *ApJ*, 620, L79
- Stern J., Faucher-Giguère C.-A., Zakamska N. L., Hennawi J. F., 2016, *ApJ*, 819, 130
- Sternberg A., Le Petit F., Roueff E., Le Bourlot J., 2014, *ApJ*, 790, 10
- Stone M., Veilleux S., Meléndez M., Sturm E., Graciá-Carpio J., González-Alfonso E., 2016, *ApJ*, 826, 111
- Sturm E. et al., 2011, *ApJ*, 733, L16
- Sugai H., Malkan M. A., Ward M. J., Davies R. I., McClean I. S., 1997, *ApJ*, 481, 186
- Sun A.-L., Greene J. E., Zakamska N. L., Nesvadba N. P. H., 2014, *ApJ*, 790, 160
- Tenorio-Tagle G., Wünsch R., Silich S., Palouš J., 2007, *ApJ*, 658, 1196
- Thompson T. A., Quataert E., Zhang D., Weinberg D. H., 2016, *MNRAS*, 455, 1830
- Tielens A. G. G. M., Hollenbach D., 1985, *ApJ*, 291, 722
- Tombesi F., Meléndez M., Veilleux S., Reeves J. N., González-Alfonso E., Reynolds C. S., 2015, *Nature*, 519, 436

Trayford J. W., Theuns T., Bower R. G., Crain R. A., Lagos C. d. P., Schaller M., Schaye J., 2016, *MNRAS*, 460, 3925

Tremmel M., Karcher M., Governato F., Volonteri M., Quinn T. R., Pontzen A., Anderson L., 2017, *MNRAS*, 470, 1121

Truelove J. K., Klein R. I., McKee C. F., Holliman J. H., II, Howell L. H., Greenough J. A., 1997, *ApJ*, 489, L179

van der Werf P. P., Genzel R., Krabbe A., Blietz M., Lutz D., Drapatz S., Ward M. J., Forbes D. A., 1993, *ApJ*, 405, 522

van der Werf P. P. et al., 2010, *A&A*, 518, L42

van Dishoeck E. F., Black J. H., 1986, *ApJS*, 62, 109

van Dishoeck E. F., Black J. H., 1988, *ApJ*, 334, 771

Veilleux S., Rupke D. S. N., Swaters R., 2009, *ApJ*, 700, L149

Veilleux S. et al., 2013, *ApJ*, 776, 27

Wang B., 1995, *ApJ*, 444, 590

Weingartner J. C., Draine B. T., 2001, *ApJ*, 548, 296

Weiß A., Neininger N., Hüttemeister S., Klein U., 2001, *A&A*, 365, 571

Weiß A., Walter F., Scoville N. Z., 2005, *A&A*, 438, 533

Weymann R. J., Carswell R. F., Smith M. G., 1981, *ARA&A*, 19, 41

Wiersma R. P. C., Schaye J., Smith B. D., 2009, *MNRAS*, 393, 99

Wiklind T., Combes F., Henkel C., 1995, *A&A*, 297, 643

Wolfire M. G., McKee C. F., Hollenbach D., Tielens A. G. G. M., 2003, *ApJ*, 587, 278

Wolfire M. G., Hollenbach D., McKee C. F., 2010, *ApJ*, 716, 1191

Wolniewicz L., Simbotin I., Dalgarno A., 1998, *ApJS*, 115, 293

Wrathmall S. A., Gusdorf A., Flower D. R., 2007, *MNRAS*, 382, 133

Wyithe J. S. B., Loeb A., 2003, *ApJ*, 595, 614

Yang B., Stancil P. C., Balakrishnan N., Forrey R. C., 2010, *ApJ*, 718, 1062

Zakamska N. L., 2010, *Nature*, 465, 60

Zhang D., Thompson T. A., Quataert E., Murray N., 2017, *MNRAS*, 468, 4801

Zubovas K., King A., 2012, *ApJ*, 745, L34

Zubovas K., King A. R., 2014, *MNRAS*, 439, 400

APPENDIX A: RESOLUTION TESTS

We ran all four parameter variations runs (see Table 1) at a standard resolution of $30 M_{\odot}$ per gas particle and a minimum gas gravitational softening of 0.1 pc. These four runs were also repeated with a factor 8 lower mass resolution ($240 M_{\odot}$ per particle), with a minimum gravitational softening of 0.2 pc. We additionally ran the low-luminosity simulation (nH10_L45_Z1) at a factor 3 higher mass resolution of $10 M_{\odot}$ per particle and a gravitational softening of 0.07 pc. We chose the low-luminosity run for the high-resolution test as it was computationally cheaper than the reference run (240 000 and 840 000 CPU hours, respectively, at our fiducial resolution), while still forming a significant molecular outflow. Also, note that the increase in mass resolution is only a factor of 3 in the high-resolution run, rather than the factor of 8 from the low to standard resolution. Any higher resolution than this would have been prohibitively expensive.

To test whether the formation of fast molecular outflows in our simulations is sensitive to the resolution, we show in Fig. A1 the mass outflow rate of molecular hydrogen versus time, as in Fig. 4, for the different resolution levels (low resolution, dashed curves; standard resolution, solid curves; and high resolution, dotted curve). The outflow rate was calculated in the high-resolution wedge, including only particles outflowing with $v_{\text{out}} > 100 \text{ km s}^{-1}$, then multiplied by 17 to give the outflow rate in the full spherical shell. We only show the three parameter variations that form molecules.

The H_2 outflow rate increases with increasing resolution in all runs, by a factor $\approx 2\text{--}4$ from low to standard resolution, and by a factor ≈ 1.5 from standard to high resolution. The H_2 outflow rate is thus not well converged with resolution. However, since it tends to increase with increasing resolution, this would further strengthen

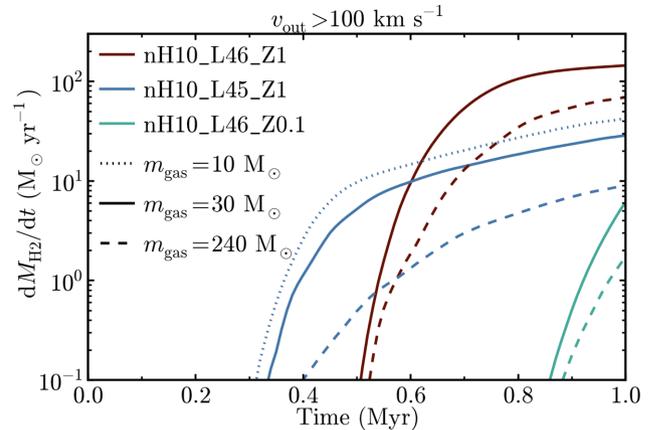


Figure A1. Mass outflow rates of H_2 versus time at low resolution (dashed curves), standard resolution (solid curves), and high resolution (dotted curve), for runs nH10_L46_Z1 (red curves), nH10_L45_Z1 (blue curves), and nH10_L46_Z0.1 (green curves). The H_2 outflow rate tends to increase with increasing resolution.

our main conclusion that fast molecular outflows can form *in situ* within AGN-driven winds.

The increase in H_2 outflow rate at high resolution is primarily due to an increase in the H_2 fraction of the outflowing material, rather than an increase in the total mass of gas in the outflow. The molecular fractions of CO, OH, and HCO^+ also increase with increasing resolution, although the differences are largest for H_2 . We note that the H_2 mass fraction $f_{\text{H}_2} = M_{\text{H}_2}/M_{\text{H}_{\text{tot}}}$ at 1 Myr in the standard resolution run reaches 0.24 (see Fig. 3) and by definition cannot exceed 1.0. Thus, our standard resolution run cannot underestimate the H_2 mass outflow rate by a large factor.

To test how our predictions for molecular emission and absorption lines are affected by numerical resolution, we repeated the radiative transfer calculations presented in Section 5 for the low- and high-resolution runs. Fig. A2 shows the H_2 outflow rate calculated from the CO 1–0 luminosity integrated over the wing of the CO line and assuming a standard ULIRG CO to H_2 conversion factor of $0.8 M_{\odot} (\text{K km s}^{-1} \text{pc}^2)^{-1}$ (top panel), and the maximum line of sight velocity from the CO spectra (bottom panel), plotted against AGN luminosity, as in Fig. 8. We show the low (squares), standard (circles), and high (stars) resolution runs, for nH10_L46_Z1 (red), nH10_L45_Z1 (blue), and nH10_L46_Z0.1 (green). The H_2 outflow rates tend to increase with increasing resolution, particularly in the low-metallicity run, although the fiducial run is reasonably well converged. The maximum velocities show good convergence in all runs.

Fig. A3 shows CO line ratios relative to the 1–0 transition, where the line fluxes are measured in Jy km s^{-1} , as in Fig. 9. The nH10_L46_Z1 runs (red curves) are well converged, whereas in the low-luminosity and low-metallicity runs, the CO line ratios are lower in the low-resolution runs. The standard and high-resolution runs of nH10_L45_Z1 are well converged.

Fig. A4 shows the CO to H_2 conversion factor, α_{CO} , calculated for the three lowest rotational CO lines. At solar metallicity (red and blue symbols), α_{CO} tends to increase with increasing resolution, by ≈ 25 per cent from low to standard resolution. The numerical convergence of α_{CO} is much poorer in the low-metallicity run, varying by up to a factor of 2 from low to standard resolution.

Fig. A5 shows the strength of the blue-shifted absorption peak of the OH 119 μm line versus the peak absorption velocity, calculated as in Fig. 11. In the nH10_L46_Z1 run, the peak absorption

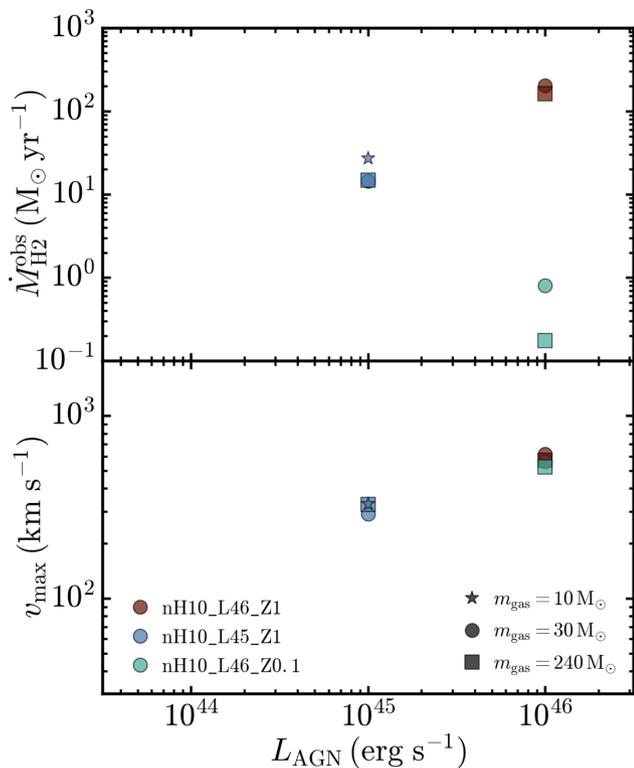


Figure A2. H₂ outflow rates calculated from CO emission (top panel) and maximum line-of-sight velocity from the CO spectra (bottom panel) plotted against AGN luminosity for runs nH10_L46_Z1 (red), nH10_L45_Z1 (blue), and nH10_L46_Z0.1 (green), at low (squares), standard (circles), and high (stars) resolution. The CO-derived H₂ outflow rate tends to increase with increasing resolution, especially in the low-metallicity run, while the maximum velocities are well converged.

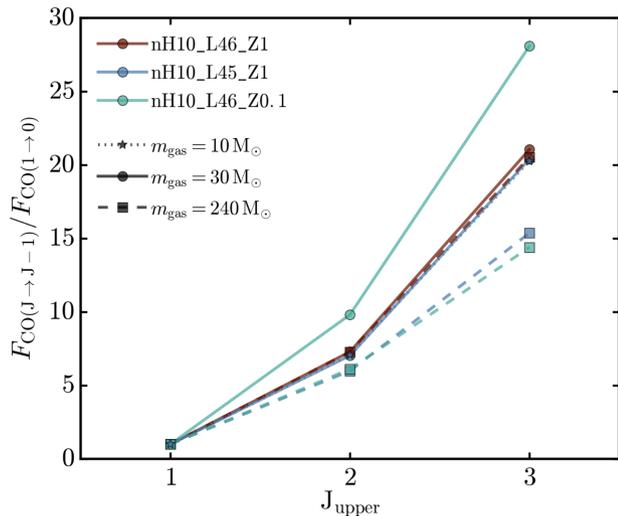


Figure A3. CO lines ratios measured from runs nH10_L46_Z1 (red), nH10_L45_Z1 (blue), and nH10_L46_Z0.1 (green), at low (squares and dashed curves), standard (circles and solid curves), and high (stars and dotted curves) resolution. The nH10_L46_Z1 runs are well converged, while the CO ratios in the low-luminosity and low-metallicity runs are lower at low resolution. The standard and high-resolution runs of nH10_L45_Z1 are well converged.

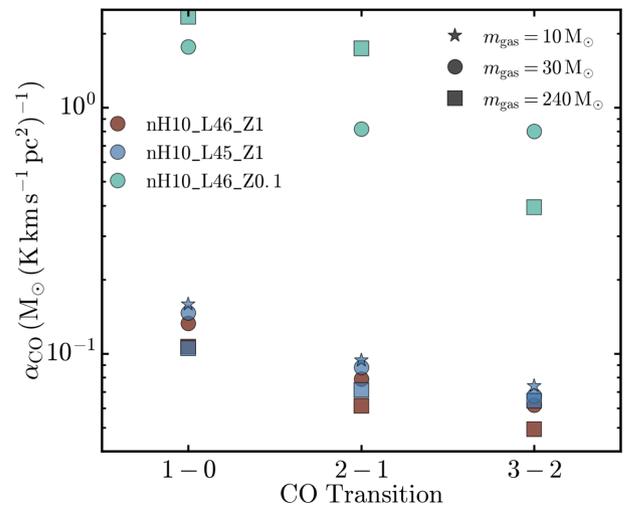


Figure A4. Resolution dependence of the CO to H₂ conversion factor, α_{CO} , for the three lowest rotational transitions of CO. In the two runs at solar metallicity, α_{CO} increases with increasing resolution, by ≈ 25 per cent from low to standard resolution. The low-metallicity run shows much poorer convergence in α_{CO} , which changes by up to a factor of 2 from low to standard resolution.

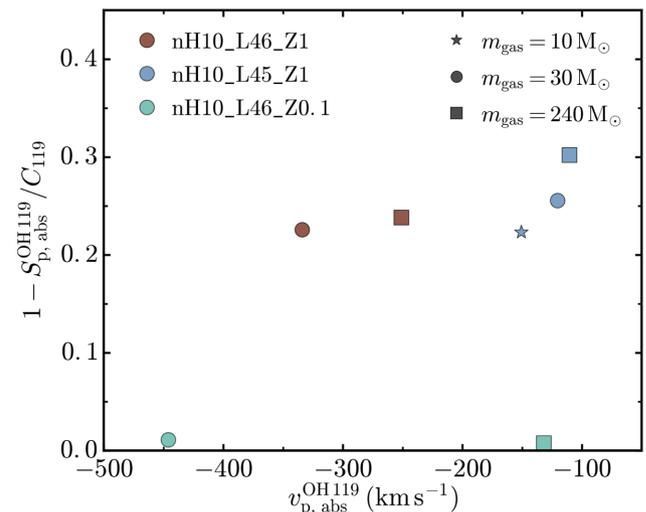


Figure A5. Resolution dependence of the absorption strength of the blue-shifted absorption feature of the 119 μm OH line versus velocity of the absorption peak. The peak velocity of the OH 119 μm absorption is poorly converged in the low-metallicity run (green), although this run only shows very weak OH absorption. In the two runs at solar metallicity, the OH absorption is weaker and more strongly blue-shifted at higher resolution, although, comparing to the black symbols in Fig. 11, we see that these two runs are still in agreement with the observations of González-Alfonso et al. (2017) at all resolution levels.

velocity becomes more strongly blue-shifted as resolution increases, decreasing from -250 km s^{-1} at low resolution to -330 km s^{-1} at standard resolution. However, the strength of the absorption peak is well converged in this run. In the low-luminosity run, the absorption peak becomes weaker [from $(1 - S_{\text{p,abs}}^{\text{OH119}} / C_{119}) = 0.30$ to 0.22] and more strongly blue-shifted (from -110 to -150 km s^{-1}) as the resolution increases from the low-resolution to the high-resolution run. However, if we compare to the black symbols in Fig. 11, we see that all runs at solar metallicity remain in good agreement with the observations of González-Alfonso et al. (2017). The low-metallicity

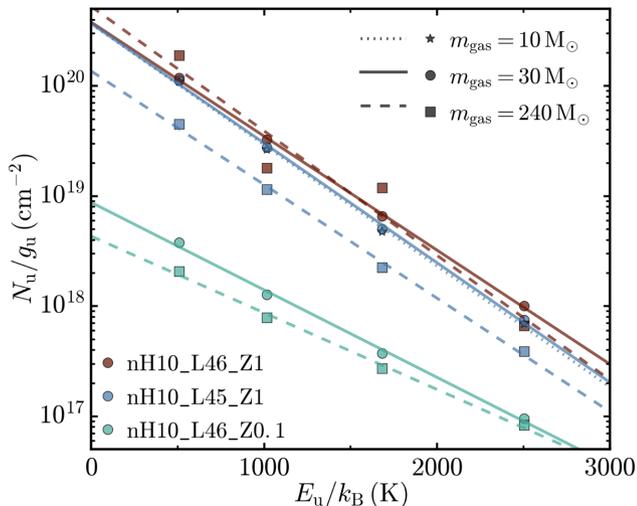


Figure A6. H_2 excitation diagram for different resolution levels. The symbols show H_2 level populations calculated from the line fluxes, plotted as the column density of the upper level, N_u , divided by the statistical weight, g_u , while the curves show single-temperature thermal distributions fitted to the level populations. The H_2 column densities increase with increasing resolution, consistent with the trend of increasing H_2 mass in the outflow with increasing resolution. However, the H_2 excitation temperatures, as determined from the slope of the best-fitting thermal distributions, are well converged.

run, on the other hand, shows much poorer convergence, as the peak absorption velocity decreases from -130 km s^{-1} at low resolution to -450 km s^{-1} at standard resolution. This is likely because this run shows very weak OH absorption, so the peak absorption velocity is poorly constrained.

Fig. A6 shows the level populations of H_2 , calculated from the H_2 line fluxes, plotted against the energy of the upper level of each transition, as in Fig. 13. The H_2 column densities increase with increasing resolution, which is consistent with the trend of increasing H_2 mass with increasing resolution. However, the H_2 excitation temperatures, as indicated by the slope of the best-fitting single-temperature thermal distributions (curves), are well converged. In the fiducial, low-luminosity and low-metallicity runs, the H_2 excitation temperature varies by 9, 6 and 13 per cent, respectively, between resolution levels.

APPENDIX B: LOCAL SHIELDING APPROXIMATION

To correctly follow the attenuation of the photoionization and photodissociation rates in the simulations, we would need to include a treatment for the full 3D multifrequency radiative transport of the UV radiation from the AGN, coupled to the non-equilibrium chemistry solver module. However, this is computationally expensive, so instead we use a local Sobolev-like approximation that assumes that each gas particle is only shielded locally, by gas within a shielding length L_{sh} that depends on the local density ρ and density gradient (see Section 2.2).

In this section we test this approximation by comparing the column densities of individual particles in the final snapshot from simulation nH10_L46_Z1 after 1 Myr to the true column densities calculated by ray tracing from the AGN to the gas particles in post-processing.

Fig. B1 shows this comparison between the Sobolev-like approximation column densities (sob) and the actual column

densities from ray tracing (actual). The top six panels compare the total hydrogen column densities (top left), along with the column densities of H I , H_2 , He I , He II , and CO . The bottom left-hand panel compares the approximate and actual dust extinctions, $A_v = 4.0 \times 10^{-22} N_{\text{H,tot}} Z/Z_\odot$, which include only gas particles below 10^5 K . The solid lines indicate the one-to-one relation between the approximate and actual column densities, while the dashed and dotted lines show factors of 10 and 100, respectively, deviations from the one-to-one relation.

We include only particles within the high-resolution wedge. We also exclude the ambient ISM gas by including only particles that are outflowing, with $v_{\text{out}} > 0 \text{ km s}^{-1}$. There is much poorer agreement between the Sobolev-like approximation and the actual column densities in the ambient ISM than in the outflowing gas. This is unsurprising, because the Sobolev-like approximation uses the density gradient, which would be zero (resulting in an infinite shielding length) if the ambient ISM were perfectly uniform. The shielding approximation that we use assumes that the gas is arranged in clumps that shield themselves, but this assumption is not valid in a uniform medium. However, we find that any molecules in the ambient ISM are destroyed when the forward shock sweeps up this gas into the outflow. Thus the molecular content of the outflow is independent of the chemical state of the ambient ISM.

In the top left-hand panel of Fig. B1 we see that, while the total hydrogen column density in the outflow computed using the Sobolev-like approximation on average follows the actual column density, there is a large scatter of an order of magnitude, with some particles underestimated by more than two orders of magnitude. To explore how these uncertainties might affect the resulting molecular outflows, we ran one simulation with column densities lowered by a factor of 10 (this run is labelled ‘lowShield’ in Fig. 6; see also the discussion in Section 4). We found that this reduced the H_2 outflow rate after 1 Myr by a factor of 4.3 compared to the reference model. However, reducing all column densities by a factor of 10 is a worst-case scenario, and would correspond to all particles lying along the upper dashed lines in Fig. B1. In practice, many particles lie below the one-to-one relation in this figure, i.e. the Sobolev-like approximation underestimates the actual column density for these particles. This would tend to underestimate the true H_2 fractions, and so this would tend to strengthen our main conclusion that molecules can form *in situ* within AGN-driven outflows.

Another uncertainty in the local shielding approximation is that the column densities of individual species are calculated by multiplying the total hydrogen column density by the abundance of that species. However, this assumes that the abundances are uniform over the shielding length. In Fig. B1 we also compare the local approximation to the actual column densities for individual species that are important for shielding. For most species, the scatter is similar to that for the total hydrogen column density. The H_2 column density does show broader scatter, exceeding two orders of magnitude. However, most particles still lie within a factor of 10 of the true value. Also, the largest scatter in N_{H_2} is below the one-to-one relation, so the true values, and hence the true H_2 fractions, would tend to be higher.

Safranek-Shrader et al. (2017) compared several local shielding approximations to a full radiative transfer calculation in simulations of a section of a galactic disc in post-processing. They showed that a density gradient-based approach, as used in our work, performs poorly in terms of reproducing the correct column densities and equilibrium molecular fractions. Of the methods that use only local quantities (i.e. excluding their six-ray approximation), they found

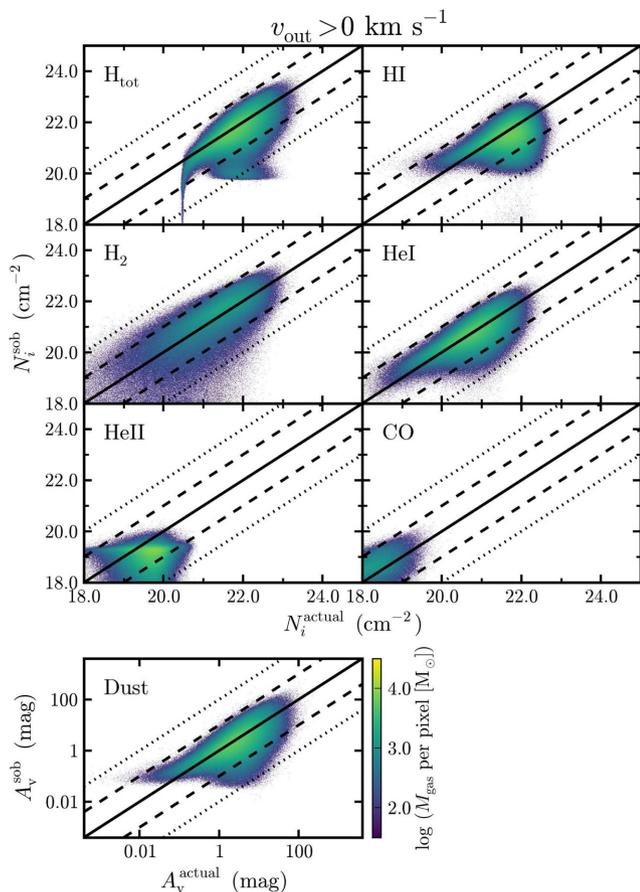


Figure B1. Column densities calculated using the Sobolev-like approximation (sob) and from ray tracing in post-processing (actual) in the simulation nH10_L46_Z1 after 1 Myr. The comparisons show the total hydrogen column density (top left), column densities of H I (top right), H₂ (second row, left), He I (second row, right), He II (third row, left), CO (third row, right), and dust extinction (bottom left). Only particles within the high-resolution wedge and outflowing with $v_{\text{out}} > 0 \text{ km s}^{-1}$ are included. Solid lines indicate the one-to-one relation, while dashed and dotted lines indicate factors of 10 and 100, respectively, above and below the one-to-one relation. While the Sobolev-like approximation follows the actual column density on average, there is a large scatter of at least an order of magnitude.

that using the Jeans length with a temperature capped at 40 K gave the best agreement with the radiative transfer results.

In the left-hand panel of Fig. B2, we compare the total hydrogen column density using the Jeans length in simulation nH10_L46_Z1

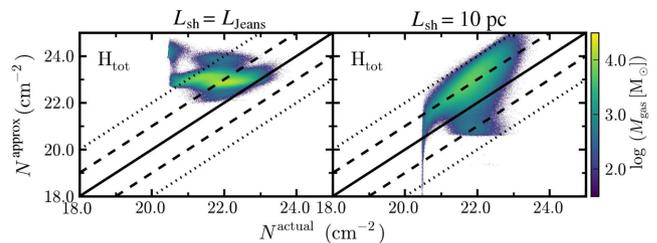


Figure B2. Total hydrogen column densities in simulation nH10_L46_Z1 calculated using the Jeans length (left-hand panel) and using a constant shielding length of 10 pc (right-hand panel), versus the actual column density calculated from ray tracing in post-processing. We include particles within the high-resolution wedge that are outflowing with $v_{\text{out}} > 0 \text{ km s}^{-1}$. Both of these alternative approximations systematically overestimate the true column density by an order of magnitude or more.

to the true column density from ray tracing. We do not cap the temperature to 40 K as in Safrank-Shrader et al. (2017), because we saw in Fig. 5 that most of the molecular gas in our simulations is at temperatures of hundreds to thousands of K. When we use the Jeans length, the column density is systematically higher than from ray-tracing, typically by more than an order of magnitude. Comparing to the top left-hand panel of Fig. B1, we see that the density gradient-based method performs better than the Jeans length in our simulations. However, the physical conditions in these fast molecular outflows are very different from those in the galactic discs simulated by Safrank-Shrader et al. (2017). Therefore, this comparison does not contradict the results of Safrank-Shrader et al. (2017), as we are looking at a very different physical regime.

If we impose a temperature cap of 40 K (not shown), as in Safrank-Shrader et al. (2017), we find that the column densities are lower and are closer to those from ray tracing. However, the slope of the $N_{\text{Htot}}^{\text{approx}}$ versus $N_{\text{Htot}}^{\text{actual}}$ relation is still flatter than the one-to-one relation, with the approximation tending to overestimate N_{Htot} below $\sim 10^{22} \text{ cm}^{-2}$ by up to an order of magnitude, and underestimate it at higher column densities.

Finally, the right-hand panel of Fig. B2 shows the total hydrogen column density assuming a constant shielding length of 10 pc, compared to the actual column density. This is the approach used by Ferrara & Scannapieco (2016), as 10 pc is the Field length at the mean density of their simulations, below which thermal conduction will suppress local thermal instabilities (Field 1965). We see that this method systematically overestimates the column density, by an order of magnitude on average, although we do not include thermal conduction in our simulations.

This paper has been typeset from a $\text{\TeX}/\text{\LaTeX}$ file prepared by the author.