

Cosmological nanolensing by dense gas clouds

Artem V. Tuntsov[★], Mark A. Walker[†]

Manly Astrophysics, 15/41–42 East Esplanade, Manly, NSW 2095, Australia

8 April 2022

ABSTRACT

We study the influence of a cosmological population of dense gas clouds on distant sources, with emphasis on quasar optical variability. In addition to gravitational lensing such clouds affect flux measurements via refraction in the neutral gas and via dust extinction, leading to a variety of possible light curves even in the low optical depth limit. We classify and illustrate the types of light curves that can arise. For sources as large as quasars we show that gravitational lensing and extinction are the dominant effects, with gas refraction playing only a minor rôle. We find that clouds with mass $\sim 10^{-4.5 \pm 0.5} M_{\odot}$ can reproduce the observed distribution of quasar variation amplitudes, but only if such clouds make up a large fraction of the closure density. In that case there may also be substantial extinction of distant optical sources, which can in principle be constrained by data on “standard candles” such as type Ia supernovae. Unfortunately that extinction is essentially grey, even when the material opacity is strongly wavelength dependent, making it difficult to distinguish from the influence of the background geometry. We propose a novel statistical test of the origin of quasar variability, based on the angular structure of the variation timescale for a large number of quasars distributed all over the sky. If quasar variability is primarily due to nanolensing that angular structure is expected to include a quadrupole term of amplitude $\sim 5\%$, which ought to be measurable with future data from the *Gaia* mission.

Key words: ISM: general – dark matter – gravitational lensing: micro – dust, extinction – quasars: general – cosmology: observations

1 INTRODUCTION

There is now strong evidence for the existence of dark matter that far exceeds the amount of visible material on large scales (e.g., Trimble 1987; de Swart et al. 2017). Over the past three decades the idea that most of the dark matter is non-baryonic has become dominant, thanks mainly to the success of non-baryonic models in describing structure formation in the universe, as traced by galaxies and by the cosmic microwave background fluctuations (Wechsler & Tinker 2018; Planck Collaboration et al. 2020). Nevertheless the hypothesised “dark matter particle” remains elusive (Feng 2010; Boveia & Dogliani 2018) and the picture thus lacks its most fundamental confirmation.

Even from that conventional perspective, however, baryonic matter is not fully accounted for either. Models of Big Bang Nucleosynthesis require twice as much baryonic matter as observed, in order to match the light-element abundances seen in the visible baryons (Fukugita et al. 1998). The “missing baryons” may perhaps be present in the form of a hot, ionised intergalactic medium (Cen & Ostriker 1999), but the evidence to date has not yielded a precise measurement of the amount of this gas (e.g. Macquart et al. 2020). Less popular, but nevertheless interesting, is the idea that a lot of

baryons may instead be locked up in dense, self-gravitating molecular clouds (Pfenniger et al. 1994; de Paolis et al. 1995). Such a population of undetected clouds would serve as a reservoir of material for ongoing star formation in galaxies (Pfenniger et al. 1994), and may be able to account for some of the observed regularities in those star-forming galaxies (Walker 1999b).

Molecular clouds that are both cold and dense are very difficult to detect (e.g. Combes & Pfenniger 1997; Gerhard & Silk 1996). At low temperatures there is very little emission from molecular hydrogen itself, whereas the low vapour pressures of heavier, polar molecules such as CO limit their emissivity. Moreover, any microwave spectral lines that do contribute significantly to the emissivity of the gas are likely to be transitions into the ground state, which tend to be very optically thick – because of the low temperature, high density and small Doppler broadening – so radiation can escape a cloud only in the wings of those lines. Consequently we expect this material to fall into the category of “CO-faint” molecular gas (Bolatto et al. 2013).

If such clouds were very dense indeed they could be compact enough to act as gravitational microlenses for sources in the Galactic Bulge or Magellanic Clouds (Henriksen & Widrow 1995); but in these settings it is also possible that they could be detected as a result of gas lensing — i.e. refraction in the molecular gas itself (Draine 1998). If the gas is assumed to be completely transparent the Galactic microlensing data are able to exclude a substantial pop-

[★] E-mail: artem.tuntsov@manlyastrophysics.org

[†] E-mail: mark.walker@manlyastrophysics.org

ulation in certain regions of the cloud mass-radius plane, but some mass-radius combinations remain largely unconstrained (Rafikov & Draine 2001), depending on the assumed equation of state.

At the other extreme one can consider gas clouds that are totally opaque, but which have no lensing effect whatsoever. Drake & Cook (2003), building on the predictions of Kerins et al. (2002), used data from the MACHO experiment (Alcock et al. 2000) to search for extinction events due to Galactic clouds. They found no clear examples of extinction events and were thus able to exclude a dynamically significant, opaque population of clouds at the upper end of the planetary mass range. Opaque clouds with low masses ($\lesssim 10^{-4} M_{\odot}$) yield only short-lived extinction events and are not strongly constrained by the Drake & Cook (2003) analysis.

More generally we expect the hypothesised cloud population to cause both refraction and extinction, but at present there are no reliable predictions of the form of the resulting light-curves. Walker & Wardle (2019) presented detailed models of the internal structure of cold, dense clouds in which molecular hydrogen is condensing into solid or liquid form. Although detailed, those models are not yet accurate enough to allow us to construct reliable lightcurve templates. And without a suitable template the Galactic microlensing data of experiments such as OGLE (Udalski et al. 2000) cannot constrain the hypothesised population — we do not know what we are looking for. Furthermore the fields targeted by the microlensing experiments tend to be chosen for their low mean extinction, and are therefore biased against finding dusty objects. This is in stark contrast to the situation with a population of primordial black holes, which are accurately modelled as point-mass gravitational lenses and whose abundance in the planetary mass range can be strongly constrained by the Galactic microlensing experiments (e.g. Carr & Kühnel 2020). We note, however, that the constraints do relax somewhat if the mass spectrum is broad Carr et al. (2021).

If there is a wider cosmological distribution of dark gas clouds they would lead to variations in the received flux of distant, compact sources such as quasars. In this arena Hawkins (1993, 1996) has long maintained that the observed variability of field quasars is best explained by a cosmologically significant population of gravitational microlenses, initially suggested to be¹ primordial black holes of mass $\sim 10^{-3} M_{\odot}$. And in a similar vein Schild (1996) inferred the presence of large numbers of $\sim 10^{-5} M_{\odot}$ objects, which he dubbed “rogue planets”, in the halo of the galaxy responsible for gravitationally lensing the multiply-imaged quasar Q0957+561.

By contrast, most of the recent work on multiply imaged quasars (e.g. Jiménez-Vicente et al. 2015) has interpreted differences amongst the light curves of the various images as due to microlensing by stars in the lens galaxy, with the dark matter assumed to be smoothly distributed on the relevant angular scales — consistent with the dark matter being elementary particles rather than macroscopic bodies. Certainly the lens galaxies do contain stars, and those stars will cause microlensing, but it has been pointed out by Hawkins (2020a,b) that the surface density of stars (at the

macro-image locations) may in some cases be too low to account for the observed variations. Furthermore, if most of the microlenses are of stellar mass then the implied quasar sizes (radii) are larger than expected on the basis of accretion disk models (e.g. Shakura & Sunyaev 1973). Source size is a key input parameter for modelling microlensing, and in Appendix A we give an overview of the available constraints on quasar sizes.

A clear theoretical perspective on the possibility of widespread microlenses was established long ago by Schneider (1993), who considered gravitational microlensing of quasars by point-mass objects randomly distributed between the Earth and the quasar. Schneider’s results demonstrated that the observed variability of field quasars was not well explained by a substantial cosmological density of massive lenses ($M \gtrsim 10^{-3} M_{\odot}$), but left open the possibility of a microlensing interpretation utilising a large population of low-mass lenses ($M \lesssim 10^{-5} M_{\odot}$). These scenarios were revisited by Zackrisson & Bergvall (2003), who showed that the constraints do depend on the assumed cosmology and source properties — source size in particular (again, please see Appendix A of this paper for an overview of quasar size constraints). Zackrisson et al. (2003) concluded that gravitational lensing by planetary mass objects could indeed reproduce a variety of quasar light curve statistics, but that it could not be the whole story and there had to be another, presumably intrinsic, variation mechanism.

In the present paper we study how quasar light curves are affected by a cosmological population of dense, molecular gas clouds of sub-stellar mass. Whereas stellar mass lenses at cosmological distances yield Einstein ring radii in the micro-arcsecond range, the much lower mass lenses that we consider yield structure on the nano-arcsecond scale; thus instead of “microlenses” we are dealing with “nanolenses”. Section 2 identifies the important scales that determine the qualitative character of the combined effect of gravitational lensing, gas refraction and extinction, and demonstrates how to classify the resulting light curves. For quasar nanolensing we find that it is usually either gravitational lensing or extinction that dominates what is seen, with gas lensing playing a lesser rôle. In Section 3 we compare our model statistics for source variability amplitude to data on quasar variability, finding that clouds with mass $\sim 10^{-4.5 \pm 0.5} M_{\odot}$ can provide a good fit to the observed distribution, providing they comprise a suitably large mean cosmological density, but our analysis is unable to constrain the cloud parameter space much further. In Section 4 we turn our attention to the systematic influence of extinction on distant sources, demonstrating that the effect is essentially grey, and can be significant for sources at redshifts $z \gtrsim 1$. Grey extinction will be difficult to identify in practice, but in Section 5 we present a new and purely kinematic effect — quadrupolar modulation of variation timescale across the sky — which will permit a test of the extrinsic origin of quasar variability. Busy readers can skip straight to Section 6 where our results are summarised.

2 PHENOMENOLOGY OF GAS CLOUD LENSING

2.1 Gas clouds: gravitating, refractive and opaque

There are three distinct physical effects that arise when radiation propagates through a population of molecular gas clouds: in addition to acting as a gravitational lens, a compact molecular cloud affects radiation via refraction in the neutral gas, and through extinction.

¹ In later papers, e.g. Hawkins (2011), the same author suggested a lensing interpretation relying on stellar mass black holes instead. Stellar mass lenses were demonstrated to be a poor match to the data by the analysis of Schneider (1993), although those constraints do depend on the assumed model parameters, particularly the source size Zackrisson & Bergvall (2003). We discuss quasar size estimates in an Appendix to this paper. Stellar mass lenses would also strongly magnify some type Ia supernovae at redshifts $z \sim 1$ (Metcalfe & Silk 1999), and the lack of any such examples in existing data constrains the cosmological density of stellar mass lenses to be $\Omega < 0.1$ (Zumalacárregui & Seljak 2018).

2.1.1 Refraction

Just like its gravitational counterpart, in the case of small deflections, as relevant for astrophysical applications, gas lensing can be described in terms of the phase delay imparted to each ray. That phase can be computed by integrating along the undeflected path:

$$\Phi_g(\mathbf{r}) = \int dz k [n(\mathbf{r}, z) - 1]. \quad (1)$$

Here k is the wavevector and $n(\mathbf{r}, z)$ is the refractive index at the position z along, and \mathbf{r} across, a fiducial optical axis. The refractive index is just a linear function of the gas density, ρ , with $n - 1 = a_\lambda \rho$ (Draine 1998); we refer to the constant of proportionality, a_λ , as the specific refractivity.

It is possible for the gas to have a non-uniform composition. That is the case, for example, in the models of Walker & Warde (2019), where precipitation of molecular hydrogen condensates makes the central regions rich in H_2 , while the surface layers are almost pure helium. Such configurations have a specific refractivity that is not uniform. For the sake of simplicity in this paper we will consider only clouds of uniform composition, so that the phase delay is proportional to the column density $\Sigma(\mathbf{r})$,

$$\Phi_g(\mathbf{r}) = k a_\lambda \Sigma(\mathbf{r}). \quad (2)$$

For a mix of molecular hydrogen and helium, with a ‘primordial’ He fraction of 24 per cent by mass, Draine (1998) gives a value

$$a_\lambda \approx 1.23 \text{ cm}^3 \text{ g}^{-1}, \quad (3)$$

with less than a 3 percent variation across the wavelength range from 440 nm to 670 nm.

In application across the full electromagnetic spectrum a much larger variation in a_λ would of course be encountered. At very long wavelengths the refractivity asymptotes to a value ($1.18 \text{ cm}^3 \text{ g}^{-1}$) only a few percent smaller than in the optical. Moving to shorter wavelengths the refractivity rapidly increases as the Lyman and Werner bands of H_2 are approached, where it fluctuates rapidly atop a systematic decline; it is negligible in the X-ray band and above.

Note that gas lensing is local – i.e., light rays are only affected by the matter they pass through. In contrast, the delay from gravitational lensing is due to a long-range force and can be related to the column density via a Poisson equation (e.g. Schneider et al. 1992)

$$\nabla_{\mathbf{r}}^2 \Phi_l(\mathbf{r}) = \frac{8\pi G k}{c^2} \Sigma(\mathbf{r}). \quad (4)$$

Appropriate boundary conditions ensure that a homogeneous mass distribution, being absorbed into the definition of angular diameter distances, does not contribute to light ray deflection.

The total phase delay, in the limit of weak gravity and small deflections that we consider, is the sum of the gaseous and gravitational contributions,

$$\Phi(\mathbf{r}) = \Phi_g(\mathbf{r}) + \Phi_l(\mathbf{r}). \quad (5)$$

We restrict ourselves to geometric (ray) optics, as appropriate in the optical, and assume that the lensing matter is concentrated in a region that is small along the line of sight compared to the angular diameter distances to the source D_s , to the lens D_l and from the lens to the source D_{ls} . In this case both gas and gravitational lensing act to bend the light rays by an angle

$$\hat{\alpha}(\mathbf{r}) = k^{-1} \nabla_{\mathbf{r}} \Phi(\mathbf{r}) \quad (6)$$

and the observed position θ of the image is related to the true position β of the source via the lens equation

$$\beta = \theta - \alpha(\theta). \quad (7)$$

The scaled deflection angle as a function of image position is related to the physical deflection on the lens plane (6) via

$$\alpha(\theta) = \frac{D_{ls}}{D_s} \hat{\alpha}(D_l \theta). \quad (8)$$

Apart from extinction, which we turn to in a moment, and redshift, which is constant on the timescales of interest to us, intensity is conserved along rays, hence variations in the observed flux are purely due to changes in the solid angle of the image compared to that of the source. Locally – i.e., for a small source – the magnification is thus given by the Jacobian of the lens mapping:

$$\mu = \det \partial \theta_i / \partial \beta_j = \frac{1}{(1 - \kappa)^2 - \gamma^2} \quad (9)$$

where convergence, κ , and shear, γ , have been introduced:

$$\kappa = \frac{D_{ls} D_l}{2k D_s} \left(\partial_x^2 \Phi + \partial_y^2 \Phi \right), \quad (10)$$

$$\gamma = \frac{D_{ls} D_l}{2k D_s} \sqrt{\left(\partial_x^2 \Phi - \partial_y^2 \Phi \right)^2 + 4 \left(\partial_{xy}^2 \Phi \right)^2}. \quad (11)$$

The convergence and shear define eigenvalues $1 - \kappa \mp \gamma$ of the transformation matrix $\partial \beta_j / \partial \theta_i$ that are independent of the coordinate system. The convergence due to the combination of gas and gravitational lensing is simply the sum of the two individual convergences and, remarkably, it is local. The local character of the gas refraction is obvious; in the case of the gravitational contribution it follows from the Poisson equation and we have

$$\kappa = \left(1 + \ell_\lambda^2 \nabla^2 \right) \frac{\Sigma}{\Sigma_{cr}}. \quad (12)$$

In this result for the beam convergence, which is dimensionless, there are two characteristic scales. The first is the familiar one for gravitational lensing, namely the critical surface density

$$\Sigma_{cr} = \frac{c^2 D_s}{4\pi G D_l D_{ls}} \approx 3.475 \times 10^2 \text{ g cm}^{-2} \left(\frac{D_l D_{ls}}{1 \text{ Mpc}} \right)^{-1}, \quad (13)$$

which a gravitational lens must exceed in order to produce multiple images. In the absence of any gas lensing ($\ell_\lambda = 0$) the beam convergence is just that due to gravitational lensing and its value is simply the value of the column-density, Σ , relative to the critical value. The second characteristic scale in equation (12) is a length scale: the gas lensing curvature radius, ℓ_λ , given by

$$\ell_\lambda^2 = \frac{a_\lambda c^2}{4\pi G} \approx \left(3.6 \times 10^{13} \text{ cm} \right)^2. \quad (14)$$

Here the numerical estimate corresponds to the particular value of specific refractivity given in equation (3). This length scale plays a key rôle because it determines the strength of the gas lens relative to that of the gravitational lens. If the cloud has a column-density curvature radius that is small (large) compared to ℓ_λ then the beam convergence due to gas lensing will be large (small) compared to that due to gravitational lensing.

Magnification of a source of uniform brightness is, as usual, just the ratio of the total solid angle of all images to that of the unlensed source. Given that gas lensing is local, it cannot lead to the formation of images outside the radius of the cloud, and it can only lead to significant magnification for sources whose angular size is much smaller than that of the cloud. In instances where gravitational refraction leads to image formation beyond the limb of the cloud the properties of those images are unaffected by gas lensing.

Provided that the cloud surface-density doesn’t greatly exceed

the mean column, $\langle \Sigma \rangle \equiv M/\pi R^2$ for cloud mass M and radius R , equation (12) allows us to estimate the importance of gas lensing and gravitational lensing, as follows. The Einstein ring radius, R_E , which is the radius at which the images form for a pure, strong gravitational lens with a source on-axis, is given by $\pi R_E^2 = M/\Sigma_{cr}$. Thus with $\Sigma \sim \langle \Sigma \rangle$ the gravitational contribution to equation (12) – i.e. the first term in parentheses on the right-hand side – is $\sim (R_E/R)^2$. The contribution of gas lensing, on the other hand – i.e. the second term in parentheses on the right-hand side – is $\sim (\ell_\lambda R_E)^2/R^4$. Thus we see that, just as R_E is a critical cloud radius for gravitational lensing, the critical radius for gas lensing is the geometric mean of R_E and ℓ_λ .

The existence of a critical column-density for gravitational lensing, Σ_{cr} (equation 13), and the curvature scale for gas lensing, ℓ_λ (equation 14), together imply a fiducial cloud mass, $M_* \equiv \pi \ell_\lambda^2 \Sigma_{cr}$. And the point (M_*, ℓ_λ) in the mass-radius plane lies at the intersection of the boundaries of three different regimes of lens behaviour, as shown in figure 1. In that figure the upper boundary of the region of strong gravitational lensing is simply a line of constant column-density – the critical surface density, Σ_{cr} , given in equation (13) – i.e. $(M/M_*) = (R/\ell_\lambda)^2$. Equivalently we can describe the upper boundary of strong gravitational lensing by the condition $R = R_E$ (the Einstein ring radius). As noted above, the corresponding critical radius for strong gas lensing is $R = \sqrt{R_E \ell_\lambda}$, and that boundary can be written in the alternative form $(M/M_*) = (R/\ell_\lambda)^4$. We note that the boundary of the region of strong gas lensing is only shown for $R < \ell_\lambda$ in figure 1, for the following reason. Clouds that are larger than ℓ_λ exhibit more gravitational refraction than gas refraction, so if gas lensing is strong and $R > \ell_\lambda$ then gravitational lensing is guaranteed to be even stronger and the main images will form outside the cloud (where gas refraction is irrelevant).

Because the critical radius for gas lensing can be much larger than the Einstein ring radius, if $\ell_\lambda \gg R_E$, it follows that for some types of clouds the optical depth to gas lensing can be much larger than the optical depth to gravitational lensing – by a factor $\ell_\lambda/R_E \gg 1$ – with correspondingly high event rates predicted for microlensing of stars in the Galaxy (Draine 1998; Rafikov & Draine 2001).

The fiducial mass just discussed depends on line-of-sight distances, but is independent of the properties of the source. However, in order for a lens to be able to substantially magnify a source we require that the lens must be larger than the (projected size of the) source. Therefore, for typical lenses with $D_l \sim D_{ls} \sim D_s/2$, the gas lensing curvature scale itself also demarcates two distinct regimes of behaviour: for large sources with $R_s \gg 2\ell_\lambda$ it is *only* possible to obtain substantial magnification with gravitational lensing, and then only with clouds that are more massive than the fiducial value M_* . The two types of sources that we are concerned with here – quasars and supernovae – both have radii $R_s \gg 2\ell_\lambda$, and it is for that reason that gas lensing is of little consequence in this paper.

We have already noted the wavelength dependence inherent in neutral gas refraction, but in this paper we will not explore that dimension. Instead we will concentrate on describing the behaviour in the optical (specifically the V-band), supplemented by some brief commentary on how the situation changes across the broader spectrum. Even with that restriction we demonstrate a wide variety of behaviour in the model. We note in passing that measurements made in V-band would correspond to shorter wavelengths – hence a larger refractivity – at the location of a cosmologically distant cloud where the lensing actually occurs. A similar point applies to the extinction introduced by a cloud.

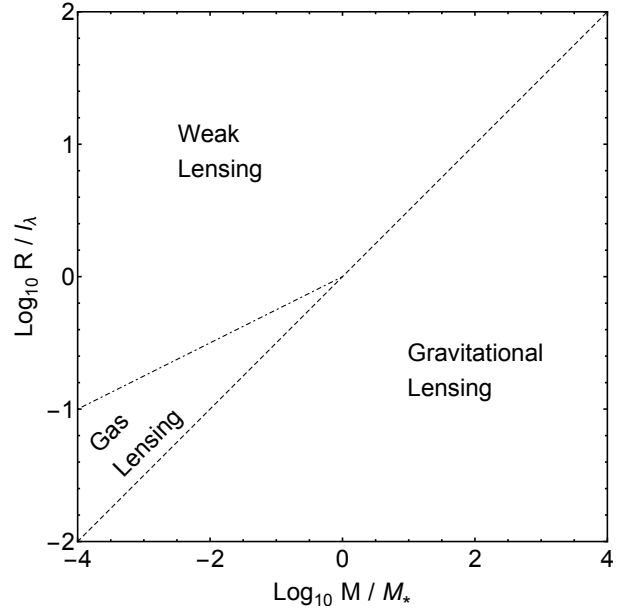


Figure 1. Classification of lens types according to their position in the mass-radius plane. In this plot radii are in units of the gas-lensing curvature scale, ℓ_λ , and masses are in units of the fiducial value $M_* = \pi \ell_\lambda^2 \Sigma_{cr}$, for which the Einstein radius $R_E = \ell_\lambda$. For any given optical configuration the transition to a strong gravitational lens corresponds to a particular surface density, Σ_{cr} , for which the critical radius is $R = R_E \propto \sqrt{M}$. The critical radius for gas lensing is $\sqrt{R_E \ell_\lambda} \propto M^{1/4}$, and is only important for $M < M_*$.

2.1.2 Extinction

We expect there to be two distinct contributions to the extinction of light: absorption and scattering by the individual atoms and molecules that make up the gas, and the extinction due to any dust particles that are present. Helium atoms and hydrogen molecules both respond strongly in the far-UV, and only comparatively weakly elsewhere. Below that band the larger contribution by far comes from H_2 . The Rayleigh scattering cross-section for H_2 is a strong function of wavelength (λ^{-4}), and in the V-band it contributes an opacity of (Dalgarno & Williams 1962) $2.1 \times 10^{-4} \text{ cm}^2 \text{ g}^{-1}$ (after allowance is made for the mass contribution of the helium). In the optical the hydrogen molecule has some weak (quadrupole) rovibrational absorption lines, but in any broad-band filter the absorption would be dominated by the low-frequency wings of the far-UV resonance lines. This absorption can be modelled as a sum of Lorentz oscillators, as described by Kettwich et al. (2015). With that model the expected absorption opacity is $1.1 \times 10^{-3} \text{ cm}^2 \text{ g}^{-1}$, in V-band, varying approximately as λ^{-2} . The resulting optical depth to extinction, τ_λ , can be written as

$$\tau_\lambda(\mathbf{r}) = \kappa_\lambda \Sigma(\mathbf{r}), \quad (15)$$

where $\Sigma(\mathbf{r})$ is the total gas column, and κ_λ is the total opacity (absorption plus scattering) of the cloud material, which we will often quote the inverse of, κ_λ^{-1} for ease of comparison with the other column density scales of the problem; as an inverse of the opacity, this will be referred to as a material transparency. The molecular contribution to opacity is unavoidable and implies an upper limit on κ_λ^{-1} of 768 g cm^{-2} in V-band.

We have already noted that in the far-UV the atomic and molecular transparencies are very low. At much shorter wavelengths still the absorption becomes negligible and one is left with only the electron scattering (Thomson), which amounts to $\kappa_{\lambda,T}^{-1} \approx 2.9 \text{ g cm}^{-2}$ for

gas which is 76% hydrogen and 24% helium, by mass. At the other extreme of wavelengths that are much longer than the optical the transparency increases rapidly; opacity is negligible in the radio.

The total opacity in the optical could be much higher than the molecular value, as a result of extinction by dust, but the contribution of the dust is uncertain even to order of magnitude. Consequently there are combinations of cloud mass and radius for which some physically plausible models would imply that essentially all the incident light is absorbed, and other models in which essentially all incident light is transmitted. To deal with the uncertainty introduced by the unknown dust opacity we proceed by bracketing the range of possibilities; we do this in two distinct, alternate ways — depending on which seems most appropriate. When considering clouds whose mass and radius are specified only in terms of M_* and ℓ_λ , respectively, the column-density in g cm^{-2} , say, is not specified and so it is inappropriate to work with specific values of the opacity. In that case we specify the central optical depth to extinction and we consider the limiting cases of totally transparent and totally opaque clouds, as shown in figure 3. On the other hand, when considering clouds in a specific context we bracket the range of possible opacities in the manner described below.

For simplicity, we assume that equation (15) also describes the extinction due to dust particles, with a single, suitably chosen value of κ_λ that is constant throughout each cloud and across the whole cloud population. In the local ISM the coefficient of proportionality between total hydrogen column and V-band extinction is measured to be $N_{\text{H}}/A_{\text{V}} \approx 1.9 \times 10^{21} \text{ cm}^{-2} \text{ mag}^{-1}$ (Bohlin et al. 1978), which corresponds to

$$\kappa_{\lambda=\text{V}}^{-1} \approx 3.6 \times 10^{-3} \text{ g cm}^{-2}. \quad (16)$$

The aforementioned numbers apply to the local, diffuse ISM and might not be representative of the material in cold, dense, self-gravitating gas clouds. In particular we note that self-gravitating clouds could in principle be much more transparent than the local diffuse ISM because any dust particles made of refractory elements could have sedimented into a dense core (Draine 1998). Even in that case, however, hydrogen condensates may form if the gas is both cold and dense (Pfenniger & Combes 1994). Condensed hydrogen has very little absorption in the optical band (Kettwich et al. 2015), but small particles of hydrogen “dust” would scatter optical light and thus cause extinction. To gauge the level of extinction that might arise from hydrogen condensates we turn to the hydrogen “snow cloud” models of Walker & Wardle (2019). In those models the “snowflake” content of the gas is highly non-uniform, being zero in the cloud core — which is too warm to permit hydrogen condensation — and near zero in the outer layers, which are almost pure helium, but substantial at intermediate radii. Nevertheless we can see from panel (c) of figure 3 in Walker & Wardle (2019) that, for the particular cloud shown there, the column density in hydrogen snowflakes ought to be roughly a fraction $\sim 10^{-3}$ of the total column, for impact parameters below 2.4 AU. The column-density of a single, micron-sized particle of solid hydrogen is $\sim 10^{-5} \text{ g cm}^{-2}$, so if we assume all the condensate to be in that form then we arrive at a fiducial $\kappa_\lambda^{-1} \sim 10^{-2} \text{ g cm}^{-2}$ — an estimate that is not far from the value appropriate to the local ISM.

Unfortunately Walker & Wardle (2019) did not attempt to model the size of the condensed particles so we do not have a value for their characteristic column-density. An upper limit can be estimated by noting that particles with very large column-densities are removed from the atmosphere on a dynamical timescale, because gravity is then the dominant force, whereas small particles can be kept in suspension by the ram pressure associated with con-

vective updraughts. If we assume that the latter have a speed that is of order one tenth of the sound-speed in the gas, balancing ram pressure with gravity at the inner-edge of the condensing envelope of the cloud in figure 3 of Walker & Wardle (2019), leads to a characteristic column-density $\sim 1 \text{ g cm}^{-2}$ for the condensed particle — corresponding to a sphere of radius 10 cm. It is therefore possible that the transparency due to snowflakes could be as high as $\kappa_\lambda^{-1} \sim 10^3 \text{ g cm}^{-2}$, which is comparable to the molecular contribution.

Given the large uncertainties in the transparency that is due to dust, the best we can do is to bracket the opaque column within a broad range:

$$768 \text{ g cm}^{-2} \gtrsim \kappa_\lambda^{-1} \gtrsim 3.6 \times 10^{-3} \text{ g cm}^{-2} \quad (17)$$

(in V-band).

As with gas lensing, extinction is local so only images formed inside the cloud are affected: if gravitational refraction leads to image formation beyond the limb of the cloud then both gas lensing and extinction are irrelevant for that image. There is thus a regime where the effect of a compact molecular gas cloud resembles that of a point-mass gravitational lens, effecting a major simplification in the light curve structure (see §2.3). Conversely, if the clouds are sufficiently nearby that their column-density is below the critical surface density (13) then the dominant image can form within the projected radius of the cloud and it is possible for transient extinction events to arise. Notwithstanding this broad-brush picture, the detailed structure of the light curves is affected to some degree by both extinction and gas lensing, even for distant/compact lenses — through the extinction of sub-dominant images, for example.

2.2 Gaussian model cloud profile

We utilise an axisymmetric gaussian model column-density profile for the individual clouds, to illustrate lensing plus extinction effects. A gaussian profile is convenient because it is amenable to analytic treatment and consequently has been often used to model discrete clouds (e.g. Clegg et al. 1998; Walker 1999a). The model projected density is

$$\Sigma(\mathbf{r}) = \Sigma_0 \exp \left[- \left(\frac{r}{r_0} \right)^2 \right], \quad (18)$$

which has two free parameters, the central column density Σ_0 and lens scale r_0 .

This leads to the lens equation between the radial position of the source y and its image x relative to the centre of the lens:

$$y = x - s \left(\frac{1 - e^{-x^2}}{x} + 2gx e^{-x^2} \right), \quad (19)$$

with positive (negative) x being interpreted as images lying on the same (opposite) side of the optical axis as the source. Both x and y are quantities projected onto the lens plane, so that the effective size of the source is scaled by the angular diameter distance ratio D_l/D_s , and expressed in units of the lens scale r_0 (or its angular equivalent $\theta_0 = r_0/D_l$). The two free parameters have been recast as the central column density in critical units (13), and the (square of the) gas lensing curvature scale (14) in units of the lens radius:

$$s \equiv \frac{\Sigma_0}{\Sigma_{\text{cr}}}, \quad g \equiv \left(\frac{\ell_\lambda}{r_0} \right)^2 = \left(\frac{\ell_\lambda}{D_l \theta_0} \right)^2. \quad (20)$$

For $s(1 + 2g) > 1$ there are two critical curves and their corresponding caustics where the magnification factor diverges. One

is the tangential critical curve (Einstein ring) that corresponds to a non-trivial solution of the equation (19) at $y = 0$, hence the corresponding caustic lies on the optical axis. The other is the radial critical curve that is the position x where dy/dx vanishes, and the corresponding position y in the source plane is the radial caustic. Far from the lens, $y \gg \sqrt{s(1+2g)}$, the mapping is nearly the identity and there is only one solution, $x \approx y$. If the source moves towards the lens and crosses the radial caustic, two additional images appear either side of the radial critical curve. As the source is brought closer to the optical axis the two extra images move away from the radial critical curve, one towards the lens centre and the other towards the Einstein ring where it merges with the primary image as $y \rightarrow 0$. For $s(1+2g) < 1$ the mapping is one-to-one throughout.

Equation (19) has no close-form solution and we use Newton's method to solve it. In the case of the primary image we start our iteration from the point $x = y$. For any additional images, y is inside the radial caustic, we start the iteration from halfway between the radial critical curve and either the lens centre or the tangential critical curve. The positions of the caustic and the radial curves are not available as close-form expressions either and are likewise computed, for given g and s , using Newton's method.

Depending on the size of the source relative to the radial caustic (if present), and its position, all or a part of the source may be multiply imaged with the extra images possibly merging at the radial critical curve. We model the source as a disc of uniform brightness, so that the total flux can be readily computed once the limb of the source is mapped into the image plane (thus avoiding a two-dimensional averaging of the magnification). To proceed we break the source into a number of small, quasi-rectangular patches by chords that are radial with respect to the centre of the lens, map their end points, and count the total area of the resulting images. Extinction depends on the position of the image and, unlike gravitational and gas lensing, modifies the brightness. The introduction of extinction necessitates integration along each chord image, rather than just mapping the boundary positions, but in the particular case of a gaussian column-density profile the result can be expressed analytically in terms of the exponential integral.

2.3 Light-curve classification for point-like sources

If we want to know the details of the light-curves that a given gas cloud might generate then we need to undertake calculations specific to that case, but once the properties of the clouds are specified the qualitative nature of the light-curves can be anticipated without detailed calculations — as we now describe. We first consider the lensing behaviour and then turn to the influence of opacity.

2.3.1 Position in the mass-radius plane

Typically one is considering hypothetical cloud models that have a well-defined mass and radius, and thus can be placed in a diagram such as figure 1. The type of lensing light-curves that could result from occultation of a source is then determined by which sector the cloud resides in — although, of course, the classification is unambiguous only when the cloud is located far from the sector boundaries.

In the Weak Lensing domain there is no significant refraction, hence no noticeable change in flux due to lensing. In the Gravitational Lensing domain the influence of the cloud is similar to that of a point-like mass, and the light-curve will be close to a Paczyński curve (Paczynski 1986). In the Gas Lensing domain there is no

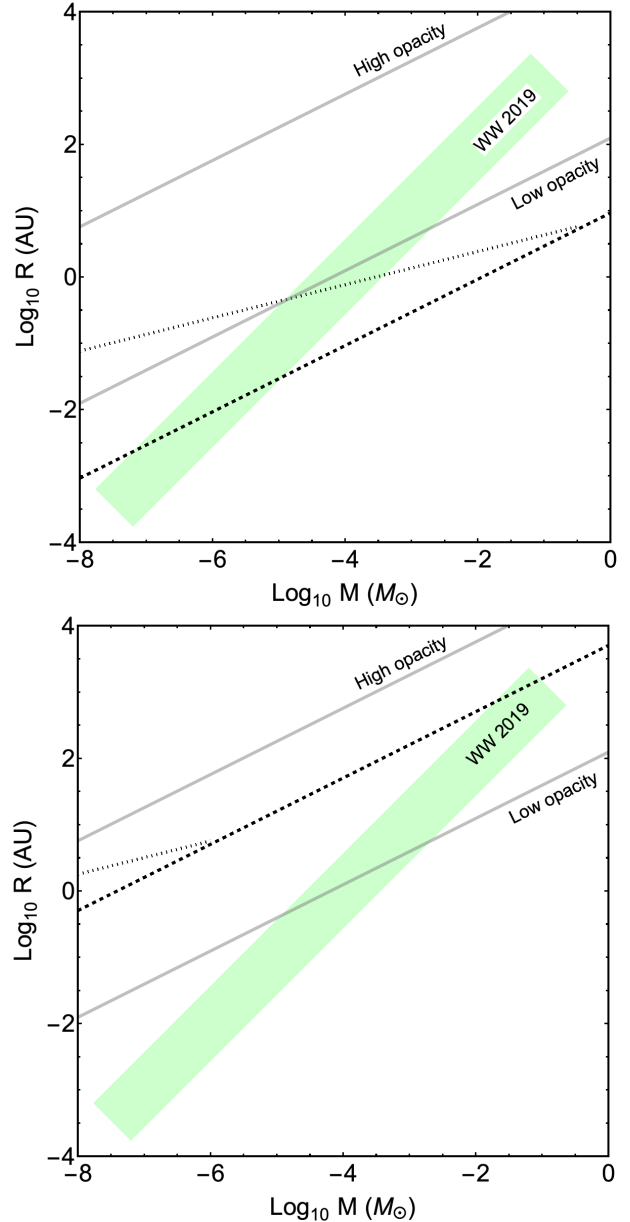


Figure 2. Partitioning of the mass-radius plane for molecular gas clouds according to their lensing and extinction characteristics. The upper (lower) panel is for Galactic (cosmological) clouds. The green shaded area roughly delineates the “snow-cloud” models shown in figure 4 of Walker & Wardle (2019). Dashed lines show the transition to a strong gravitational lens for $n = 3/2$ polytropes; dotted lines show the corresponding transition to strong gas lensing. The lower (upper), solid, grey line corresponds to a column $\Sigma = \kappa_{\lambda}^{-1}$ for the maximum (minimum) material transparency considered in §2.1.2. Clouds lying to the left (right) of the appropriate transparency line have a column-density below (above) the opaque column, κ_{λ}^{-1} , and thus are transparent (opaque).

unique functional form for the light-curves, but the generic structure of the flux maps is a demagnification region where the source is behind the limb of the cloud and a magnification region where it is on-axis. Examples of gas lensing behaviour can be found in the next section for the specific case of clouds having a gaussian column-density profile, and in the literature (Draine 1998; Rafikov & Draine 2001) for the case of polytropic clouds.

If nothing more is known about the cloud than its mass and radius then figure 1 itself should be used for classification, but we often have further information about the physics of the structures under consideration and in that case a clearer differentiation can be achieved by using a more precise construction. For example the gas cloud models shown in figure 4 of Walker & Wardle (2019) are not far from $n = 3/2$ polytropes, because they comprise a large core that is precisely polytropic and only a small, non-polytropic envelope. For any pure polytropic model the relationships between point-wise properties and their average values can be quantified, so we can determine values for both contributions to the beam convergence in equation (12) in terms of the cloud mass and radius.

The gravitational lensing beam convergence is simply Σ/Σ_{cr} , and for $n = 3/2$ polytropes the peak (central) value of the column-density is $\Sigma_{max} \approx 4.15 \langle \Sigma \rangle = 4.15 M / (\pi R^2)$; thus in this case the transition to a strong gravitational lens occurs at a cloud radius that is approximately twice as large as for gravitational lensing by the cloud as a whole. Polytropic clouds with $n > 3/2$ can show much larger differences in that respect, because their surface density profiles are increasingly centrally peaked as n increases.

Similar considerations apply to the onset of strong gas lensing. For polytropic clouds with $n = 1.5, 2.0, 2.5, \dots, 4.5$ that transition was quantified by Draine (1998) as a gas lens strength $S = S_c(n)$, i.e. a critical value of the strength parameter, depending on the polytropic index. That description can be easily related to our parameterisation, because Draine’s lens strength is $S = (3/4)(M/M_*)(\ell_\lambda/R)^4$. The critical value for $n = 3/2$ is $S_c(n = 3/2) \approx 0.026$ (Draine 1998), and decreasing rapidly with increasing n .

Analogous to figure 1, these estimates allow us to delineate domains in the mass-radius plane where the clouds can be classified as: weak lenses; strong gas lenses; or, strong gravitational lenses. Figure 2 shows that classification for $n = 3/2$ polytropes, for both Galactic and cosmologically-distant clouds, overlaid with the theoretical gas cloud models shown in figure 4 of Walker & Wardle (2019).

Also shown in figure 2 are the boundaries between opaque and transparent clouds, for the smallest and largest material transparencies decided on in §2.1.2. The effect of these boundaries must be considered in combination with the nature of the lensing, in order to anticipate the character of the light-curves that would result from an occultation, as follows. In the case of strong gravitational lensing the dominant images are formed beyond the limb of the cloud, where the optical depth is zero regardless of the cloud’s material opacity, and extinction can have only a minor influence on the light-curves.

If the cloud is not a strong gravitational lens then images will form inside the limb and will thus be subject to some extinction; but it need not be a large effect, even if much of the cloud is opaque. The key point to notice is that if the gas lensing is sufficiently strong then during an occultation the dominant images will form close to the limb of the cloud — because the refraction angle is zero at the limb, but increases rapidly as one moves inwards, so rays can only reach the observer if they pass near the limb. And near the limb the column of gas and dust is relatively low. Depending on the details — i.e. the material opacity, the column-density profile, and the strength of the gas lensing — the dominant images could then be subject to either high- or low-levels of extinction. Thus we see that the material opacity is guaranteed to have a major effect on the light-curve only if both gaseous and gravitational refraction are weak enough that the dominant images traverse the central, opaque region of the cloud. In that case, of course, occultations by opaque clouds will generate extinction events.

With those points in mind, by examining figure 2 we arrive at the following conclusions. If the material transparency is low then all of these models are totally opaque in V-band; and even if the material transparency is high then only the larger clouds are transparent. Given that our “high transparency” case corresponds to the absolute maximum that is physically possible (i.e. only the molecular contribution to the opacity), at the low mass end of the distribution the Walker & Wardle (2019) models are guaranteed to be opaque. Naturally the lensing characteristics differ greatly between Galactic and cosmological contexts. In the Galactic case the larger clouds are seen to form only weak lenses, so depending on the material transparency those clouds could result in occultations that can be classified as pure extinction events, or occultations with very little effect on the measured source flux. By contrast small clouds within the Galaxy are seen to act as strong gas lenses, so that significant magnification (and also demagnification) is expected — as per Draine (1998); Rafikov & Draine (2001). However, as already noted, these models must also be opaque and the resulting light-curves should manifest both extinction and gas lensing. Even in the Galactic context, gravitational lensing starts to become important for the smallest of the Walker & Wardle (2019) model clouds, pushing the dominant image beyond the limb of the cloud so that extinction plays a minor role in determining the appearance of occultations. Figure 2 also demonstrates that in the cosmological context essentially all of the models from figure 4 of Walker & Wardle (2019) are very strong gravitational lenses and thus would yield light-curves very similar to those of a point-mass gravitational lens (e.g. Paczynski 1986).

2.3.2 Example flux maps for gaussian lenses

In this paper we utilise a model with a gaussian column-density profile — because the simulations presented in the next section require rapid solution of the lens equation — which does not have a finite extent (radius). But it is straightforward to classify the regimes of lensing and extinction behaviour based on any two of the parameters $\{M, \Sigma_0, r_0\}$. In the previous section (§2.2) we used $\{\Sigma_0, r_0\}$ (equations 20) to classify the lensing behaviour, but in §3 it is more convenient to use $\{M, \Sigma_0\}$ instead. In terms of those parameters the criterion for a strong gravitational lens remains the same, namely $\Sigma_0 > \Sigma_{cr}$, whereas the criterion for a strong gas lens becomes $M/M_* < 2(\Sigma_0/\Sigma_{cr})^2$.

Any given cloud occulting a source can generate a variety of light-curves, depending on the impact parameter of the event (b) and the effective transverse velocity of the cloud relative to the source (v_\perp). For axisymmetric cloud models the effect of these kinematic parameters can be readily imagined because they describe the (projected) radial position of the source, y , as a function of time, t , via $y^2 = b^2 + v_\perp^2 t^2$. Given the appropriate flux map for the cloud — i.e. the source flux as a function of y — the entire family of possible light-curves can be anticipated. We note that in the absence of any extinction the flux map could alternatively be called the “magnification map”, which terminology is familiar in the context of gravitational lensing.

Figure 3 illustrates a variety of flux maps for point-like sources seen through a lens with a gaussian profile, depending on the strength of the gas lensing, gravitational lensing and extinction. Modulo the logarithmic scaling of the horizontal axes, these plots can be thought of as half of the light-curve (the egress) for events with zero impact parameter. The cloud mass is $10^{-4} M_*$ in all cases, with the central column varying by a factor of 10 between consecutive panels from $10^{-3} \Sigma_{cr}$ (top left) to $10^2 \Sigma_{cr}$ (bottom right).




Figure 3. The flux of a point-like source, normalised to the flux in the absence of an occulting cloud, as a function of (the Log of the) source position, y , from equation (19), expressed in units of r_0 . In all cases the cloud mass is $10^{-4}M_*$, but with different values of the central column-density as indicated in the lower right-hand corner of each panel. The solid, blue curve is appropriate to a completely transparent cloud, and the dashed, red curve is for a cloud with a central optical depth to extinction of 1,000. We note that the vertical line at $y \approx 100 r_0$ in the top-right panel is a real feature – it is the radial caustic of the lens mapping – but it is very narrow and would appear washed out in the case of non point-like sources.

And in each panel the map is shown for both $\kappa_\lambda \rightarrow 0$ (transparent case; blue curve) and a lens with $\kappa_\lambda^{-1} = 10^{-3}\Sigma_0$ (opaque case; red, dashed curve). For the lowest column-density shown the lensing is weak and thus in the transparent case there is almost no discernible effect on the source flux; in the opaque case occultations yield pure extinction events. At a column of $10^{-2}\Sigma_{cr}$ the cloud is already a strong gas lens, manifesting a radial caustic at $y \approx 0.3 r_0$ in the transparent case, whereas in the opaque case the map remains close to that of the opaque, weak lens. The region of demagnification at $y \sim r_0$ is noteworthy: gas lensing can produce demagnification, but gravitational lensing cannot. Increasing the column further to $0.1\Sigma_{cr}$ makes for a very strong gas lens with the radial caustic pushed out to $y \approx 10^2 r_0$; the map for the opaque case is still dominated by extinction, but the tangential caustic has a non-negligible influence near the origin. When the central column reaches Σ_{cr} the cloud is just at the transition to strong gravitational lensing, while the gas lensing is very strong indeed and is the dominant influence on the shape of the flux map. At this point the gas lensing is so strong that the dominant images are formed at large radii; there even the opaque cloud imposes very little extinction and so the flux map of the opaque cloud is very similar to that of the transparent cloud. As the central column increases further the differences between the opaque and transparent cases become too small to discern with the eye, as the dominant images form at still larger radii. At a central column of $10\Sigma_{cr}$ the flux map is that of a hybrid gas and gravitational lens; the net demagnification is now limited to modest

values because it is offset by the magnification due to the gravitational refraction. For the largest central column shown in figure 3 ($100\Sigma_{cr}$) the demagnification region has completely vanished and the flux map is very close to that of a point mass gravitational lens: the largest differences relative to a point mass are at the level of a few percent, manifest at $y \approx 20 r_0$.

All of the maps shown in figure 3 are for point-like sources and if we consider extended sources then there is substantial additional variety in the resulting flux maps. However, the effect of a non-zero source size is simply to convolve the (two-dimensional) point-source flux map with the intensity profile of the source, so the influence can be readily conceived as a simple smearing of the point-source flux map. For sufficiently large sources the effects of any one lens can be made arbitrarily small, but if there are many clouds occulting a source then it is their aggregate effect that is of interest — e.g. the total extinction, as discussed in §4.2.

3 COSMOLOGICAL NANOLENSING OF QUASARS

Hawkins (1993) suggested that gravitational microlensing by a cosmological population of black holes could explain the observed optical variability of quasars. That idea was explored theoretically by Schneider (1993), using simulations of microlensing by point-mass objects, with source redshifts and temporal sampling matched to the data underlying Hawkins & Veron (1993). By comparing the

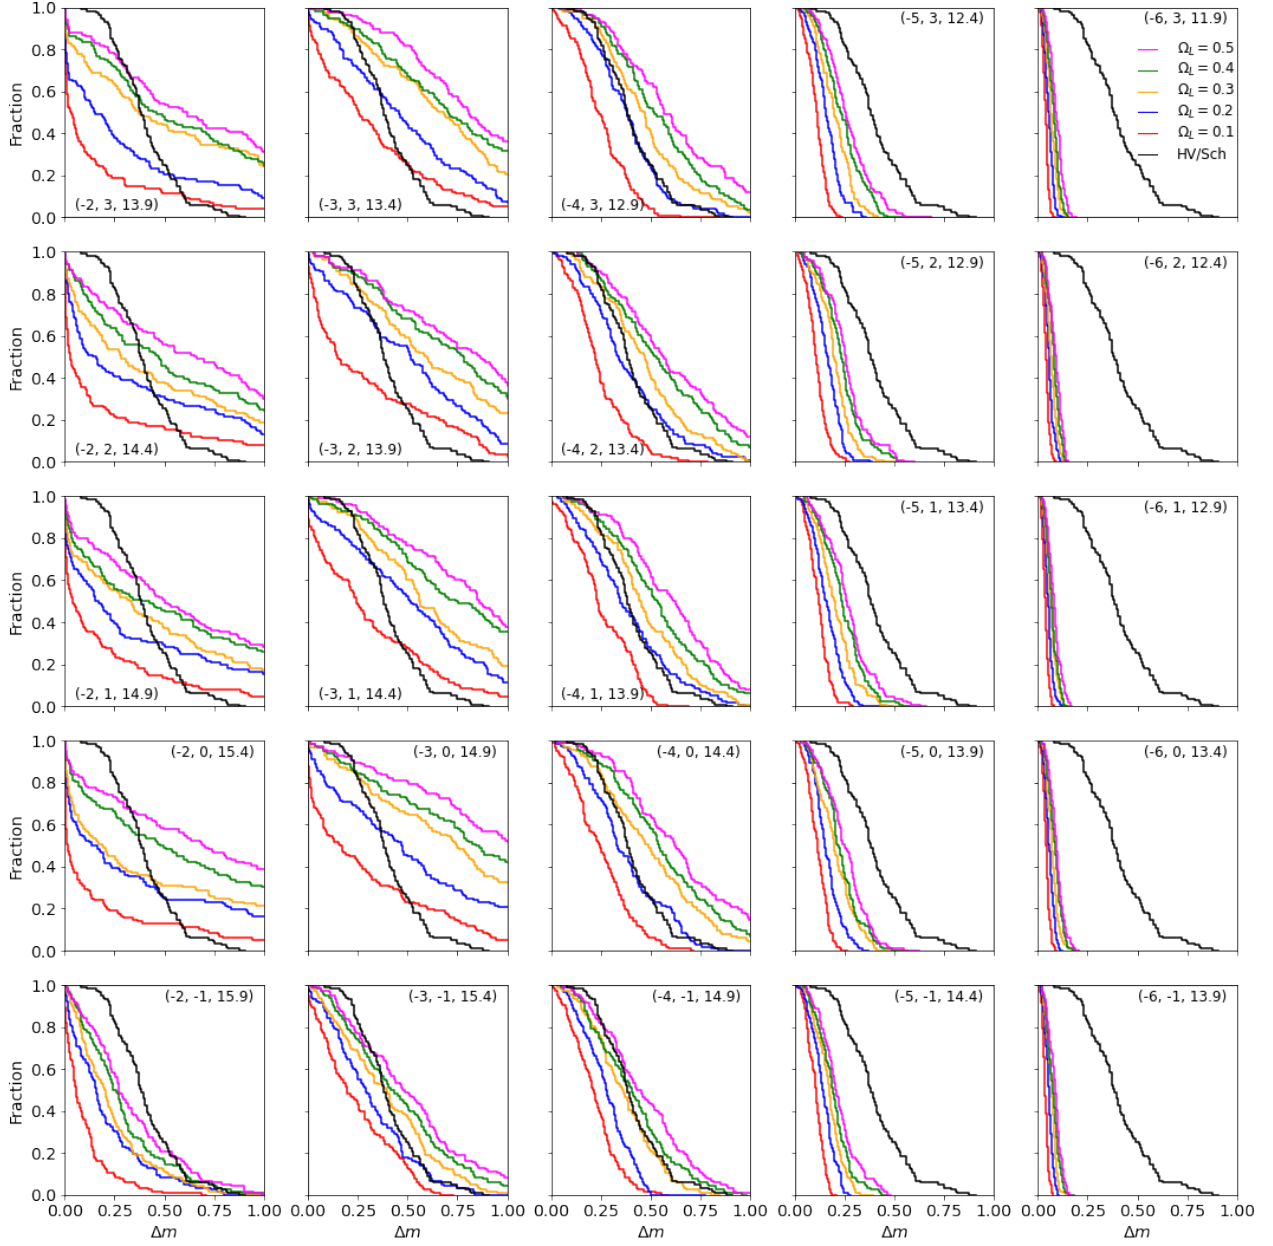


Figure 4. Fraction of the synthetic light curves changing by a given magnitude threshold for a selection of cloud models with gaussian column-density profiles, characterised by the parameters $(\log_{10} M, \log_{10} \Sigma_0, \log_{10} r_0)$ as indicated in each panel. We note that M and Σ_0 define the model, and an approximate value of $\log_{10} r_0$ (cm) is quoted in addition, for the convenience of readers. The layout of the grid is as follows: columns correspond to clouds of fixed mass, with $\log_{10} M (M_\odot) = \{-2, -3, -4, -5, -6\}$, left to right; rows correspond to clouds of fixed central column-density, with $\log_{10} \Sigma_0 (\text{g cm}^{-2}) = \{3, 2, 1, 0, -1\}$, top to bottom. Lines of different colours represent different values of the cosmological density Ω_l that the clouds make up, as per the legend in the top-right panel. We note that the two largest values, $\Omega_l = 0.4, 0.5$ are actually greater than the total matter content of the background cosmological model that we are using, and cannot be physically realised. For comparison, the distribution in the [Hawkins & Veron \(1993\)](#) sample, as read off Figure 5 in [Schneider \(1993\)](#) is shown in black in each panel. This figure is appropriate to clouds that are as transparent as possible, with a material transparency of $\kappa_\lambda^{-1} = 768 \text{ g cm}^{-2}$ (see §2.1.2).

variability amplitudes (i.e. full variation from minimum to maximum) in simulated light curves against the data, [Schneider \(1993\)](#) was able to rule out a large cosmological population of point-like lenses with masses in the range $10^{-3} M_\odot \lesssim M \lesssim 3 \times 10^{-2} M_\odot$, under certain assumptions about the underlying cosmological model and intrinsic quasar properties.

In this section we repeat the analysis of [Schneider \(1993\)](#), but for the case of non-point-like lenses. Specifically: we consider com-

pact gas clouds, and we include the effects of gas lensing and extinction. As our background geometry we adopt the Planck cosmological model ([Planck Collaboration et al. 2016](#)) for the computation of angular-diameter distances and volume elements. However, in other respects we follow in detail the approach described by [Schneider \(1993\)](#); in particular we adopt the magnification-multiplication approximation when calculating the total microlensing magnification due to multiple lenses along the line of sight. Although inadequate




Figure 5. Same as Figure 5, but with low material transparency, $\kappa_{\lambda}^{-1} \approx 3.6 \times 10^{-3} \text{ g cm}^{-2}$, as would be appropriate for the local, Galactic ISM.

to describe gravitational lensing in the case of near-unit optical depth, it is nevertheless suitable for computing the statistics of significant changes in the magnification on short time scales – years, say – which are dominated by changes in the relative position with respect to the lens closest to the line of sight. The multiplication approximation is a useful and physically motivated method to interpolate from one such lens to another.

In our simulations, clouds are distributed at random with a specified, uniform mean comoving number-density. We compute the magnification factor numerically for clouds that are closer than either 15 length scales (r_0) or 15 Einstein radii (whichever is the greater) to the nearest edge of the source. For all other clouds the magnification is very close to unity and very little error is incurred by utilising the analytic result for a point mass lens; we therefore

utilise that approximation for these clouds. For practical reasons noted in §2 we model the source as a disc of uniform brightness and assume a radius of $R_s = 3 \times 10^{15} \text{ cm}$, which is three times larger than used by Schneider (1993); this choice is discussed in Appendix A.

To generate light-curves, the sources are all assumed to move at 400 km s^{-1} and lenses are assigned random velocities drawn from a gaussian distribution with an r.m.s. of 400 km s^{-1} per component; the observer is stationary. The vast majority of our synthetic light curves resemble a series of Paczyński curves (Paczynski 1986), similar to point-mass gravitational lenses, and a minority of light curves manifest extinction events. Following Schneider (1993), we compare our simulated light curves to the data through the variability amplitude distribution, having matched the redshift distribution

and temporal sampling of our simulations to those of the [Hawkins & Veron \(1993\)](#) sample. We include all sources, regardless of the average or minimum magnification in their light curve, when deriving the amplitude distribution, which corresponds to the ‘volume-limited’ case (as opposed to ‘flux-limited’ one) in the calculations undertaken by [Schneider \(1993\)](#).

Figures 4 and 5 present the fraction of light curves with variability amplitude above a given threshold, for a variety of cloud models. Each panel in these figures corresponds to a particular combination of cloud mass and central column-density, with the various curves within each panel showing results for different values of the assumed cosmological density in clouds of that type, Ω_l . Unsurprisingly, figures 4 and 5 show that for each cloud model the variability amplitudes increase systematically with Ω_l . We note that our adopted cosmology has a total matter content that is less than the two largest values ($\Omega_l = 0.4, 0.5$) used for the simulations shown in these figures, so the corresponding universes cannot be physically realised. Figure 5 is appropriate to clouds having a transparency similar to that of the local Galactic ISM ([Bohlin et al. 1978](#)), while figure 4 is appropriate to much more transparent gas, where the material transparency is set at the upper limit set by molecular absorption (as discussed in §2.1.2).

Figure 5 demonstrates a strong change in behaviour, at fixed cloud mass, as the central column-density, Σ_0 , decreases. For high central columns the behaviour is essentially that of point-mass gravitational lenses, while for very low central columns the variability is dominated by extinction and exhibits very large amplitudes for clouds of all masses.

Much like [Schneider \(1993\)](#), our results readily lend themselves to excluding substantial populations of certain types of lenses. Indeed, as noted above, for the majority of our models – those with sufficiently high central columns $\Sigma_0 \gg \Sigma_{cr}$ – the effect of gas clouds is much the same as that of point-like lenses, and consequently our constraints are broadly similar to those of [Schneider \(1993\)](#); [Zackrisson & Bergvall \(2003\)](#); [Zackrisson et al. \(2003\)](#). In other words: a cosmologically significant population of dense gas clouds with masses $M \gg 10^{-4} M_\odot$ can be excluded because such a population would result in variation amplitudes that are much larger than observed. This incompatibility arises because high-mass gravitational lenses have Einstein radii that are larger than the sources, so high magnifications commonly occur.

Also in conflict with the data are low-transparency gas clouds with low central column-densities, $\Sigma_0 \lesssim 1 \text{ g cm}^{-2}$, for clouds of all masses within the range investigated (figure 5). For these clouds the very large variation amplitudes seen in our simulations are primarily due to extinction. We caution that these large extinction fluctuations are in part attributable to the very extended (infinite!) radius of the gaussian column-density profile that we are using as a model; and that aspect of the model is, of course, unphysical.

Less objectionable are the very low mass clouds, $M \ll 10^{-5} M_\odot$, which (excepting the low column-density, low transparency clouds noted above) all produce little variability. These low-mass lenses have Einstein ring radii that are much smaller than the source radii, so significant magnification does not occur. We cannot exclude a high abundance of such clouds. But on the other hand there doesn’t seem to be much motivation to consider such a population, given that they cannot explain the observed quasar variability.

At the interface between the benign-but-uninteresting low-mass clouds and the high-mass clouds, which produce too much variability, figures 4 and 5 show that there are some combinations of cloud mass, column-density and abundance for which our sim-

ulations produce variability amplitude distributions that are similar to the observed quasars. These parameter combinations are ‘preferred’ in the sense that they have the potential to explain the optical variability of quasars.

In determining the parameters of our preferred models it is important to bear in mind that the observed distribution is not corrected for measurement uncertainties. Nor are those uncertainties included in our simulations, because the information provided with the data is not sufficient to allow us to do that. Crudely speaking the effect of measurement errors of σ magnitudes is to put a floor of $\sim 2\sigma$ on the measured variability amplitude of all sources, and this will shift the observed probability distribution by the same amount. Examining the black curve that is plotted in each panel of figures 7 and 8 suggests that the photometric errors in the data might be as much as $\sigma \sim 0.05 \text{ mag}$ — a figure that is consistent with the estimates in [Hawkins \(1996\)](#). Making allowance for that possibility, our simulations would best match the data if $M \lesssim 10^{-4} M_\odot$ (with $\Omega_l \gtrsim 0.1$) and $M \gtrsim 10^{-5} M_\odot$ (with $\Omega_l \lesssim 1$). This preferred range applies to all of the central column-densities we investigated in the case of high transparency gas clouds (figure 4), and for central columns $\Sigma_0 \gtrsim 10 \text{ g cm}^{-2}$ in the case of low transparency material (figure 5).

4 EXTINCTION DUE TO DENSE CLOUDS

The previous section confined attention to the variability introduced by a population of dense gas clouds, but there is also an effect on the average flux of distant sources — even for dense clouds with Einstein radii in excess of their physical size. Parts of the image plane with strong extinction inevitably map onto some areas of the source plane, regardless of the effects of gravitational and gas lensing, and therefore the total flux is lower than it would have been if the lenses were completely transparent.²

4.1 The reddening curve

Reddening is a key signature of dust extinction. But for sufficiently dense and compact gas clouds, the extinction that is introduced is very nearly grey even when the specific opacity of the material is strongly wavelength dependent. The reason is that transmission through the central part of a highly opaque lens is effectively zero over a broad range of wavelengths. In this case the position of the outer edge of the opaque region changes with wavelength, but when most of the lens is opaque there is little room for this change.

We can define a wavelength-dependent opaque radius where the optical depth through the lens is unity $r_\lambda : \Sigma(r_\lambda) = \kappa_\lambda^{-1}$, which, for the gaussian lens of eq. (18) with $\Sigma_0 > \kappa_\lambda^{-1}$, is

$$r_\lambda = r_0 \sqrt{\log \kappa_\lambda \Sigma_0}. \quad (21)$$

The (image plane – please see below) effective optical depth to extinction of a *population* of such lenses with number density n and line-of-sight depth D is

$$\tau_\lambda^{pop} = nD \int dr 2\pi r \left[1 - e^{-\kappa_\lambda \Sigma(r)} \right] \approx \pi n D r_\lambda^2 \quad (22)$$

$$= \pi n D r_0^2 \log \kappa_\lambda \Sigma_0 \quad (\text{gaussian}) \quad (23)$$

where the approximation replaces the $[1 - e^{-\tau(r)}]$ factor with a step function at position r_λ as appropriate if τ increases rapidly around

² This effect is mentioned in footnote 2 of [Metcalf & Silk \(1999\)](#).




Figure 6. The reddening curve (black line) for diffuse gas with a total-to-selective extinction ratio $R_V = 3.1$, as appropriate to the local, diffuse ISM (Cardelli et al. 1989). Also shown are the reddening curves for high-column density gas clouds, with gaussian profiles and various values of Σ_0 (as per the legend), having the same material opacity as the local, diffuse ISM.

$\tau = 1$. In fact, for a gaussian lens the integral can be calculated exactly with $\log \kappa_\lambda \Sigma_0$ acquiring a correction of $\tilde{\gamma} - \text{Ei}(-\kappa_\lambda \Sigma_0)$. Either way this is a very slow function of κ_λ – and, thereby, λ – when Σ_0 is well above the material transparency κ_λ^{-1} . Figure 6 compares the familiar reddening curve for the local, diffuse Galactic ISM (Cardelli et al. 1989) with that of compact gas clouds of the same composition, for various assumed central column-densities; it is evident that the reddening effect of the dust is greatly reduced.

For completeness, in the cosmological case with a constant comoving number density of clouds $n_{\text{comov}} = \Omega_I \rho_{\text{cr}} / M$, the surface density nD would need to be replaced by an integral

$$nD \rightarrow N(z_s) = n_{\text{comov}} D_\tau(z_s) = \frac{3H_0^2 \Omega_I}{8\pi G M} D_\tau(z_s) \quad (24)$$

$$\text{with } D_\tau(z_s) \equiv \int_0^{z_s} \frac{d^3 V_{\text{comov}}(z)}{dz d^2 \Omega} \frac{dz}{D_A^2(z)}, \quad (25)$$

e.g. employing the $d^3 V_{\text{comov}}(z)/dz d^2 \Omega$ volume as conveniently realised in the Python programming language with the `differential_comoving_volume()` method of the `astropy.cosmology.FLRW` class (Astropy Collaboration et al. 2018).

An important qualitative point to note is that the gaussian column-density profile used to construct the curves in figure 6 extends to infinite radius. More realistic cloud models (e.g. Walker & Wardle 2019) have a finite extent, and with sufficiently high central optical depths such models may exhibit even less reddening than shown here.

4.2 Extinction estimate for large sources

The optical depth defined on the image plane, τ^{pop} , is an appropriate measure of extinction when gravitational and gas lensing is

weak but should be corrected (which we will mark by a hat, $\hat{\tau}^{\text{pop}}$) to account for the relative areas of differently extinguished portions of the image plane when mapped onto the source plane. In the case of the gas clouds considered in this paper this mapping is correlated with the degree of extinction, making correction non-trivial. Here it should be reiterated that the individual images have a very small angular separation and would be correspondingly difficult to resolve; in practice, then, only the total source flux is measured.

In the case of a sufficiently large source, with radius $R_s \gg R_E (\Sigma_{\text{cr}}/\Sigma_{\text{pop}})^{1/2}$, we expect there to be many lenses projected onto the source, resulting in a net magnification (relative to an empty beam) of $\langle \mu \rangle \simeq (1 - \Sigma_{\text{pop}}/\Sigma_{\text{cr}})^{-2}$, as would have been expected of a smooth matter distribution. The only difference with the latter is that there would be a dark spot at the location of each of the opaque clouds, and therefore the source plane effective optical depth to extinction is equal to its image plane, geometric value (22, 23),

$$\hat{\tau}_\lambda^{\text{pop}} \approx \tau_\lambda^{\text{pop}} \approx \left\langle \frac{\Sigma_{\text{pop}}}{\Sigma_{\text{cr}}} \frac{r_\lambda^2}{R_E^2} \right\rangle. \quad (26)$$

If the cloud abundance is sufficiently large the optical depth to extinction may be non-negligible.

4.3 Supernova dimming

The foregoing points are relevant to the interpretation of “standard candles” at cosmological distances, such as type Ia supernovae (e.g. Riess et al. 1998; Perlmutter et al. 1998). The low levels of reddening that are observed in distant SNe Ia (e.g. Aguirre 1999b; Goobar et al. 2002; Riess et al. 2004) provide tight constraints on the extinction contribution from dust associated with intervening diffuse gas if it is similar to the local, Galactic ISM. In the case of dense gas clouds, however, we have shown that the extinction is very nearly grey even if the clouds are made of material similar to the local ISM (figure 6), so that low levels of reddening do not exclude high levels of extinction. That raises the question of whether dense gas clouds might actually be having a significant effect on the shape of the magnitude-redshift relation for type Ia SNe? In this section we address that question.

Attempts to place constraints on the extinction contributed by dense clouds are hampered by a lack of knowledge about their properties. Some simple physical models of the internal structure of dense, Galactic clouds have recently been constructed (Walker & Wardle 2019), but as yet there are no reliable models for the formation and evolution of such objects. Some hypothetical evolutionary scenarios can be constrained using the supernovae photometry. For example: clouds which form at high redshift and do not subsequently evolve would yield a redshift-dependent extinction much like that of a smooth background distribution of dust with constant comoving density (c.f. Goobar et al. 2002). In models of this type the extinction increases rapidly with redshift, because the density scales as $(1+z)^3$, and would be easy to recognise; such models are excluded by the available data on SNe Ia (Riess et al. 2004). However, this very simple scenario does not offer a plausible description of the evolution of the dense clouds under consideration at $z \gtrsim 1$. We note, in particular, that at $z \gtrsim 1$ the microwave background radiation will play an increasingly important rôle in the thermodynamics of the outer, cooler regions of the cloud, where nucleation and precipitation of hydrogen dust (i.e. particles of solid H_2) takes place (Walker & Wardle 2019).

4.3.1 Extinction fluctuations

One constraint that any acceptable model should satisfy is that the *variations* in extinction, across different lines of sight at the same redshift, must be small, because of the small photometric scatter around the mean that is observed for type Ia SNe. This requirement parallels the situation with gravitational lensing of type Ia SNe: the small observed dispersion in fluxes places an upper limit on the mass of the elementary lumps of matter, such that the luminosity distance is well approximated by that of a homogeneous universe (Metcalf & Silk 1999; Zumalacárregui & Seljak 2018; Helbig 2020). The latter condition is satisfied if we expect to have a number $\langle N \rangle \gg 1$ elementary lumps of matter within the “beam” — i.e. inside the cone defined by the observer, at the vertex, and the photosphere of the source as the base. This is commonly referred to as the “filled beam” case, in contrast to the “empty beam” circumstance $\langle N \rangle \ll 1$ that is expected when the elementary lumps of matter are sufficiently massive. In the context of gravitational lensing the mass within the beam translates directly into beam convergence, κ (see §2), and in the filled beam case we expect small fractional fluctuations in the beam convergence

$$\frac{\delta\kappa}{\kappa} \sim \frac{1}{\sqrt{\langle N \rangle}} \ll 1. \quad (27)$$

The correspondence to extinction fluctuations is now obvious, because the expected optical depth to extinction is $\sim \langle N \rangle \times r_\lambda^2 / R_s^2$, and the fluctuations around the mean are smaller by a factor $\sqrt{\langle N \rangle}$. We note that the increased extinction caused by having $\sim \sqrt{\langle N \rangle}$ additional clouds in the beam (relative to the mean) is offset by the additional beam convergence (hence magnification) that they introduce — i.e. the photometric fluctuations of the two effects are totally anticorrelated.

The cloud mass and abundance, together with the assumed background geometry, determine the expected number of clouds in the beam. Whatever background is assumed, a useful reference point can be established by placing the entire matter content in the form of dense clouds of a single mass; for the (Planck) cosmology that we adopted in §3 that circumstance corresponds to $\langle N \rangle = 1$ at a cloud mass of $\sim 10^{-4} M_\odot$ (for $R_s = 3 \times 10^{15}$ cm at $z = 1$).

4.3.2 Can extinction be a large effect?

We have already given general expressions for the amount of extinction that is introduced by a population of dense clouds. There is, however, a specific question that is of particular interest in connection with the magnitude-redshift relation of type Ia SNe: can the extinction be large enough to reconcile the observations with a non-accelerating universe? In this section we address that question.

The observed dimming amounts to a difference, relative to a flat, matter dominated universe, of about 0.5 magnitudes at $z \simeq 1$. To produce this much extinction requires each cloud to obscure quite a large area. And the opaque radius cannot be larger than the cloud radius, so in order to produce a dimming as large as is observed we require a cloud population that has a covering fraction of at least 0.5 at $z = 1$. In the case of a flat universe dominated by a population of dense clouds that translates directly to a requirement that the individual clouds have an average column-density

$$\frac{M}{\pi R^2} \lesssim 0.2 \text{ g cm}^{-2}. \quad (28)$$

We emphasise that this is a limit on the average column-density of a single cloud — the peak (central) column-density of a cloud can

be much larger — so it is not a region of parameter space that is excluded by the results of §3 (figure 5 in particular). It is in fact difficult to relate the constraint on mean column (28) to a gaussian cloud model, because that model extends to infinite radius.

For the convenience of readers we translate our column-density limit (28) into a lower limit on the cloud radius of

$$R \gtrsim 6 \left(\frac{M}{10^{-5} M_\odot} \right)^{1/2} \text{ AU}, \quad (29)$$

where we have scaled to the low-mass end of the preferred range identified in §3 as being able to reproduce the observed, optical variability of quasars.

We are not aware of any fundamental reason why dense gas clouds should not have radii greater than the limit just derived, so it is in principle possible that extinction by dense clouds could have a substantial influence on the observed magnitude-redshift relation of type Ia SNe. There are, however, two caveats to add. First, the limit given in equation (28) is orders of magnitude smaller than the characteristic value (140 g cm^{-2}) estimated for clouds within galaxies by Walker (1999b), from consideration of the destructive collision rate between clouds. It might be possible to reconcile these two very different estimates for the mean column if the mass of each cloud is strongly concentrated towards the centre. In that case complete destruction would require a physical collision with an impact parameter that is small compared to the cloud radius, leading to a much larger estimate of the mean column. The second caveat is that for $M \sim 10^{-5} M_\odot$ our lower limit on cloud radius is only slightly smaller than the largest model of that mass obtained in theoretical structural modelling (figure 5 of Walker & Wardle 2019).

4.3.3 Other constraints on supernova dimming by dense clouds

The possible influence of grey extinction on distant SNe has been a topic of interest for the last twenty years, and a variety of different constraints have been established. In all cases those constraints have been formulated for models that are very different to ours and need to be reconsidered, as follows.

First, non-zero reddening is present to some degree in all models of dust extinction and has been used as a constraint by many authors (e.g. Aguirre 1999b; Goobar et al. 2002). Unfortunately we do not know the extinction characteristics for the dust within the dense gas clouds under consideration here, nor do we know the details of their internal structure, so we cannot derive firm quantitative constraints. However, as shown in §4.1 (particularly figure 6) the ratio of total-to-selective extinction for dense clouds is much greater than that of the constituent material, so that reddening becomes a much less useful constraint. A combination of large dust grains located within dense, highly opaque clouds would yield almost completely grey extinction and would be very difficult to exclude on the basis of colour.

Secondly, the absorption of starlight by intergalactic dust is expected to lead to a diffuse far-IR background as the dust reradiates the energy at longer wavelengths (Aguirre & Haiman 2000). However, for the cold, dense gas that would make up the sort of objects we consider, the dust particles are likely to be predominantly solid H_2 snowflakes (Walker & Wardle 2019), with optical absorption cross-sections that are negligible compared to their scattering cross-sections (Kettwich et al. 2015). Consequently the level of far-IR background associated with a given level of optical extinction is decreased by orders of magnitude and this constraint becomes ineffective.

Thirdly, dust particles scatter X-rays through substantial angles, and the lack of extended (arcminute) haloes around distant quasars has been used to set limits on the amount of intergalactic dust (Petric et al. 2006). However, X-rays scattered by dust that is located within dense gas clouds will be subject to photoelectric absorption by that gas, causing strong suppression of the scattering halo. Photoelectric absorption declines rapidly with increasing photon energy, so the suppression would be much weaker at high energies, but unfortunately most of the photon flux emitted by X-ray sources is usually to be found at low energies. Consequently the extinction constraints that can be derived from X-ray scattering haloes are also much weaker when that extinction is associated with dense gas clouds.

Finally, we are aware of one type of constraint on extinction that does not lose its effectiveness when the dust is in dense gas clouds: comparison between luminosity distances and angular diameter distances (e.g. More et al. 2009). In a transparent universe, luminosity distance and angular diameter distance differ only by a factor $(1+z)^q$ – where the index q depends on the type of luminosity under consideration (e.g. $q = 2$ for bolometric luminosity) – thus providing a direct test for transparency. The resulting constraints depend strongly on how the question is framed: (cosmological) model-dependent analyses indicate small extinction coefficients (e.g. $\lesssim 0.02$ mag Gpc $^{-1}$, Goobar et al. 2018), whereas (cosmological) model-independent analysis is much less constraining ($\lesssim 0.2$ mag Gpc $^{-1}$, Liao et al. 2015). Clearly it is the latter, looser constraint that is the relevant one when it is the cosmological model itself that is at issue, so we are unable to exclude extinction by dense clouds as the cause of the dimming of distant SNe Ia.

5 QUADRUPOLE TEST

The extent to which objects along the line of sight contribute to quasar variability has been a controversial subject. To help clarify the issue a number of tests have already been undertaken, with results that in the main favour a microlensing interpretation: chromaticity of the variability statistics (Hawkins 2003); lack of a cosmological time dilation (Hawkins 2010); statistical symmetry in the light curves (Hawkins 1996); and, the relationship between variations in the continuum and broad emission lines (Hawkins 2011). Despite these successes the microlensing interpretation of quasar variability is not widely accepted — perhaps because the framing of these tests relies, to some degree, on our imperfect understanding of quasar structure and evolution.

Here we propose a new test — one that is kinematic in nature and is therefore less model dependent than the tests which have been implemented to date. The timescale of the extrinsic variability — be that gravitational or gas lensing or extinction — depends on the effective transverse velocity of the line of sight, which is a weighted sum of the transverse velocities of the source, lens and observer. Now the transverse projection of the observer’s three-dimensional velocity varies systematically over the sky, being zero in the direction parallel (or anti-parallel) to the observer motion, and is maximised at 90 degrees to that axis. Hence, there should be two poles of slowest variation on the sky and a band of fastest variation halfway between the poles; that is the test we are proposing, which we call ‘quadrupole’ after its dominant pattern (as higher multipoles might be present depending on which quantity — e.g., timescale, variation rate *etc.* — is actually measured). The influence of the observer’s velocity on microlensing event rates has long been recognised (e.g. Mosquera & Kochanek 2011; Mediavilla et al.

2016; Neira et al. 2020), but to our knowledge it has not previously been understood to offer a test for an extrinsic origin of quasar variability.

The measured dipole component of the cosmic microwave background anisotropy corresponds to a speed of $V_o = 370$ km s $^{-1}$ (Planck Collaboration et al. 2020) with respect to the CMB. Interpreted as a peculiar velocity this speed is unremarkable (Tully et al. 2013) and is widely interpreted as predominantly kinematic in origin (Planck Collaboration et al. 2014), although uncomfortable discrepancies have been found for the amplitude of a related signal in the number counts of radio (Singal 2019) and infrared sources (Secrest et al. 2021). The peculiar velocities of individual lenses and sources are not known, but they are not expected to align with the observer velocity so the quadrupole pattern due to the observer’s motion will be superimposed on a monopole. Moreover the quadrupole is small compared to the monopole, as we now show.

The expression for the effective velocity is (Kayser et al. 1986):

$$\mathbf{v}_{\text{eff}} = \mathbf{v}_l - \left(\frac{1+z_l}{1+z_s} \right) \frac{D_l}{D_s} \mathbf{v}_s - \frac{D_{ls}}{D_s} \mathbf{v}_o, \quad (30)$$

where \mathbf{v}_o refers to the transverse (to the line of sight) component of the observer three-dimensional velocity such that

$$v_o = V_o \sin \theta, \quad (31)$$

with θ being the angle the line of sight makes with the observer’s motion. Assuming the first two terms in (30) are drawn from a gaussian distribution with zero mean, so is their combination and the scaled effective velocity squared has a non-central χ^2 distribution with two degrees of freedom. The mean of that distribution is

$$\langle v_{\text{eff}}^2 \rangle_{v_{s,l}} = 2\sigma^2 + \left(\frac{D_{ls}}{D_s} \right)^2 v_o^2, \quad (32)$$

where σ^2 is the (per spatial component) variance of the first two terms in equation (30):

$$\sigma^2 = \sigma_l^2 + \left(\frac{1+z_l}{1+z_s} \right)^2 \left(\frac{D_l}{D_s} \right)^2 \sigma_s^2. \quad (33)$$

The source velocity dispersion ought to be lower as the sources are seen at earlier stages in the development of cosmic structure, when the peculiar velocities are expected to be smaller, and because quasars are located near the centres of massive galaxies. The effect of the source velocity variance, σ_s^2 , is itself reduced by redshift and lever-arm ratios and will be neglected in the following.

The relative amplitude of the quadrupole effect

$$\eta \equiv \frac{\Delta \langle v_{\text{eff}} \rangle_{v_{s,l}}}{\langle v_{\text{eff}} \rangle_{v_{s,l}}} \simeq \left(\frac{D_{ls}}{D_s} \right)^2 \frac{V_o^2}{4\sigma^2} \quad (34)$$

depends explicitly on the position of the lenses along the line of sight and the appropriate effective value would depend on the effect responsible for variability. We assume that gravitational lensing, rather than extinction, is the variability mechanism of interest — as is expected for high column-density clouds. Lensing is most effective when the critical density $\Sigma_{cr}(z_l, z_s)$ is lowest, which, at moderate source redshifts, occurs roughly halfway between the source and the observer. Thus the factor in parentheses above is close to 1/2, yielding

$$\langle \eta \rangle_{l.o.s.} \simeq \frac{V_o^2}{16\sigma^2} \quad (35)$$

as our estimate. A more accurate estimate would require detailed numerical modelling.

For a (one-dimensional) lens velocity dispersion of $\sigma \sim 400 \text{ km s}^{-1}$, as appropriate in the field, the relative amplitude of the variation across the sky is therefore small: $\eta \sim 0.05$, requiring a precision of better than one percent in different sky directions. With the variations due to lensing being stochastic, the measured timescale itself is a random variable and the required precision can only be achieved by averaging over many sources. The variation in the timescale is related to those in the lens position along the line of sight, intrinsic spatial autocorrelation scale of the illumination pattern and relative velocity. Because lensing is ineffective for lenses close to the source and observer, this filter selects lenses at a particular distance (about halfway to the source) with moderate variation thereof. The intrinsic autocorrelation scale depends on the parameters of the lens, so the contribution of this effect into the variation will be modest if there is a preferred range of cloud parameters in the Universe. In this case the variation will be dominated by that of the effective velocity given by

$$\frac{\langle (\Delta v_{\text{eff}}^2) \rangle}{\langle v_{\text{eff}}^2 \rangle} = \frac{\sqrt{1 + (D_{ls}/D_s)^2 v_o^2 / \sigma^2}}{1 + (D_{ls}/D_s)^2 v_o^2 / 2\sigma^2} \approx 1 \quad (36)$$

implying averaging variability timescales in 10^4 s sources for the required accuracy or \sim per cent. This is likely a lower limit because it assumes that variability arising from any mechanism other than gravitational lensing can be neglected. On the other hand, detection could be aided by demanding that the anisotropy is aligned with that of the cosmological dipole. Ultimately, however, the required level of averaging will be determined empirically as a number of sources necessary to make the average stable from one patch of the sky to an adjacent one at the required level.

The above estimate suggests that the currently available quasar light curve data, such as the Southern Stripe 82 sample of the SDSS DR7 (MacLeod et al. 2012), are inadequate for the proposed test. However, high quality photometry of over a million quasars (Gaia Collaboration et al. 2020), distributed all over the sky, will be made available in the next data release of the *Gaia* mission; these data have the potential for a highly significant detection of the timescale-quadrupole, or else a rejection of the nanolensing interpretation of quasar variability. We caution that the peculiar – and, most importantly, position-dependent – time sampling of *Gaia* targets will require careful accounting of the influence of that sampling on the derived timescale estimates.

6 SUMMARY AND CONCLUSIONS

We have described the effects of a cosmological population of compact gas clouds on the radiation from distant sources. Gas clouds affect light in three different ways: through extinction, gas refraction and gravitational lensing, leading to a wide variety of possible light curves. The general character of the light curves can be anticipated for any circumstances and we demonstrated how to classify them according to the specified cloud properties. In the particular case of a gaussian column-density profile we showed the effect of a cloud of fixed mass but with central column-densities ranging over five orders of magnitude, including examples of weak lensing, strong gas lensing, and strong gravitational lensing — for both transparent and highly opaque clouds.

Despite the potential importance of gas lensing – which can substantially increase the lensing cross-section in some circumstances – we find that neutral gas refraction has little influence on observations of quasars and type Ia supernovae at maximum

light (or later). This result stands in contrast to Draine (1998), and Rafikov & Draine (2001), who studied lensing of stars by gas clouds in and around the Galaxy. The origin of that difference is the much bigger sources that we consider — which, at $R_s \sim (1 - 3) \times 10^{15} \text{ cm}$, are large compared to the gas lensing curvature radius ($\ell_\lambda \approx 3.6 \times 10^{13} \text{ cm}$), which in turn is large compared to the radii of dwarf stars ($R_s \sim 10^{11} \text{ cm}$). The key point is that we require a cloud size $r_0 < \ell_\lambda$ for gas lensing to be stronger than gravitational lensing, so if $R_s \gg \ell_\lambda$ then that gas lens cannot greatly magnify the source.

For clouds with high central column-densities ($\Sigma_0 \gtrsim 1 \text{ g cm}^{-2}$ for low-opacity clouds; $\Sigma_0 \gtrsim 10 \text{ g cm}^{-2}$ for high-opacity clouds) our simulations yield quasar light-curves that are largely independent of that column — all being similar to the point-mass case. This similarity is unsurprising because when gravitational lensing is important the dominant images form outside the cloud itself, so the details of its structure (and the associated gas refraction and extinction) have little influence. In this circumstance the mass of the individual lenses and their contribution to the mean cosmological density (Ω_l) are the key variables that determine the appearance of the light-curves.

We find that high mass clouds ($10^{-4} M_\odot \ll M \lesssim 10^{-2} M_\odot$) yield a great deal of large-amplitude quasar variability on timescales of a few years, and a high cosmological density of such objects can be excluded — as originally pointed out by Schneider (1993). At the low end of our studied cloud mass range, below $10^{-5} M_\odot$, we find only small variations in the magnification due to gravitational lensing, for sources as large as quasars; a high density of such clouds is permitted by the data, but cannot explain the variations that are seen. And between these extremes there is a narrow range of cloud mass, $M \sim 10^{-4.5 \pm 0.5} M_\odot$, for which gravitational lensing by a population of gas clouds can explain much of the observed quasar variability, given a suitable mean cosmological density in such objects. All of these mass estimates would change systematically if we were to repeat our simulations using a quasar optical size that differs from our adopted value of $R_s = 3 \times 10^{15} \text{ cm}$, with the scaling being $M \propto R_s^2$ (for $R_s \sim 3 \times 10^{15} \text{ cm}$).

If the material transparency is high the “sweet spot” – wherein lensing can explain the observed variability – applies to clouds across the full range of central column-densities that we studied ($0.1 \text{ g cm}^{-2} \lesssim \Sigma_0 \lesssim 1000 \text{ g cm}^{-2}$). But if the material transparency is as low as in the local ISM then gaussian clouds that have low central columns ($\Sigma_0 \lesssim 1 \text{ g cm}^{-2}$) are excluded because they introduce unacceptably large extinction fluctuations. Indeed such low-column, low-transparency gaussian clouds are in conflict with the quasar data for the full range of cloud masses studied in this paper.

The quasar variability data exclude large-amplitude extinction fluctuations but say little about the mean extinction. It is standard practice to use reddening as a gauge of extinction; however, we showed that the extinction introduced by dense gas clouds is very nearly colour-neutral — even when the material opacity is strongly chromatic. This situation arises because of saturation: the large column-densities under consideration are totally opaque across a wide range of wavelengths, and the overall level of extinction of a source just reflects the fraction of its solid angle that is obscured by these opaque regions. As has been pointed out previously (Aguirre 1999a), the possible presence of grey extinction introduces ambiguity into the interpretation of data on standard candles such as type Ia supernovae (e.g. Riess et al. 1998; Perlmutter et al. 1998). In this regard the dense gas clouds we have studied are particularly

insidious because their total-to-selective extinction ratio is naturally very high indeed, and consequently colour is ineffective at discriminating between geometry and extinction. We have also noted how two other proposed constraints on intergalactic dust – based on X-ray scattering (e.g. [Petric et al. 2006](#)), and on the properties of the cosmic infrared background ([Aguirre & Haiman 2000](#)) – are similarly ineffective. A direct test of cosmic transparency is available by comparing luminosity distances with angular diameter distances; however, when formulated in a framework that is independent of the cosmological model the resulting constraint is at present quite weak ([Liao et al. 2015](#)).

Finally, we described a new statistical test which can differentiate between intrinsic and extrinsic origins of quasar variability. Extrinsic variations depend on the relative transverse velocity between the source, lens and observer, and the observer contribution to that velocity varies systematically over the sky. Because it relies only on kinematics, our test does not require the lenses to be gas clouds. The expected form of the variation over the sky is a quadrupolar pattern in the distribution of the average rate of external variations. With the source and lens velocity being randomly distributed, the amplitude of this variation is second order in the observer velocity and it is further suppressed by the lever-arm ratio that enters into the effective velocity. For gravitational lensing this results in an estimated amplitude of $\eta \sim 0.05$, implying that the average variation timescales in different patches of the sky need to be measured to better than one percent. This is only possible by averaging timescales from a large number of light curves – at least 10^4 are needed to overcome the dispersion in the effective velocity itself. The currently available compilations of light curves are inadequate for this task, but the situation will improve when per-epoch *Gaia* photometry of one million quasars becomes available. The complex sampling pattern of *Gaia*, varying systematically across the sky, must be precisely accounted for when undertaking the proposed test.

DATA AVAILABILITY

The data underlying this article will be shared on reasonable request to the corresponding author.

ACKNOWLEDGEMENTS

Helpful feedback on the quadrupole test was provided by Rudy Schild, Geraint Lewis and Vikram Ravi, and on the extinction of SNe Ia by Dominik Schwarz — all much appreciated. Thanks also to Geraint for comments on the manuscript, and to the referee, Mike Hawkins, for some good suggestions for improvements. This research made use of Astropy,³ a community-developed core Python package for Astronomy ([Astropy Collaboration et al. 2013, 2018](#)).

REFERENCES

- Aguirre A. N., 1999a, *ApJ*, **512**, L19
Aguirre A., 1999b, *ApJ*, **525**, 583
Aguirre A., Haiman Z., 2000, *ApJ*, **532**, 28
Alcock C., et al., 2000, *ApJ*, **542**, 281
Astropy Collaboration et al., 2013, *A&A*, **558**, A33

- Astropy Collaboration et al., 2018, *AJ*, **156**, 123
Blackburne J. A., Pooley D., Rappaport S., Schechter P. L., 2011, *ApJ*, **729**, 34
Blandford R. D., McKee C. F., 1982, *ApJ*, **255**, 419
Bohlin R. C., Savage B. D., Drake J. F., 1978, *ApJ*, **224**, 132
Bolatto A. D., Wolfire M., Leroy A. K., 2013, *ARA&A*, **51**, 207
Boveia A., Doglioni C., 2018, *Annual Review of Nuclear and Particle Science*, **68**, 429
Cardelli J. A., Clayton G. C., Mathis J. S., 1989, *ApJ*, **345**, 245
Carr B., Kühnel F., 2020, *Annual Review of Nuclear and Particle Science*, **70**, 355
Carr B., Kohri K., Sendouda Y., Yokoyama J., 2021, *Reports on Progress in Physics*, **84**, 116902
Cen R., Ostriker J. P., 1999, *ApJ*, **514**, 1
Clegg A. W., Fey A. L., Lazio T. J. W., 1998, *ApJ*, **496**, 253
Combes F., Pfenniger D., 1997, *A&A*, **327**, 453
Cornachione M. A., Morgan C. W., Millon M., Bentz M. C., Courbin F., Bonvin V., Falco E. E., 2020, *ApJ*, **895**, 125
Dai X., Kochanek C. S., Chartas G., Kozłowski S., Morgan C. W., Garmire G., Agol E., 2010, *ApJ*, **709**, 278
Dalgarno A., Williams D. A., 1962, *ApJ*, **136**, 690
Draine B. T., 1998, *ApJ*, **509**, L41
Drake A. J., Cook K. H., 2003, *ApJ*, **589**, 281
Edelson R., et al., 2015, *ApJ*, **806**, 129
Edelson R., et al., 2017, *ApJ*, **840**, 41
Fausnaugh M. M., et al., 2016, *ApJ*, **821**, 56
Feng J. L., 2010, *ARA&A*, **48**, 495
Fukugita M., Hogan C. J., Peebles P. J. E., 1998, *ApJ*, **503**, 518
Gaia Collaboration et al., 2020, arXiv e-prints, p. arXiv:2012.02036
Gerhard O., Silk J., 1996, *ApJ*, **472**, 34
Goobar A., Bergström L., Mörtzell E., 2002, *A&A*, **384**, 1
Goobar A., Dhawan S., Scolnic D., 2018, *MNRAS*, **477**, L75
Hawkins M. R. S., 1993, *Nature*, **366**, 242
Hawkins M. R. S., 1996, *MNRAS*, **278**, 787
Hawkins M. R. S., 2003, *MNRAS*, **344**, 492
Hawkins M. R. S., 2010, *MNRAS*, **405**, 1940
Hawkins M. R. S., 2011, *MNRAS*, **415**, 2744
Hawkins M. R. S., 2020a, *A&A*, **633**, A107
Hawkins M. R. S., 2020b, *A&A*, **643**, A10
Hawkins M. R. S., Veron P., 1993, *MNRAS*, **260**, 202
Helbig P., 2020, *The Open Journal of Astrophysics*, **3**, 1
Henriksen R. N., Widrow L. M., 1995, *ApJ*, **441**, 70
Homayouni Y., et al., 2019, *ApJ*, **880**, 126
Huchra J., Gorenstein M., Kent S., Shapiro I., Smith G., Horine E., Perley R., 1985, *AJ*, **90**, 691
Jiang Y.-F., et al., 2017, *ApJ*, **836**, 186
Jiménez-Vicente J., Mediavilla E., Kochanek C. S., Muñoz J. A., 2015, *ApJ*, **799**, 149
Kayser R., Refsdal S., Stabell R., 1986, *A&A*, **166**, 36
Kerins E., Binney J., Silk J., 2002, *MNRAS*, **332**, L29
Kettwich S. C., Anderson D. T., Walker M. A., Tuntsov A. V., 2015, *MNRAS*, **450**, 1032
Kochanek C. S., 2004, *ApJ*, **605**, 58
Liao K., Avgoustidis A., Li Z., 2015, *Phys. Rev. D*, **92**, 123539
MacLeod C. L., et al., 2012, *ApJ*, **753**, 106
Macquart J. P., et al., 2020, *Nature*, **581**, 391
Mediavilla E., Jiménez-Vicente J., Muñoz J. A., Battaner E., 2016, *ApJ*, **832**, 46
Metcalf R. B., Silk J., 1999, *ApJ*, **519**, L1
More S., Bovy J., Hogg D. W., 2009, *ApJ*, **696**, 1727
Morgan C. W., Kochanek C. S., Morgan N. D., Falco E. E., 2010, *ApJ*, **712**, 1129
Morgan C. W., Hyer G. E., Bonvin V., Mosquera A. M., Cornachione M., Courbin F., Kochanek C. S., Falco E. E., 2018, *ApJ*, **869**, 106
Mortonson M. J., Schechter P. L., Wambsganss J., 2005, *ApJ*, **628**, 594
Mosquera A. M., Kochanek C. S., 2011, *ApJ*, **738**, 96
Mudd D., et al., 2018, *ApJ*, **862**, 123
Neira F., Anguita T., Vernardos G., 2020, *MNRAS*, **495**, 544

³ <http://www.astropy.org>

- Paczynski B., 1986, *ApJ*, **304**, 1
- Perlmutter S., et al., 1998, *Nature*, **391**, 51
- Petric A., Telis G. A., Paerels F., Helfand D. J., 2006, *ApJ*, **651**, 41
- Pfenniger D., Combes F., 1994, *A&A*, **285**, 94
- Pfenniger D., Combes F., Martinet L., 1994, *A&A*, **285**, 79
- Planck Collaboration et al., 2014, *A&A*, **571**, A27
- Planck Collaboration et al., 2016, *A&A*, **594**, A13
- Planck Collaboration et al., 2020, *A&A*, **641**, A1
- Poindexter S., Kochanek C. S., 2010a, *ApJ*, **712**, 658
- Poindexter S., Kochanek C. S., 2010b, *ApJ*, **712**, 668
- Pooley D., Blackburne J. A., Rappaport S., Schechter P. L., 2007, *ApJ*, **661**, 19
- Pringle J. E., 1981, *ARA&A*, **19**, 137
- Rafikov R. R., Draine B. T., 2001, *ApJ*, **547**, 207
- Riess A. G., et al., 1998, *AJ*, **116**, 1009
- Riess A. G., et al., 2004, *ApJ*, **607**, 665
- Schechter P. L., Wambsganss J., 2004, in Ryder S., Pisano D., Walker M., Freeman K., eds, Vol. 220, *Dark Matter in Galaxies*. p. 103 ([arXiv:astro-ph/0309163](https://arxiv.org/abs/astro-ph/0309163))
- Schild R. E., 1996, *ApJ*, **464**, 125
- Schneider P., 1993, *A&A*, **279**, 1
- Schneider P., Ehlers J., Falco E. E., 1992, *Gravitational Lenses*, doi:10.1007/978-3-662-03758-4.
- Secrest N. J., von Hausegger S., Rameez M., Mohayaee R., Sarkar S., Colin J., 2021, *ApJ*, **908**, L51
- Shakura N. I., Sunyaev R. A., 1973, *A&A*, **500**, 33
- Singal A. K., 2019, *Phys. Rev. D*, **100**, 063501
- Trimble V., 1987, *ARA&A*, **25**, 425
- Tully R. B., et al., 2013, *AJ*, **146**, 86
- Udalski A., Zebrun K., Szymanski M., Kubiak M., Pietrzynski G., Soszynski I., Wozniak P., 2000, *Acta Astron.*, **50**, 1
- Walker M. A., 1999a, *MNRAS*, **306**, 504
- Walker M. A., 1999b, *MNRAS*, **308**, 551
- Walker M. A., Wardle M. J., 2019, *ApJ*, **881**, 69
- Wechsler R. H., Tinker J. L., 2018, *ARA&A*, **56**, 435
- Wyithe J. S. B., Webster R. L., Turner E. L., 2000, *MNRAS*, **315**, 51
- Yu Z., et al., 2020, *ApJS*, **246**, 16
- Zackrisson E., Bergvall N., 2003, *A&A*, **399**, 23
- Zackrisson E., Bergvall N., Marquart T., Helbig P., 2003, *A&A*, **408**, 17
- Zumalacárregui M., Seljak U., 2018, *Phys. Rev. Lett.*, **121**, 141101
- de Paolis F., Ingrosso G., Jetzer P., Roncadelli M., 1995, *Phys. Rev. Lett.*, **74**, 14
- de Swart J. G., Bertone G., van Dongen J., 2017, *Nature Astronomy*, **1**, 0059

APPENDIX A: QUASAR OPTICAL CONTINUUM SIZE

Source size plays a critical rôle in the modelling of nanolensing variations because it determines the spatial smoothing scale of the flux (magnification) map. In the present paper we use a representative source radius of $R_s = 3 \times 10^{15}$ cm, which is three times larger than the value used by Schneider (1993). Our chosen value is close to that expected from standard, thin disc theory for a $10^9 M_\odot$ black hole accreting at a fair fraction of the Eddington rate (e.g. Morgan et al. 2010). In this Appendix we describe three different ways of estimating the size of the optical continuum emission region in quasars.

A1 Photometric estimate for standard accretion disc

Models of thin accretion discs are simplest in the case where the radiation is optically-thick thermal emission, which is locally well approximated by the Planck function, and at large radii, R , the temperature, T , varies as a power-law: $T \propto R^{-3/4}$ (Shakura & Sunyaev 1973; Pringle 1981). This functional form is expected to be a good approximation outside the hottest portion of the disc, which in turn is

close to its inner edge at 1–6 gravitational radii (depending on the angular momentum of the hole), i.e. $10^{14.2-15}$ cm $\times (M_{\text{BH}}/10^9 M_\odot)$.

In the case of locally black-body emission and a power-law variation of temperature with radius, an observing frequency ν picks out a characteristic radius, R_ν , such that

$$\frac{h\nu}{kT} = \left(\frac{R}{R_\nu} \right)^{2/m}. \quad (\text{A1})$$

In terms of that radius, the luminosity spectral density, L_ν , inferred by an observer at infinity, and polar angle θ relative to the disc normal, is given by

$$L_\nu = 16\pi^2 R_\nu^2 \cos \theta \frac{h\nu^3}{c^2} \mathcal{I}(m), \quad (\text{A2})$$

where

$$\mathcal{I}(m) \equiv \int_0^\infty \frac{x dx}{\exp(x^{2/m}) - 1} = \frac{1}{2} \Gamma(1+m) \zeta(m), \quad (\text{A3})$$

with Γ and ζ being the Gamma (factorial) Function and the Riemann Zeta Function, respectively. The integral converges for $1 < m < \infty$. For a standard thin disc the appropriate power-law is $m = 8/3$ and the integral evaluates to $\mathcal{I}(8/3) \approx 2.576$.

It is important to note that equation (A2) does not explicitly depend on the mass of the black hole, its accretion rate, or the viscosity in the disc. Consequently the luminosity can be inferred from the measured source flux and redshift alone, within the adopted cosmological model, and thus we can immediately determine the corresponding value of R_ν — up to a factor of order unity that is associated with the unknown disc inclination.

In §3, for reasons of computational efficiency, we used a uniform surface-brightness source model. In detail the nanolensing flux variations do depend on the surface-brightness profile of the source, so there can be only an approximate equivalence between that description and an accretion disc model. However, it is known that the source size — as gauged by the half-light radius — is the most important factor in determining the magnification statistics (Mortonson et al. 2005). For a disc having a power-law temperature profile with $m = 8/3$ the half-light radius is $R_{1/2} \approx 2.436 R_\nu$, whereas a uniform disc has a radius, R_s , that is $\sqrt{2}$ larger than its half-light radius, thus the radius of the equivalent uniform disc is $R_s \approx 3.444 R_\nu$ ($m = 8/3$).

Figure A1 shows the value of R_s appropriate to the observed B -band magnitudes and redshifts of the Hawkins & Veron (1993) sample of quasars in the same (Planck) cosmological model that we used in §3. In all cases we have taken a representative value of $\cos \theta = 0.5$. From this figure we conclude that a suitable radius for a uniform disc model for this sample of sources is $R_s = 3 \times 10^{15}$ cm.

In the case of other physical models that predict different power-law variations of temperature with radius, the equivalent radius of a uniform disc scales approximately as $R_s \propto m^2$ in the vicinity of $m = 8/3$ (valid over the range $1.5 \lesssim m \lesssim 4.5$).

A2 Continuum reverberation mapping

Reverberation mapping is a technique that has its origins in studies of the broad line region (BLR) of quasars. The physical picture is that the emission lines arise from photoionisation and should therefore vary in response to fluctuations in the UV continuum flux from an accretion disc. The temporal impulse response function of the BLR, as recorded by an observer far away, depends on its geometry, so photometric monitoring of both continuum and line




Figure A1. The radius of the equivalent disc of uniform surface-brightness having the same half-light radius as a locally black-body disc ($R_s = 3.444 R_v (m = 8/3)$), for each of the quasars in the [Hawkins & Veron \(1993\)](#) sample. These values are computed using equation (A2), assuming no extinction and average inclination of the accretion disc to the line of sight ($\cos \theta = 0.5$). All radii are for the observed *B*-band, which corresponds to progressively shorter wavelengths in the rest frame of the quasar at higher redshifts.

fluxes can constrain the spatial distribution of the emission line clouds ([Blandford & McKee 1982](#)).

Continuum reverberation mapping is a related idea: high energy photons are imagined to be absorbed by the accretion disc itself (e.g. because the disc is flared), and the absorbed power is subsequently reradiated in the form of lower energy photons. In this setting photometric monitoring in a variety of continuum passbands can constrain the geometry of the continuum emitting regions. A detailed physical description of this mechanism is currently lacking, but even so one can examine data to see whether they correspond well to the idea, and to establish quantitative constraints.

Continuum reverberation mapping measurements of nearby active galactic nuclei in NGC 4151 and NGC 5548 (which have well-measured black hole masses) confirmed the predicted scaling of accretion disc size with wavelength (equation A1 with $m = 8/3$), but returned lag amplitude about a factor of 3 higher than expected ([Fausnaugh et al. 2016](#); [Edelson et al. 2015](#)). At the same time these studies called into question the validity of the temporal response model – often referred to as the “lamp post model” – underlying the method ([Edelson et al. 2017](#)). A similar size discrepancy was reported from a statistical analysis of high signal-to-noise lag measurements for a large number of quasars in the Pan-STARRS data ([Jiang et al. 2017](#)).

Recent continuum reverberation studies, however, have not supported the idea of large continuum sources. High quality light curves from the Sloan Digital Sky Survey (nearly a hundred quasars, [Homayouni et al. 2019](#)) and Dark Energy Survey (over twenty quasars, [Mudd et al. 2018](#); [Yu et al. 2020](#)) returned lag estimates – hence source sizes – that are fully consistent with the size of

standard, thin accretion discs. Furthermore it was pointed out by [Homayouni et al. \(2019\)](#) that earlier statistical work, which focused on objects with high signal-to-noise lag measurements, was strongly biased towards large lags and therefore large disc sizes. In summary: at present there is no evidence from continuum reverberation mapping for quasar sizes larger than expected in the standard, thin accretion disc theory.

A3 Microlensing of multiply-imaged quasars

Quasars that are gravitationally lensed by foreground galaxies can offer insights into the properties of the source. When there are multiple, distinct images one can obtain information about the intrinsic changes in source flux – which, up to a lag, are the same for all images – and also about the microlensing, which differs amongst the different paths. As noted earlier, the character of the microlensing signal is strongly affected by the size of the source.

A3.1 Flux ratio anomalies

Smooth models of the gravitational potential of the lens galaxy can explain the positions of the lensed images, but the observed flux ratios often differ from the corresponding predictions and are thus termed “anomalous” (e.g. [Schechter & Wambsganss 2004](#)). In reality the lens galaxy potential has structure on small scales – e.g. due to individual stars – and so the observed flux ratio anomalies have a natural interpretation in terms of small-scale structure in the lensing magnification map.

In a study of ten quadruply imaged quasars [Pooley et al. \(2007\)](#) found that the flux ratio anomalies are much smaller in the optical than in the X-ray. They interpreted this as the quasar optical emission region being larger than the X-ray emission region, so that the magnification map is more strongly smoothed in the optical. Using simulations [Pooley et al. \(2007\)](#) determined the optical source size to be one-tenth to one-third of the typical Einstein ring radius of the microlenses. These results were reinforced with an augmented sample of 12 quasars by [Blackburne et al. \(2011\)](#) — who also find a source size scaling with wavelength that is much flatter, on average, than the standard accretion disc theory predicts.

In order to translate these results into an absolute source size one needs to know the typical Einstein ring radius of the microlenses, whereas the data themselves do not determine that radius. [Pooley et al. \(2007\)](#) and [Blackburne et al. \(2011\)](#) assumed stellar mass microlenses, which led them to conclude that optical quasars are much larger than standard, thin accretion discs. If the dark matter is indeed an elementary particle then the assumption of stellar mass microlenses is reasonable. However, in this paper we are considering the possibility of a universe with a high density of sub-stellar objects, and in that context one also expects small-scale structure in the magnification map arising from nanolenses.

A3.2 Light-curve analysis

The degeneracy just mentioned can in principle be lifted in studies of microlensing light-curves, because the temporal coordinate translates to a spatial scale through the effective transverse speed. However, that quantity is itself not known a priori and the best that can be done is to construct a probability distribution based on reasonable physical assumptions.

One lensed quasar, Q2237+0305 ([Huchra et al. 1985](#)), has received a great deal of attention in the literature — because of the

uniquely low redshift of the lens galaxy, and the correspondingly short timescale of its microlensing variations. This lens galaxy is a massive spiral and all four images lie within one kpc of the centre, so we are seeing the quasar through a high surface density of stars. In this case one might imagine that we are on firm ground if we assume the microlenses to be of stellar mass. In fact [Wytke et al. \(2000\)](#) excluded the possibility that a significant fraction of the total surface density could be in planetary mass objects, on the grounds that the resulting microlensing rate would be much higher than observed. By the same token, however, it remains possible that the observed microlensing could include a significant contribution from planetary mass objects: even a small surface density in such objects could yield a higher event rate than the population of stellar microlenses. In connection with this possibility it is worth noting that a virialised dark halo of dense gas clouds is expected to have a core in its density profile ([Walker 1999b](#)), as a result of physical collisions, and should indeed make only a small contribution to the surface-density near the centre *even if dense gas comprises all of the dark matter*.

Contemporary analyses of microlensing in Q2237+0305 utilise Bayesian inference to identify viable combinations of the source size, microlens mass and effective transverse speed. Early work (e.g. [Kochanek 2004](#)) returned source sizes consistent with standard accretion discs, whereas recent results have favoured larger sources (e.g. [Poindexter & Kochanek 2010a,b](#)). However, to date none of these analyses have considered the circumstance just mentioned, in which a small admixture of planetary mass objects contributes significantly to the observed microlensing signal. Consequently it is unclear what these studies have to say about the appropriate quasar size to use in the present paper.

In large part the analyses of microlensing in other multiply-imaged quasars have been fashioned after the pioneering Q2237+0305 studies (notably the work of [Kochanek 2004](#)), and there are many papers which favour quasar sizes significantly larger than standard accretion discs (e.g. [Dai et al. 2010](#); [Morgan et al. 2018](#); [Cornachione et al. 2020](#)). However, just like Q2237+0305, all of these studies are based on monomodal distributions of the microlens mass, whereas a bimodal distribution is required if the dark matter consists of planetary mass lumps.