

A quasi-periodic modulation of the iron line centroid energy in the black hole binary H1743–322

Adam Ingram,¹★ Michiel van der Klis,¹ Matthew Middleton,² Chris Done,³
Diego Altamirano,⁴ Lucy Heil,¹ Phil Uttley¹ and Magnus Axelsson⁵

¹*Anton Pannekoek Institute for Astronomy, University of Amsterdam, Science Park 904, NL-1098 XH Amsterdam, the Netherlands*

²*Institute of Astronomy, Cambridge University, Madingley Road, Cambridge CB3 0HA, UK*

³*Center for Extragalactic Astronomy, Department of Physics, University of Durham, South Road, Durham DH1 3LE, UK*

⁴*Department of Physics & Astronomy, University of Southampton, Southampton, Hampshire SO17 1BJ, UK*

⁵*Department of Physics, Tokyo Metropolitan University, Minami Osawa 1-1, Hachioji, Tokyo 192-0397, Japan*

Accepted 2016 May 19. Received 2016 May 19; in original form 2016 April 15

ABSTRACT

Accreting stellar-mass black holes often show a ‘Type-C’ quasi-periodic oscillation (QPO) in their X-ray flux and an iron emission line in their X-ray spectrum. The iron line is generated through continuum photons reflecting off the accretion disc, and its shape is distorted by relativistic motion of the orbiting plasma and the gravitational pull of the black hole. The physical origin of the QPO has long been debated, but is often attributed to Lense–Thirring precession, a General Relativistic effect causing the inner flow to precess as the spinning black hole twists up the surrounding space–time. This predicts a characteristic rocking of the iron line between red- and blueshift as the receding and approaching sides of the disc are respectively illuminated. Here we report on *XMM–Newton* and *NuSTAR* observations of the black hole binary H1743–322 in which the line energy varies systematically over the ~ 4 s QPO cycle (3.70σ significance), as predicted. This provides strong evidence that the QPO is produced by Lense–Thirring precession, constituting the first detection of this effect in the strong gravitation regime. There are however elements of our results harder to explain, with one section of data behaving differently than all the others. Our result enables the future application of tomographic techniques to map the inner regions of black hole accretion discs.

Key words: accretion, accretion discs – black hole physics – relativistic processes – X-rays: individual: H1743–322.

1 INTRODUCTION

Accreting stellar-mass black holes routinely exhibit ‘Type-C’ low-frequency quasi-periodic oscillations (QPOs) in their X-ray flux, with a frequency that evolves from ~ 0.1 to 30 Hz as the X-ray spectrum transitions from the power-law-dominated hard state to the thermal disc-dominated soft state (e.g. Wijnands, Homan & van der Klis 1999; van der Klis 2006). The thermal disc component is well understood as originating from a geometrically thin, optically thick accretion disc (Novikov & Thorne 1973; Shakura & Sunyaev 1973), and the power-law emission, which displays breaks at low and high energy, is produced via Compton up-scattering of seed photons by a cloud of hot electrons located close to the black hole (Thorne & Price 1975; Sunyaev & Truemper 1979). The low- and high-energy breaks are associated respectively with the seed photon temperature and the electron temperature. The exact geometry of

this electron cloud is uncertain, and is probably changing through this transition. In the *truncated disc model* (Ichimaru 1977; Esin, McClintock & Narayan 1997; Poutanen, Krolik & Ryde 1997), the disc truncates in the hard state at some radius larger than the innermost stable circular orbit (ISCO), with the inner regions forming a large scaleheight, hot accretion flow (hereafter the *inner flow*) which emits the Comptonized spectrum. The Comptonized spectrum becomes softer and the disc component becomes more prominent in the spectrum as the truncation radius moves inwards (e.g. Done, Gierlinski & Kubota 2007). Alternatively, the hot electrons may be located in a corona above the disc (e.g. Galeev et al. 1979; Haardt & Maraschi 1991) or at the base of a jet (e.g. Markoff, Nowak & Wilms 2005), or perhaps some combination of these alternatives. The X-ray spectrum also displays reflection features, formed by Comptonized photons being scattered back into the line of sight by the disc. The most prominent features of the reflection spectrum are the iron $K\alpha$ line at ~ 6.4 keV, formed via fluorescence, and the reflection hump peaking at ~ 30 keV, formed via inelastic scattering from free electrons (e.g. Ross & Fabian 2005; García et al. 2013).

★ E-mail: a.r.ingram@uva.nl

The shape of the reflection spectrum, and in particular the iron line which is narrow in the rest frame, is distorted by orbital motion of the disc material and gravitational redshift (Fabian et al. 1989).

The QPO arises from the immediate vicinity of the black hole,¹ but its physical origin has long been debated. Suggested QPO models in the literature generally consider either some instability in the accretion flow or a geometric oscillation. Instability models consider, for example, oscillations in mass accretion rate or pressure (e.g. Tagger & Pellat 1999; Cabanac et al. 2010) or standing shocks in the disc (e.g. Chakrabarti & Molteni 1993). Geometric models mostly consider relativistic precession. Due to the frame dragging effect, a spinning black hole drags the surrounding space–time around with it, inducing Lense–Thirring precession in the orbits of particles out of the equatorial plane (Lense & Thirring 1918). Stella & Vietri (1998) and Stella, Vietri & Morsink (1999) were the first to suggest that low-frequency QPOs could be driven by Lense–Thirring precession, noting that the expected precession frequency of a test mass at the truncation radius is commensurate with the QPO frequency. Schnittman, Homan & Miller (2006) considered a precessing ring in the disc, and corrugation modes in the disc caused by the frame dragging effect have also been studied (e.g. Wagoner, Silbergleit & Ortega-Rodríguez 2001). Ingram, Done & Fragile (2009) suggested that the entire inner flow precesses whilst the disc remains stationary, motivated by the simulations of Fragile et al. (2007). This model explains why the QPO modulates the Comptonized emission much more than the disc emission, and predicts that the QPO should be stronger in more highly inclined sources as observed (Schnittman et al. 2006; Heil, Uttley & Klein-Wolt 2015; Motta et al. 2015). It also makes a distinctive prediction: as the inner flow precesses, it will illuminate different azimuths of the disc such that an inclined observer sees a blue/redshifted iron line when the approaching/receding sides of the disc are illuminated (Ingram & Done 2012). The precession model therefore predicts that the line energy changes systematically with QPO phase. This is a difficult effect to measure, since phase resolving the QPO poses a technical challenge. Miller & Homan (2005) used a simple flux selection to obtain suggestive but inconclusive results for GRS 1915+105. Ingram & van der Klis (2015, hereafter IK15) developed a more sophisticated technique to discover spectral pivoting and a modulation in the iron line flux in the same source, but data quality prevented unambiguous measurement of a line energy modulation. Recently, Stevens & Uttley (2016) developed a similarly sophisticated QPO phase-resolving technique, which involves cross-correlating each energy channel with a reference band. Using this technique, they found a modulation in the disc temperature of GX 339–4, interpreted as reprocessed radiation from a precessing inner flow or jet. However, they too lacked the data quality to measure a line energy modulation.

In this paper, we further develop the QPO phase-resolving method of IK15, conducting fitting in the Fourier domain rather than the time domain so that the error bars are independent. We use this method to analyse a long-exposure observation of the black hole binary H1743–322 in the hard state. We summarize the observations in Section 2, describe our phase-resolving method in Section 3 and present the results of fitting a phenomenological model to the phase-resolved spectra in Section 4. We discuss our findings in Section 5 and outline our conclusions in Section 6.

¹ The light crossing time-scale puts a hard upper limit of $\sim 300 R_g$ (where $R_g = GM/c^2$), but the true size scale is likely $\lesssim 60 R_g$ (e.g. Axelsson, Hjalmarsson & Done 2013).




Figure 1. Long-term light curve summarizing the *XMM–Newton* and *NuSTAR* observations analysed in this paper. Throughout, the data are referred to as labelled in this plot. The rise in count rate from *XMM–Newton* orbit 2a to 2b is due to a small change in the instrument setup (from pn thick to pn thin) as a new PI took over the observation.

2 OBSERVATIONS

The *X-ray Multi-Mirror Mission* (*XMM–Newton*; Jansen et al. 2001) observed H1743–322 for two full orbits of the satellite around the Earth in late September 2014. The first orbit (obs ID 0724400501) lasted from $\sim 18:45$ on September 21 until $\sim 08:45$ on September 23. The second orbit lasted from $\sim 20:10$ on September 23 until $\sim 08:35$ on September 25. The second orbit is split into two obs IDs, with the first ~ 70 ks classified as 0724401901 and the final ~ 50 ks, which had a different PI, as 0740980201. In this paper, we split up each orbit into two separate observations to allow for potential evolution of spectral and timing properties over such long exposures (and also the small change in instrumental setup as the PI changed). Hereafter, we refer to these four *XMM–Newton* observations as orbits 1a, 1b, 2a and 2b. The *Nuclear Spectroscopic ARray* (*NuSTAR*; Harrison et al. 2013) observed the source from $\sim 18:20$ on September 23 until $\sim 08:50$ on September 25 (obs ID 80001044004). Fig. 1 shows long-term 4–10 keV light curves for all exposures and illustrates our naming convention for the *XMM–Newton* data. Spectral and timing analyses of the *XMM–Newton* data have been previously presented by Stiele & Yu (2016) and De Marco & Ponti (2016), whereas the *NuSTAR* data are reported on here for the first time.

2.1 Data reduction

2.1.1 XMM–Newton

We used the *XMM–Newton* Science Analysis Software (SAS) version 14.0 to reduce data from the EPIC-pn (European Photon Imaging Camera) in timing mode. We generated calibrated and concatenated event lists using `EPPROC` with the default settings for timing mode as of SAS v14.0 (`runepreject=yes withxrlcorrection=yes runepfast=no withdrpha=yes`). We extracted all products from a region $32 \leq \text{RAWX} < 44$, $\text{RAWY} \geq 23$ and use only single and double events (`PATTERN ≤ 4`), whilst ignoring bad pixels (`FLAG == 0`). We generated response and ancillary files using `RMFGEN` and `ARFGEN`, and rebinned all spectra to have at least 20 counts per channel using `SPECGROUP`. We extracted background spectra from the region $3 \leq \text{RAWX} \leq 5$, $\text{RAWY} \geq 23$ and find that the source contributes 98.5 per cent of the total counts (this number is likely even higher in reality, since source counts can contaminate the background spectrum in timing mode; Done & Diaz Trigo 2010). Since the source dominates, we did not perform a background subtraction when extracting light curves. Inspection of the long-term 10–12 keV light curve reveals that none of the exposure is affected by proton flares.




Figure 2. 4–10 keV power spectrum for *XMM-Newton* orbits 1a (black), 1b (red), 2a (green) and 2b (blue), and 4–10 keV co-spectrum between the *NuSTAR* FPMA and FPMB (magenta). For all data sets, we see a strong Type-C QPO with two clearly detected harmonics. The QPO frequency increased from ~ 0.2 to ~ 0.25 Hz over the ~ 300 ks elapsed from the start of orbit 1a to the end of orbit 2b. Error bars are 1σ .

We extract light curves in 20 energy bands. We focus our phase-resolved analysis on the 4–10 keV region, so extract one broad light curve for energies < 4 keV and one broad light curve for energies > 10 keV (both of which will be ignored for the analysis), leaving 18 high signal-to-noise channels in the region of interest. These channels are broad enough to achieve good statistics, and are trivially broader than the full width at half-maximum (FWHM) of the instrument response. We used the `FTOOL RBNRMF` in order to rebin the spectral response file into these 20 energy bands.

2.1.2 NuSTAR

We used the *NuSTAR* analysis software, `NUSTARDS v1.4.1`. We extracted products from the cleaned event list with the `FTOOL NUPRODUCTS`, using a 120 arcsec circular source extraction region and a 90 arcsec circular background extraction region taken from an area not contaminated by source counts. We find that the source contributes 99.7 per cent of the total counts, and consequently we did not perform a background subtraction when extracting light curves. The background is negligible up to ~ 50 keV, above which it dilutes the rms and phase lags by a small amount. We extract light curves in 19 energy bands. We concentrate on the energy range 4–60 keV, and so bin into two broad channels for energies < 4 keV and one broad channel > 60 keV (with these three channels to be ignored in the analysis), leaving 16 high signal-to-noise channels in the range of interest. As for *XMM-Newton*, we rebinned the spectral response file using `RBNRMF`.

2.2 Power spectra

Fig. 2 shows 4–10 keV power spectra calculated for *XMM-Newton* orbits 1a (black), 1b (red), 2a (green) and 2b (blue) and *NuSTAR* (magenta). The *XMM-Newton* power spectra are calculated in the standard way, with a constant Poisson noise level subtracted (van der Klis 1989; Uttley et al. 2014). For *NuSTAR* we instead calculate the co-spectrum between the two (independent) focal plane modules,

FPMA and FPMB (Bachetti et al. 2015), since the *NuSTAR* dead time of $\tau_d \approx 2.5$ ms imprints instrumental features on the Poisson noise in a power spectrum calculated in the standard way. The co-spectrum is the real part of the cross-spectrum and includes no Poisson noise contribution. We also correct for the suppression of variability caused by the *NuSTAR* dead time using the simple formula (Bachetti et al. 2015)

$$\frac{\text{rms}_{\text{det}}}{\text{rms}_{\text{in}}} \approx \frac{1}{1 + \tau_d r_{\text{in}}} = \frac{r_{\text{det}}}{r_{\text{in}}}, \quad (1)$$

where r_{det} and r_{in} are respectively the detected and intrinsic count rates. For this observation, the ratio of detected to intrinsic variability is $\text{rms}_{\text{det}}/\text{rms}_{\text{in}} = 0.8462$ (recorded in the *NuSTAR* spectral files as the keyword ‘DEADC’). The power spectra in Fig. 2 are normalized such that the integral of the power spectrum over a given frequency range gives the variance of the corresponding time series over that range, and are plotted in units of frequency \times power.

All power spectra display QPOs with a strong fundamental (first harmonic) and overtone (second harmonic) evidenced by two large, harmonically related peaks. We see that the QPO fundamental frequency evolved from ~ 0.205 to ~ 0.25 Hz over the ~ 300 ks duration of the two *XMM-Newton* orbits. We also see that the 4–10 keV (dead time corrected) *NuSTAR* co-spectrum agrees very well with the simultaneous *XMM-Newton* orbit 2 data for the same energy band. For our analysis, we treat each of the five data sets shown in Fig. 2 separately to allow for the evolution in source properties over such a long exposure, and also to allow for the different responses of the two instruments, and the small change in the *XMM-Newton* instrumental setup during orbit 2.

2.3 Energy spectra

As a preliminary analysis, we jointly fit the spectra of both *XMM-Newton* orbit 2a and the simultaneous (FPMA) *NuSTAR* observation with a simple absorbed power-law plus Gaussian iron line model, considering only 4–10 keV for both. Throughout this paper, we account for interstellar absorption using the model `TBABS`, with hydrogen column density $N_{\text{H}} = 1.35 \times 10^{22} \text{ cm}^{-2}$ and the relative abundances of Wilms et al. (2000). We use `XSPEC v12.8.2` for all spectral fitting (Arnaud 1996). We achieve a best fit with reduced $\chi^2 = 551.14/529 = 1.04$, without applying any systematic error. There is no evidence for direct disc emission in the > 4 keV bandpass, and the *NuSTAR* spectrum above 10 keV reveals a reflection hump. In this paper, we focus on phenomenological modelling of the 4–10 keV region for our QPO phase-resolved analysis, modelling continuum and iron line with a power law and Gaussian, respectively. We consider this bandpass because it is shared between *XMM-Newton* and *NuSTAR*; it is above the energies for which direct disc emission is relevant and below energies for which the reflection hump is important. Clearly, a Gaussian function is not a physical model for the iron line, but we wish to characterize the QPO phase dependence of the iron line profile without making physical assumptions. We will focus on physical spectral modelling in a future paper.

We find a discrepancy in the power-law index measured for these two spectra (1.286 ± 0.003 for *XMM-Newton* and 1.509 ± 0.004 for *NuSTAR*). The Gaussian representing the iron line has a larger equivalent width in the *NuSTAR* spectral fit (~ 65 eV) than in the *XMM-Newton* data (~ 47 eV), and lower centroid energy in the *NuSTAR* data (~ 6.41 keV) than in the *XMM-Newton* data (~ 6.61 keV), but the line width is consistent. This cross-calibration discrepancy poses a problem for time-averaged spectral analysis.

However, our analysis is differential: it focuses on the variation of spectral parameters with QPO phase, and is therefore far more robust to cross-calibration issues. We demonstrate in the following two sections that the variability properties are consistent between the two observatories, and that the differential variation in each of the spectral parameters with QPO phase is consistent. For our phase-resolved spectral analysis, we allow the time-averaged power-law index and line energy to be different between the two observatories, but tie their *differential* properties between the two observatories.

3 PHASE-RESOLVING METHOD

We use the phase-resolving method of IK15, with some small changes designed to increase signal-to-noise and circumvent the *NuSTAR* dead time, and also some more significant changes to allow us to reliably calculate statistical significances for our fits. The essence of the IK15 phase-resolving method is to measure the Fourier transform (FT) of the QPO as a function of energy E , for each harmonic for which this is possible. For the j th harmonic, this can be written as

$$W_j(E) = \mu(E)\sigma_j(E)e^{i\Phi_j(E)}, \quad (2)$$

where $\mu(E)$ is the mean count rate, $\sigma_j(E)$ is the fractional rms in the j th QPO harmonic and $\Phi_j(E)$ is referred to in IK15 as the *phase offset* of the j th harmonic, all as a function of energy. It is clear from Fig. 2 that the QPO for the observations considered here has two strong harmonics; therefore, we calculate the QPO FT for $j = 1$ and 2. We must also consider the case of $j = 0$; i.e. the DC component (standing for direct current). This is simply the mean count rate, such that $W_0(E) = \mu(E)$.

As for the phase offsets, $\Phi_j(E)$, we can calculate the cross-spectrum between each energy channel and a reference band in order to measure the phase lag for each harmonic, $\Delta_j(E)$, as a function of energy. That is, we can measure by how many radians the j th harmonic of each energy channel lags the j th harmonic of the reference band. What we can *not* measure using the cross-spectrum is the phase difference between the harmonics. By measuring this phase difference, we can calculate the phase offsets of the first two harmonics using the formulae

$$\begin{aligned} \Phi_1(E) &= \Phi_1 + \Delta_1(E) \\ \Phi_2(E) &= 2[\Phi_1 + \psi] + \Delta_2(E). \end{aligned} \quad (3)$$

Here, ψ is the phase difference between the two harmonics in the reference band and Φ_1 is the arbitrary reference phase of the first harmonic, which we set to $\Phi_1 = \pi/2$ following IK15. Note that there is a version of the above formula in IK15 (equation 8 in that paper), which differs slightly from equation (3) presented here. The version presented here is correct and the mistake is in IK15. Note that we only need to measure the phase difference between the harmonics in one band (it is obviously advantageous to measure this for the reference band which has far more photons than the individual channels). The phase difference between the harmonics as a function of energy is given by $\psi(E) = \psi - \Delta_1(E) + \Delta_2(E)/2$. We stress that the use of a broad reference band does *not* smear out the data in some way, as is a common misconception. For unity coherence, changing the reference band affects *only* the constant offset of the lag spectrum and also the signal-to-noise (see e.g. Uttley et al. 2014).

The method of IK15 involves taking the inverse FT of equation (2) to give an estimate of the QPO waveform in each energy channel. This method is intuitive, since it gives a way of estimating the




Figure 3. Spectra from two selected QPO phases, plotted as a ratio to an absorbed power-law continuum model. The blue circles correspond to a QPO phase a quarter of a cycle later than the red triangles (11/16 cycles compared with 7/16 cycles). We show *XMM-Newton* data averaged between orbits 2a and 2b. We see shifts in the iron line energy between the two selected QPO phases, and the hard X-ray coverage of *NuSTAR* (inset) additionally reveals that the reflection hump is enhanced relative to the line when the line is blueshifted. Error bars are 1σ .

spectrum as a function of QPO phase. The inverse FT, however, introduces correlations in the errors between different QPO phases. Here, we first summarize the time domain approach and then describe a new Fourier domain approach that circumvents the problem of correlated error bars associated with the time domain method.

3.1 Phase-resolved spectra in the time domain

We can inverse FT equation (2) to estimate the QPO waveform for each energy band

$$w(E, \gamma) = \mu(E) \left\{ 1 + \sqrt{2} \sum_{j=1}^2 \sigma_j(E) \cos[j\gamma - \Phi_j(E)] \right\}, \quad (4)$$

where γ is QPO phase. Plotting this instead as count rate versus photon energy for a given QPO phase gives phase-resolved spectra. We describe in the following subsections how we measure $\mu(E)$, $\sigma_j(E)$ and $\Delta_j(E)$, focusing mainly on the modifications we have made to the IK15 method in order to maximize signal-to-noise and correct for the *NuSTAR* dead time. We propagate the errors in equation (4) using a Monte Carlo simulation.

We first reconstruct phase-resolved spectra in the time domain using equation (4). We consider 16 QPO phases (i.e. 16 values of γ), and throughout we analyse each data set defined in Fig. 1 separately, resulting in five independent data sets. Fig. 3 shows examples of the phase-resolved spectra, plotted as a ratio to an absorbed power-law continuum model (folded around the telescope response matrix). The continuum model has been fitted ignoring 5.5–8 keV, where the iron line is prominent and >10 keV, where the reflection hump is prominent. For this plot, we only consider the *NuSTAR* data and *XMM-Newton* orbits 2a and 2b, which were simultaneous with the *NuSTAR* observation. For plotting purposes, we have averaged together data from orbits 2a and 2b, even though we treat them as two separate data sets in our analysis. Red triangles correspond to a QPO phase of $\gamma = 7/16$ cycles and blue circles to a QPO phase of $\gamma = 11/16$ cycles; i.e. the blue points are a quarter of a cycle after

the red points. We see that the line energy changes over the course of a QPO cycle, and the *NuSTAR* data reveal that the reflection hump becomes more prominent when the line energy is higher (blue triangles). In the following section, we model the iron line with a Gaussian and the continuum with an absorbed power law in order to characterize this QPO phase dependence of the line energy. However, in order to robustly assess the statistical significance of the line energy modulation, we fit the same model in the Fourier domain, as described in the following subsection.

3.2 Phase-resolved spectra in the Fourier domain

It is straightforward to fit the 16 phase-resolved spectra, as expressed in equation (4), with a phenomenological spectral model to determine if the best-fitting spectral parameters vary systematically with QPO phase. However, assessing the statistical significance of the spectral parameters is complicated by correlations between the errors for different QPO phases. For this reason, this is not the method we use to determine significances. Instead, we perform the fits in the Fourier domain, which provides a different representation of the same information. The QPO FT, $W_j(E)$, from equation (2), is in units of count rate and, as a complex quantity, can be expressed in terms of amplitude, $\mu(E)\sigma_j(E)$, and phase, $\Phi_j(E)$, or in terms of real and imaginary parts, $\Re\{W_j(E)\} = \mu(E)\sigma_j(E)\cos[\Phi_j(E)]$ and $\Im\{W_j(E)\} = \mu(E)\sigma_j(E)\sin[\Phi_j(E)]$, respectively. The real and imaginary parts of $W_j(E)$, and the different harmonics, are statistically independent from one another. Thus, standard statistical methods can be applied if we fit a model to $W_j(E)$ rather than $w(E, \gamma)$. Here, we first fit spectral models to the phase-resolved spectra in the time domain to gain insight, before constructing a model for the QPO FT. We can exploit the linearity of the FT to define a model, $\tilde{W}_j(E)$, and fold around the telescope response to get the observed $W_j(E)$, as for a normal spectrum. Specifically, for the l th energy channel

$$W_j(E_l) = \int_0^\infty \tilde{W}_j(E) R(l, E) dE, \quad (5)$$

where $R(l, E)$ is the telescope response for the l th energy channel. We perform a joint fit to real and imaginary parts to preserve this linearity (which would be lost if we were to instead fit for amplitude and phase). This results in a joint fit of five spectra: the real and imaginary parts of the first and second harmonics, and the real part of the DC component (the imaginary part is trivially zero).

3.3 Phase difference between harmonics

We first measure the phase difference, ψ , between the two QPO harmonics. This phase difference represents the number of QPO cycles by which the second harmonic (first overtone) lags the first harmonic (fundamental), converted to radians (i.e. multiplied by 2π). It is defined on the interval $0-\pi$ rad, since there are two cycles of the second harmonic for each cycle of the fundamental. We split the full band light curve into segments of duration 32 s. Each segment contains 512 time bins of duration $dt = 0.0625$ s, and roughly eight QPO cycles. *XMM-Newton* orbits 1a, 1b, 2a and 2a contain respectively 2135, 2134, 2455 and 1535 segments with good telemetry, and the *NuSTAR* observation contains 2217 segments. For each segment, we calculate the phase difference ψ following IK15. In Fig. 4, we plot a histogram of these ψ values for each data set (the colour scheme is the same as defined in Fig. 1), revealing a strong peak for all data sets. These histograms have two peaks purely because ψ is cyclical and we show two cycles. We measure the peak of each histogram following IK15 to obtain the




Figure 4. Phase difference between the two QPO harmonics, with different data sets represented using the same colour scheme as Fig. 1. For all data sets, we measure the phase difference between harmonics ψ for many 32 s segments (see the text for details). This plot is a histogram of those measurements and shows that there is a well-defined average phase difference between the harmonics, which we measure by determining the peak of the plotted distribution.

average phase difference between harmonics. For *XMM-Newton* orbits 1a, 1b, 2a, 2b, we measure $\psi/\pi = 0.309 \pm 0.005$, 0.336 ± 0.005 , 0.336 ± 0.005 and 0.347 ± 0.006 . For *NuSTAR*, we take the average of the independent measurements for the FPMA and FPMB to get $\psi/\pi = 0.332 \pm 0.005$. Note that, even though the *NuSTAR* observation is simultaneous with orbit 2 of *XMM-Newton*, the measured ψ are not required to agree because the full band light curves of *XMM-Newton* and *NuSTAR* cover a different energy range. The agreement we see between observatories tells us that the phase difference has little energy dependence here.

3.4 Energy dependence of QPO amplitude

We measure the fractional rms amplitude of the two QPO harmonics as a function of energy, $\sigma_j(E)$, for all five data sets. IK15 did this by calculating the power spectrum for each energy channel. For the *XMM-Newton* data here, we instead calculate the covariance spectrum to increase signal-to-noise (Wilkinson & Uttley 2009). We follow the standard procedure for calculating the covariance and its error (Uttley et al. 2014). Our reference band is the full *XMM-Newton* band minus the channel of interest so as to avoid correlating a time series with itself. For each energy channel, we calculate the cross-spectrum between that channel and the reference band, and also the power spectrum of the reference band. The covariance is the modulus squared of the cross-spectrum divided through by the power spectrum of the reference band. Since the light curves from each energy channel are well correlated, the covariance gives a good measure of the power spectrum with smaller statistical errors (Wilkinson & Uttley 2009). For *NuSTAR*, we circumvent the dead time by calculating the co-spectrum between the FPMA and FPMB light curve for each energy channel instead of the power spectrum. Following IK15, we fit our power spectral estimates (covariance and co-spectrum for *XMM-Newton* and *NuSTAR*, respectively) in each energy channel with a sum of Lorentzian functions. We calculate the fractional rms of each QPO harmonic from the integral of the corresponding Lorentzian function. A dead time correction




Figure 5. Fractional rms as a function of energy for three selected data sets, represented using the same colour scheme as Fig. 1. The points above the grey dashed line correspond to the fundamental (first harmonic) and the points below the dashed line are for the second harmonic. The other data sets are omitted for clarity. We see tentative features around the iron line. Error bars are 1σ .

of $\text{rms}_{\text{det}}/\text{rms}_{\text{in}} = 0.8462$ also must be applied to the *NuSTAR* data. Fig. 5 shows the resulting calculation of rms as a function of energy for three of the five data sets, following the colour scheme of Fig. 1. We show only three data sets to avoid overcrowding the plot. The points above the dashed line are for the first harmonic, and below the dashed line are for the second harmonic.

For our Lorentzian fits, we use four Lorentzian functions, one for each QPO harmonic and two to fit the broad-band noise. We tie the centroid of the second harmonic component to be double that of the first harmonic and force the two QPO Lorentzians to have the same quality factor ($Q = \text{centroid frequency}/\text{FWHM}$). The centroid and quality factor of the QPO fundamental component are free to vary with energy, but we measure no significant energy dependence for either of these quantities. We tried many variations on the model to test for the robustness of the fit. We tried using more and less broad-band noise Lorentzians, allowing the QPO components to have different quality factors, relaxing the centroid frequency ratio of 2, fixing the widths and/or centroid frequencies of the QPO Lorentzians to equal those measured for the full band and so on. We even tried simply integrating the power spectral estimates over the widths of the QPO components instead of fitting a model. In all cases, we obtained consistent results, indicating that our fits are robust.

3.5 Phase lags between energy bands

We calculate the phase lag between each energy channel and a broad reference band for both QPO harmonics, $\Delta_i(E)$. For *XMM-Newton*, we use the same reference band as described above for the covariance spectrum. We calculate the cross-spectrum for each channel and average this over the width of each QPO harmonic, as defined by the Lorentzian fitting described in the previous section. The phase lag for each QPO harmonic is the argument of this averaged cross-spectrum. For *NuSTAR*, we again utilize the two independent focal plane modules. We use the full FPMB band as the reference band and calculate the cross-spectrum between this and each channel of interest in FPMA. We also calculate an independent set of




Figure 6. Phase lag of each energy channel relative to a reference band for three selected data sets, represented using the same colour scheme as Fig. 1. The reference band is the full band of the respective instrument, and therefore is slightly different between *XMM-Newton* and *NuSTAR*. This creates the small offset seen in the second harmonic. As with Fig. 5, the other data sets are omitted for clarity. Error bars are 1σ .

cross-spectra using FPMA as the reference band and FPMB for the subject bands. For each energy channel, we average together these two independent measurements of the cross-spectrum to increase signal-to-noise. Fig. 6 shows the lag spectra for the same three data sets as the previous plot. We see a small offset in the second harmonic between *XMM-Newton* and *NuSTAR*. This is simply because the lag spectra are calculated for *NuSTAR* using a different reference band, and the lag of the second harmonic depends on energy. This will introduce a small offset between *XMM-Newton* and *NuSTAR* when it comes to plotting best-fitting spectral parameters against QPO phase. As it turns out, this offset is small enough to ignore completely, but even if it were large, it would be fairly simple to correct for since it is just a constant offset. We calculate the phase offsets $\Phi_1(E)$ and $\Phi_2(E)$ using equation (3).

3.6 Step-by-step summary

The steps of the **IK15** method can be summarized as follows:

- (i) measure the phase difference between the QPO harmonics in a broad reference band (see Section 3.3);
- (ii) for both QPO harmonics, measure the rms variability as a function of energy (see Section 3.4);
- (iii) for both QPO harmonics, measure the phase lag between each energy channel and the reference band (see Section 3.5);
- (iv) combine these measurements in order to calculate the QPO FT (equation 2);
- (v) inverse FT to obtain a waveform for each energy channel (equation 4).

For the Fourier domain method, we stop at step 4 and fit a model directly to the QPO FT, whereas the time domain method also includes step 5.

4 RESULTS

Fig. 3 shows the spectrum for two selected QPO phases plotted as a ratio to an absorbed power-law continuum model, with the blue circles representing the spectrum a quarter of a cycle later than the red triangles. We see a shift in line energy between the two QPO phases. In this section, we fit the iron line with a Gaussian function to characterize the phase dependence of the centroid energy and assess its statistical significance.

4.1 Time domain fits

We first fit the phase-resolved spectra in the time domain with an absorbed power-law plus Gaussian model, in the energy range 4–10 keV. We consider the five data sets separately, which allows us to compare results for independent analyses. We initially tie all spectral parameters to remain constant during the QPO cycle and test if the fit is improved when we allow each parameter to vary freely with QPO phase. For all three data sets, we achieve the minimum reduced χ^2 value by allowing the Gaussian centroid energy (E_{line}), Gaussian flux (N_G) and the power-law index (Γ) and normalization (N_{cont}) to vary with QPO phase. The fit is not improved by allowing the Gaussian width to vary with QPO phase. We plot the best-fitting line centroid energy against QPO phase (light blue circles) in Fig. 7. We do not plot error bars here, since the errors are correlated between QPO phases in the time domain fits. All data sets show a modulation in the line energy. For all but orbit 1b of *XMM-Newton*, the line energy modulation has the same distinctive shape, with maxima at ~ 0.2 and ~ 0.7 cycles. The modulations in Γ , N_G and N_{cont} (not pictured) are also consistent between these four data sets.

It is puzzling that orbit 1b disagrees with the other data sets. This data set also exhibits different modulations in N_G and Γ from the others (this can be seen in Fig. 10, which is explained in detail in the following sections). To investigate this further, we split up orbit 1 into four quarters such that the first two quarters together make up orbit 1a and the final two quarters together make up orbit 1b. We find, as expected, that the first two quarters both show the same modulation in line shape seen for orbit 1a. The fourth quarter (i.e. the second half of orbit 1b) shows a peak in line energy at 0.7 cycles but not at 0.2 cycles, so is different from orbit 1a but only slightly. This fourth quarter also shows a modulation in Γ consistent with orbit 1a. It is the third quarter that differs so radically from all of the other data sets. This shows a peak in line energy at ~ 0.4 cycles, and also exhibits a different (but very weak) modulation in Γ from orbit 1a. The fact that the two halves of orbit 1a display the same modulations as each other, and as orbit 1a treat as a whole, gives us confidence in the robustness of the method, and implies that something different really is happening in this third quarter of orbit 1.

4.2 Fourier domain fits

We now fit the same phenomenological spectral model directly to the QPO FT derived from the data. This will allow us to assess the statistical significance of the line energy modulation. We construct a model for the QPO FT by representing the spectral parameters as periodic functions of QPO phase, γ . For example, the line energy is

$$E_{\text{line}}(\gamma) = E_0 + A_{1E} \sin[\gamma - \phi_{1E}] + A_{2E} \sin[2(\gamma - \phi_{2E})], \quad (6)$$

where E_0 , A_{1E} , A_{2E} , ϕ_{1E} and ϕ_{2E} are model parameters. We see that E_0 is the mean line energy, and all variability in the line energy as a function of QPO phase is captured by the amplitudes and




Figure 7. Iron line centroid energy as a function of QPO phase for each of the five data sets (as labelled). Light blue circles are the results of our time domain fits and the probability maps are the results of our Fourier domain fits. For the time domain fits, we fit an absorbed power-law plus Gaussian model to spectra corresponding to 16 QPO phases. For the Fourier domain fits, we consider the same model, with parameters varying periodically with QPO phase, and FT the model to fit to the data in Fourier space. We determine the significance of the modulations (as labelled) from the Fourier domain fits, and create the probability maps using a Monte Carlo Markov chain (see the text for details). The maps are normalized such that they peak at unity, and the colours are defined in the key. Note that $P/P_{\text{max}} = 0.1$, below which the colour scale looks rather white, corresponds to the $\sim 2.15\sigma$ confidence contour.

phases of the sine waves. The other potentially varying spectral parameters (Gaussian width and normalization, power-law index and normalization) are also modelled in the same manner with five parameters each. Our model calculates the resulting spectrum for 16 QPO phases and then calculates the FT for a grid of energy bins. We then fold, for each harmonic, the real and imaginary parts of this FT around the telescope response matrix (equation 5) and fit to the observed QPO FT (equation 2).

4.2.1 Separate fits

As with the time domain fits, we fit the five data sets separately, expecting to see exactly the same results as before (since the FT of a function is simply a different representation of the same function), but with more manageable statistics. We again find a best fit with modulations in the line energy (i.e. $A_{1E} > 0$ and $A_{2E} > 0$) and flux and the continuum normalization and power-law index, and again




Figure 8. Left: QPO FT as a function of energy for *XMM-Newton* orbit 1a. Real and imaginary parts of the first and second harmonics are as labelled. Here, we plot the data (black points) unfolded around the instrument response matrix, assuming the best-fitting model (red lines for the first harmonic and blue lines for second harmonic), in units of energy squared \times specific photon flux (i.e. the *eeuf* option in *XSPEC*). Right: QPO FT for the anomalous *XMM-Newton* orbit 1b (grey triangles and blue lines) compared with *XMM-Newton* orbit 1a (black circles and red lines), only considering the second harmonic (real and imaginary parts as labelled). We see clear differences in the shape for both real and imaginary parts. To demonstrate that *XMM-Newton* orbit 1a is representative of all the other data sets, we also plot the best-fitting model for orbit 2b. Error bars are 1σ .

the fits are not improved by allowing the Gaussian width to vary with QPO phase. As an example of our fits, we plot in Fig. 8 (left) the QPO FT for orbit 1a (black points) along with the best-fitting model (lines). Here, the data are unfolded around the instrument response assuming the best-fitting model and are in units of energy squared \times specific photon flux (i.e. the *eeuf* option in *XSPEC*). The best-fitting model for the first harmonic is plotted in red and the second harmonic in blue. We see features in the data and model around the iron line, which result in the model from modulations of the line energy and equivalent width. The second harmonic shows the clearest features, with an excess at ~ 6.2 keV in the real part and a dip at the same energy in the imaginary part, surrounded by two peaks either side. We fit jointly for the real and imaginary parts of both harmonics, and also for the mean spectrum (the real part of the DC component) which is not pictured here.

We plot the $E_{\text{line}}(\gamma)$ function derived from our Fourier domain fits, visualized as a probability map, in Fig. 7. The best-fitting function $E_{\text{line}}(\gamma)$ can be plotted by substituting the best-fitting values for E_0 , A_{1E} , A_{2E} , ϕ_{1E} and ϕ_{2E} into equation (6). Here, we also take into account the probability distributions of these five parameters by running a Monte Carlo Markov chain (MCMC) in *XSPEC* and then, for each step in the chain, calculating the $E_{\text{line}}(\gamma)$ function. We then create a histogram to plot the posterior distribution for the function. Details of the chain and of the calculation of these histograms are presented in Appendix A. As expected, the frequency domain results plotted in Fig. 7 agree with the time domain fits (light blue circles), but we are now able to visualize the uncertainty on the best-fitting line energy modulation (see the key).

We calculate the significances quoted in Fig. 7 by comparing the χ^2 from the best-fitting model for each data set with the minimum χ^2 achieved for the same data set when the line energy amplitudes are fixed to $A_{1E} = A_{2E} = 0$. This null hypothesis model has 4 more degrees of freedom than the best-fitting model, because it is insensitive to the phase parameters, ϕ_{1E} and ϕ_{2E} . We compare the best fit to the null hypothesis using an *f*-test, converting *p*-values to sigmas in the standard way (e.g. 1σ corresponds to $p = 0.317$).

4.2.2 The case of orbit 1b

As with the time domain fits, it is striking that all data sets except for orbit 1b show the same characteristic trend, with peaks in line energy at ~ 0.2 and ~ 0.7 cycles. Looking at the QPO FT reveals that the difference between orbit 1b and all the other data sets is in the second harmonic. Fig. 8 (right) shows the FT of the second harmonic (real and imaginary parts as labelled) for orbit 1a (black circles) and orbit 1b (grey triangles). The best-fitting models for orbits 1a and 1b are plotted in red and blue, respectively. We see very different behaviour between the two data sets. Where orbit 1a shows a dip (real part at ~ 7 keV), orbit 1b shows an excess. Where orbit 1a shows an excess (imaginary part at ~ 7 keV), orbit 1b shows a dip. All other data sets display similar behaviour to orbit 1a. To illustrate this, we plot the best-fitting model for orbit 2b in magenta. This has a slightly different normalization, but the same characteristic shape as orbit 1a.

We check if these differences can result from our assumptions when measuring the fractional rms as a function of energy. For the many different methods of measuring this described in Section 3.4, we measure QPO FTs consistent with before and therefore obtain results consistent with Fig. 7. We therefore conclude that the method produces robust results and that orbit 1b really does seem to be doing something different from the other data sets.

4.2.3 Joint fits

Since all data sets show a modulation in line energy, we combine them into a joint fit to compare with the null hypothesis: $A_{1E} = A_{2E} = 0$. We first leave out the anomalous data set, orbit 1b. We see in Fig. 7 that the two maxima in line energy measured for *NuSTAR* slightly lead those measured for the simultaneous *XMM-Newton* orbit 2. This is because we used a different reference band for *NuSTAR*. Since this constant offset turns out to be very small in this case, we are able to ignore it. We therefore tie the modulations in line energy, line flux and power-law index to be the same for all

four considered data sets, but allow the power-law normalization to differ for different data sets. We note that the modulation in power-law normalization is very similar for each data set (even including orbit 1b), but is well constrained enough for small differences in data sets to be highly statistically significant. We tie the power-law index between *XMM-Newton* and *NuSTAR* using the formula $\Gamma_{\text{NuSTAR}}(\gamma) = \Gamma_{\text{XMM}}(\gamma) + \Delta\Gamma$, in order to account for the cross-calibration discrepancy. Similarly, we tie the line energy between observatories using the formula $E_{\text{line, NuSTAR}}(\gamma) = CE_{\text{line, XMM}}(\gamma)$. We use $\Delta\Gamma = 0.236$ and $C = 0.970$. We obtain a good fit (reduced $\chi^2 = 287.13/279 = 1.029$) with the differential properties of the spectral parameters consistent with before.

When we also include the orbit 1b data in our fit, tying all parameter modulations except for the power-law normalization across all data sets, we obtain a fit with reduced $\chi^2 = 370.32/364 = 1.017$. When we allow ϕ_{1E} and ϕ_{2E} to be different for orbit 1b compared with all the other data sets (as seems to be the case from Fig. 7), the fit improves with reduced $\chi^2 = 361.82/362 = 1.000$. An *f*-test determines that this is a 2.43σ improvement, indicating that the line energy modulation in orbit 1b is likely different from the other data sets. Also freeing A_{1E} and A_{2E} for orbit 1b does not further improve the fit, so we keep these amplitudes tied across all data sets. When we also allow the second harmonic amplitudes of the N_G and Γ modulations to be different for orbit 1b, the fit again improves, with reduced $\chi^2 = 354.24/360 = 0.984$. An *f*-test indicates that this is a 2.29σ improvement. This is our best-fitting model.

Our best-fitting parameters for the line energy modulation are presented in Table 1. Fig. 9 (left) is a contour plot resulting from varying A_{1E} and A_{2E} (using the *steppar* command in *XSPEC*). The contours represent $\Delta\chi^2 = 2.3$ (black), 6.18 (red), 11.83 (green) and 19.33 (blue). These χ^2 levels correspond to 1σ , 2σ , 3σ and 4σ confidence for 2 degrees of freedom. We see that a fairly large part of parameter space can be ruled out with 4σ confidence. The null hypothesis model ($A_{1E} = A_{2E} = 0$) now has 6 more degrees of freedom than the best-fitting model, because the null hypothesis model is insensitive to ϕ_{1E} and ϕ_{2E} for orbit 1b, plus the same two parameters for the other data sets. We compare the best fit achieved by fixing $A_{1E} = A_{2E} = 0$ ($\chi^2 = 380.68/366$) with our global best

Table 1. Best-fitting line energy parameters for our joint fit. Errors are 1σ .

Parameter	Best fit
A_{1E} (keV)	$0.0446^{+0.023}_{-0.020}$
ϕ_{1E} (cycles)	$0.373^{+0.076}_{-0.13}$
A_{2E} (keV)	$0.119^{+0.026}_{-0.026}$
ϕ_{2E} (cycles)	$0.0497^{+0.019}_{-0.018}$
E_0 (keV)	$6.60^{+0.019}_{-0.018}$

fit ($\chi^2 = 354.24/360$) using an *f*-test, which rules out the null hypothesis with 3.70σ confidence.

Fig. 9 (right) shows χ^2 plotted against A_{1E} (black) and A_{2E} (red). The dashed line depicts $\Delta\chi^2 = 9$, which corresponds to 3σ for 1 degree of freedom. We see that A_{2E} in particular is fairly well constrained, with 3σ confidence limits of approximately $0.04 < A_{2E} < 0.21$. The best fit achieved when fixing $A_{1E} = 0$ has reduced $\chi^2 = 359.58/363$ and the best fit achieved with $A_{2E} = 0$ is $\chi^2 = 375.63/363$. Comparing these to the global best fit yields significances of 1.46σ for the first harmonic and 3.89σ for the second harmonic.

Fig. 10 (left) shows the probability map for all four variable spectral parameters for our final joint fit. Here, for parameters which are not tied across all data sets (such as ϕ_{1E} and ϕ_{2E}), we plot the values corresponding to orbit 1a. Note that the functions $E_{\text{line}}(\gamma)$, $N_G(\gamma)$ and $\Gamma(\gamma)$ are tied across *all* data sets except for orbit 1b. We see no statistically significant modulation in the iron line flux, but we do see a modulation in the power-law index which lags the line energy modulation by ~ 0.1 cycles. In Fig. 10 (right), we make the same plot for the case of orbit 1b. Even though the statistics are of course worse, the parameter modulations are strikingly different. The line energy and power-law index are both consistent with being constant, but the iron line flux varies with a large amplitude and high statistical significance. Clearly, there is something very different about orbit 1b. We have checked for proton flares, absorption events and various instrumental issues, but find no contribution from these effects, so are forced to conclude that this anomalous behaviour




Figure 9. χ^2 contour plots from our Fourier domain fits considering all data sets. We show the amplitude of the first and second harmonic of the line energy modulation A_{1E} and A_{2E} . Left: two-dimensional contour plot. The cross denotes the best fit, and the black, red, green and blue lines correspond to 1σ , 2σ , 3σ and 4σ confidence contours respectively for 2 degrees of freedom. Right: one-dimensional plot for A_{1E} (black) and A_{2E} (red). The grey dashed line is $\Delta\chi^2 = 9$ (3σ for 1 degree of freedom).




Figure 10. Probability maps from our Fourier domain fits for line energy (E_{line}), line flux (N_G), power-law index (Γ) and power-law normalization (N_{cont}). Significances for each parameter are as labelled and colours are as defined in Fig. 7. Left: the results of our joint fits, considering all data sets. For parameters which are not tied across all data sets (see the text for details), we plot the values corresponding to orbit 1a. Right: the results for only the anomalous orbit 1b. We see clear differences from the other data sets, with perhaps the most striking being the large modulation in iron line flux.

during orbit 1b is intrinsic to the source. We note that the iron line width is larger during orbit 1b (0.51 ± 0.05 keV) than for the other data sets combined (0.43 ± 0.02 keV). We can also see in the 4–10 keV power spectrum (Fig. 2) a slight increase in the amplitude of the fundamental from orbit 1a (black) to 1b (red), but a very slight decrease in the amplitude of the second harmonic. Also, the broad-band noise above ~ 2 Hz changes a little between orbit 1a and 1b. These differences may be indicative of there being a slightly different geometry during orbit 1b.

5 DISCUSSION

We have further developed the QPO phase-resolving method of IK15 and applied it to, in total, ~ 260 ks of *XMM-Newton* data and ~ 70 ks of *NuSTAR* data from the 2014 outburst of H1743–322. We measure a statistically significant (3.7σ) modulation of the iron line centroid energy with QPO phase by combining five independent data sets. We see in Fig. 7 that, for four of the five data sets, the line energy modulation has the same distinctive shape, with maxima at ~ 0.2 and ~ 0.7 QPO cycles. Surprisingly, one data set (*XMM-Newton* orbit 1b) does not show the same trend. Here we discuss the implications of the measured modulation and speculate as to why orbit 1b differs from the other data sets.

5.1 Interpretation: precession

Our result provides strong evidence that the Type-C QPO observed here is driven by systematic changes in the accretion geometry over the course of a QPO cycle. The only mechanism by which the line energy can vary without a geometric change is through shifts in the rest-frame line energy driven by changes in the disc ionization

state. An increase in disc ionization increases the iron line energy and *suppresses* the flux in the reflection hump relative to the line (e.g. Ross & Fabian 2005; García et al. 2013), in conflict with what we observe (Fig. 3). Also, the observed ~ 6.4 to ~ 6.8 keV change in line energy would require a factor of ~ 200 change in illuminating flux over a QPO cycle to originate purely from variations of disc ionization (see e.g. fig. 1 in Matt, Fabian & Ross 1993), which is implausible for all data sets except for orbit 1b, which show a change in line flux smaller than a factor of 2 (see Fig. 10, left). This indicates that the line energy variation is driven, at least in part, by changes in the relativistic distortions to the iron line profile, and therefore by a geometric variation over a QPO cycle. This ties in with recent population studies (Heil, Uttley & Klein-Wolt 2015; Motta et al. 2015) which show that systems observed with a more edge-on disc display systematically higher amplitude Type-C QPOs.

Shifts in the line energy are predicted to arise if the QPO originates from Lense–Thirring precession of the hot inner flow (Ingram et al. 2009; Ingram & Done 2012). As the inner flow precesses, it preferentially illuminates different disc azimuths, giving rise to a blue/redshifted iron line when the approaching/receding disc material is irradiated. For a geometry in which a single bright patch rotates about the disc surface, completing one cycle per QPO cycle, we would observe one maximum and one minimum in line energy per QPO cycle. Instead, we observe two maxima, neither of which coincides with a peak in continuum flux. This can be explained if we consider two bright patches rotating about the disc surface, as illustrated in Fig. 11. In this picture, the inner flow (orange) precesses, but the disc (grey) is held stationary by viscosity (Bardeen & Petterson 1975). The disc transitions into the hot inner flow at the truncation radius. In this schematic, the disc is irradiated by both the front and back of the flow (see the multi-coloured patches), as




Figure 11. Schematic representation of the precessing inner flow model. The inner flow (orange) extends out to $\sim 20\text{--}30 R_g$ and is misaligned with both the disc (grey) and black hole equatorial plane (horizontal). The flow precesses around the (vertical) black hole spin axis such that the front of the flow faces us in (a), to our left in (b) and so on. The front and back of the flow irradiate the disc, illustrated here by the multi-coloured patches. As the flow precesses, these irradiated patches rotate over the disc surface, prograde with disc orbital motion (white arrows). The colours of the irradiated patches encode energy shifts due to disc orbital motion and gravitational redshift.

we may expect to happen if the inner flow is sufficiently thin for its underside to be above the disc mid-plane (or for a very large misalignment between the disc and inner flow). The calculations of Ingram & Done (2012) considered an inner flow with very large vertical extent, and therefore only predicted one bright patch on the disc, as the underside of the flow was never above the disc mid-plane. The Doppler shifts experienced by photons reflected from respectively approaching and receding disc material are illustrated in Fig. 11 by the colour scheme of the irradiated patches. Precession of the flow as illustrated in Fig. 11 predicts a rocking of the iron line shape twice per precession cycle as different disc azimuths are illuminated first by the front of the flow, then half a cycle later by the back. The maximum line energy will occur when the approaching and receding sides of the disc are illuminated (Figs 11 b and c), since Doppler boosting means that the blueshifted part of the line (the so-called ‘blue horn’) will dominate over the redshifted part (the so-called ‘red wing’).

Non-relativistic precession mechanisms are unlikely. Classical precession is expected around an oblate spinning star but not for a black hole (Stella & Vietri 1998). Magnetic precession can result when the magnetic field of a spinning star intersects the accretion flow (Shirakawa & Lai 2002), but not for astrophysical black holes

which, without electric charge, have no way to generate their own magnetic field. Radiation pressure can cause variable warping in the outer disc through non-linear growth of perturbations, but only at disc radii $\gtrsim 160 R_g$ (Pringle 1996; Frank, King & Raine 2002), where orbital motion is too slow to explain the large observed energy shifts in the line. It is therefore likely that we are specifically witnessing Lense–Thirring precession. We note that Lense–Thirring precession of the reflector (the disc) rather than the illuminator (the inner flow) could potentially reproduce the observed line energy modulation (Schnittman et al. 2006; Tsang & Butsky 2013), although we note that the QPO modulates the power-law spectrum emergent from the inner flow much more strongly than the thermal disc emission visible at low energies. We also note that the observed line energy modulation could potentially result from a precessing jet (Kalamkar et al. 2016).

For Lense–Thirring precession of the entire inner flow, the precession period depends on the inner and outer radii of the inner flow, the radial surface density profile of the inner flow (Fragile et al. 2007; Ingram et al. 2009), as well as the mass and dimensionless spin parameter, $a = cJ/GM^2$, of the black hole. Assuming a constant surface density, a canonical black hole mass of $10 M_\odot$ and a spin of $a = 0.2$ (Steiner et al. 2012; Ingram & Motta 2014), the ~ 4 s period implies a truncation radius of $\sim 20\text{--}30 R_g$.

5.2 Implications

Lense–Thirring precession arises (due to the General Relativistic frame dragging effect) only in orbits with their rotational axis misaligned with the black hole spin axis. This may occur for accreting material in binary systems in which the black hole spin axis is misaligned with the axis of binary orbital motion (as the result of an asymmetric natal supernova kick; Fragos et al. 2010). Quite how the accretion flow reacts to this misalignment is a challenging theoretical question, which will be informed by our result. For a classical thin disc, the inner regions have long been thought to align with the black hole and the outer regions with the binary (Bardeen & Petterson 1975), but the location of the transition between orientations has remained uncertain. Recent simulations (Krolik & Hawley 2015) find this radius to be $\sim 8\text{--}9 R_g$, which is small enough to be within the disc truncation radius of $\sim 20\text{--}30 R_g$ indicated here by setting the precession period equal to the QPO period. This implies that the inner flow is being fed by material from the truncated disc out of the black hole equatorial plane (as in Fig. 11). Grid-based General Relativistic magnetohydrodynamic simulations of accretion flows in which the vertical extent is large compared with viscosity indicate that the entire hot inner flow can precess in this situation due to strong coupling through pressure waves (Fragile et al. 2007), in line with what is illustrated in Fig. 11. Alternatively, calculations using an α -prescription viscosity with a large misalignment angle and/or low viscosity show evidence for the disc breaking into discrete, independently precessing rings (Nixon & King 2012). This phenomenon has been seen in smoothed particle hydrodynamics simulations (Nixon et al. 2012; Nealon, Price & Nixon 2015), but not as yet in the grid-based simulations (Morales Teixeira et al. 2014; Zhuravlev et al. 2014). Such differential precession could also potentially give rise to the line energy shifts observed here, via the same mechanism of illumination of different disc azimuths. More sophisticated phase-resolved spectral modelling and additional high-quality data in future will allow tomographic mapping of the inner flow geometry, further informing numerical simulations.

Recently, van den Eijnden, Ingram & Uttley (2016) found evidence in observations of GRS 1915+105 that some form of

differential precession could indeed be at play (although likely not as extreme as that suggested by Nixon & King 2012). They show that, in observations displaying an energy-dependent QPO frequency (Qu et al. 2010; Yan et al. 2012), the phase of the band with the higher QPO frequency increases faster than that of the band with the lower QPO frequency. This confirms that the frequency difference is intrinsic to the source, and can be explained if, for example, the inner regions of the flow are precessing slightly faster than the outer regions. Although there is no energy dependence of the QPO frequency in the observations we analyse here, H1743–322 does show an energy dependence of the QPO frequency for observations with much higher ($\gtrsim 3$ Hz) QPO frequencies (Li et al. 2013).

Our result has implications for black hole spin measurements. Spin estimates obtained through disc spectral fitting often assume that the black hole spin aligns with the binary orbit (e.g. Kolehmainen & Done 2010; Steiner et al. 2012), which is incompatible with the precession model. Indeed, recent spectral modelling of Cygnus X-1 in the soft state implies a $\gtrsim 13^\circ$ misalignment (Tom-sick et al. 2014). The iron line method provides an independent measure of inclination, but assumes that the disc extends down to the ISCO, whereas the precession model assumes an evolving truncation radius. If the truncation radius really is moving, the shape of the line energy modulation should change with QPO frequency (Ingram & Done 2012), which can be tested in future. We also note that the spectral pivoting and line energy modulation detected here are *non-linear* changes in spectral shape, which could bias studies of the time-averaged spectrum. The biases are likely small, but should be quantified in future with tomographic modelling, since iron line fitting is sensitive to fairly small spectral distortions. For the case of active galactic nuclei (AGN), it is unclear if a misaligned accretion flow is expected in the absence of a binary partner.² If there is precession in AGN, it will not create a bias through non-linear variability, since the precession time-scale would be longer than a typical integration time.

5.3 Alternative interpretations

As an alternative to precession, axisymmetric variations in the accretion geometry can cause changes in the iron line shape. Since the disc rotational velocity and gravitational redshift both depend on radius, variation of the disc inner radius throughout a QPO cycle can cause shifts in the line energy. For the same reasons, changes in the radial dependence of disc irradiation, perhaps caused by changes in the vertical extent of the illuminating source, can also drive changes in the line shape. However, it is very difficult to explain how such mechanisms could give rise to two maxima in line energy per QPO cycle. None the less, in future we will explicitly test the precession model described above against the data presented here, and compare it to simple axisymmetric alternatives.

De Marco & Ponti (2016) recently suggested that the soft lag measured in the 0.1–1 Hz frequency range for the *XMM–Newton* data is a reverberation lag corresponding to a $\sim 100 R_g$ path length. However, this frequency range is dominated by the QPO. Both soft and hard lags are routinely observed for QPOs (e.g. Qu et al. 2010), and the QPO lag is often very different from that measured for the broad-band noise (e.g. Wijnands et al. 1999), and not compatible with a reverberation lag (Stevens & Uttley 2016). Moreover, in our Fig. 6 we show that the two QPO harmonics have different

lags, and so averaging them together has little physical meaning. Reverberation lags are still expected to be present of course, but will yield much smaller soft lags than the QPO.

5.4 Anomalous data set: orbit 1b

As for the anomalous data set, orbit 1b, this is puzzling in the context of any QPO model. The modulations in line energy and flux and also power-law index are consistent between all the other data sets. Orbit 1b shows different modulations in all three of these parameters,³ as can be seen in Fig. 10. The most striking is perhaps the large-amplitude, and highly statistically significant (4σ), modulation in the line flux in orbit 1b. This may be indicative of a different geometry during orbit 1b. Such a geometrical change needs to explain the increased iron line flux, the increased variability in line flux and also the increased width of the iron line (0.51 keV for orbit 1b and ~ 0.43 keV for the other observations). It also needs to be consistent with the only subtle differences in other diagnostics (such as the full band power spectrum and the time-averaged power-law photon index) and the change needs to plausibly happen over a $\lesssim 60$ ks time-scale. The increased line flux implies a greater fraction of continuum photons intercept the disc, which will broaden the line somewhat by increasing the disc ionization. The increased variability in line flux suggests that this fraction *varies* more than for the other data sets. This could occur if the misalignment angle between the disc and the black hole spin axes, β , is somehow larger, since the misalignment between the disc and inner flow varies between 0 and 2β in the precession model (Veledina, Poutanen & Ingram 2013; Ingram et al. 2015). This extra variability in illuminating photons could make line energy variations due to ionization changes significantly more important than for the other data sets. We see in Fig. 10 (right, second panel) that the line flux, and therefore the flux irradiating the disc, varies by a factor of ~ 8 over a QPO cycle for orbit 1b. This means that the ionization parameter ($\xi \propto$ illuminating flux) should also vary by a factor of 8. In fig. 1 of Matt, Fabian & Ross (1993), we see that varying the ionization parameter from $\xi \sim 100$ to $\xi \sim 800$ changes the line rest-frame energy from ~ 6.4 to ~ 6.7 keV. This modulation in the rest-frame line energy should be in phase with the line flux, and therefore in anti-phase with the line energy modulation seen in the other data sets. It is unfortunate that *NuSTAR* was not observing during orbit 1b, otherwise this hypothesis could have been tested by tracking the reflection hump. Alternatively (or perhaps additionally), our view may be obstructed by some material in our line of sight during orbit 1b, which is plausible given the likely high inclination of H1743–322. The variable illumination of the line-of-sight material will give rise to variable ionization, which will imprint itself on to the phase-resolved spectra.

6 CONCLUSIONS

We find that the iron line centroid energy in H1743–322 is modulated on the QPO period with a statistical significance of 3.7σ . We also find that this modulation has a non-zero second harmonic with a statistical significance of 3.94σ . Shifts of the line energy over a QPO cycle are a distinctive prediction of the Lense–Thirring precession model (Ingram et al. 2009), in which the inner accretion

² Also, it is notoriously difficult to detect a Type-C QPO analogue due to the very long period expected through mass scaling (Vaughan & Uttley 2005).

³ The power-law normalization is trivially very similar across all data sets, because QPO phase is defined from the reference band flux, which tracks the power-law normalization to a good approximation, given that the power-law index varies only with small amplitude.

flow precesses due to the frame dragging effect. Our observation is a typical example of a Type-C QPO, implying that this class of QPOs in general are driven by Lense–Thirring precession, and therefore supporting studies that measure black hole mass and spin using the period of the Type-C QPO in combination with that of high-frequency QPOs (Ingram & Motta 2014; Motta et al. 2014; Fragile, Straub & Blaes 2016). There are still, however, unanswered questions. We have simply employed phenomenological modelling to track the iron line here, but more physical modelling using a self-consistent reflection model will provide further insight. We will perform this modelling in a future paper, as well as testing alternative models to precession. The largest question mark concerns the anomalous data set, orbit 1b, which exhibits different parameter modulations to all other data sets (which all agree with one another).

In future, high-quality observations of the same source displaying a QPO with a higher frequency will provide further insight. The precession model predicts the disc inner radius to be smaller for higher QPO frequencies, and therefore we expect the line energy dependence on QPO phase to have a different shape. Studies such as this will be greatly enhanced by new instrumentation. Detectors with a very large collecting area will allow us to perform similar studies without needing to stack over very long exposures as is necessary here. Also, X-ray polarimetry will provide an extra dimension, particularly when combined with phase-resolved spectroscopy (Ingram et al. 2015). The precession model predicts that the polarization angle changes with QPO phase, and that the extrema in polarization angle coincide with maxima in the line energy.

ACKNOWLEDGEMENTS

AI acknowledges support from the Netherlands Organisation for Scientific Research (NWO) Veni Fellowship, grant number 639.041.437. He thanks Martin Heemskerk for assistance making the plots in this paper. He also acknowledges useful conversations with Dom Walton, Matteo Bachetti and Brian Grefenstette about *NuSTAR* data reduction, and useful conversations with Andrew King about disc warping. MJM appreciates support via an STFC Ernest Rutherford Fellowship. CD acknowledges support from the STFC consolidated grant ST/L00075X/1. DA acknowledges support from the Royal Society. MA is an International Research Fellow of the Japan Society for the Promotion of Science. We thank the anonymous referee for constructive comments.

REFERENCES

Arnaud K. A., 1996, in Jacoby G. H., Barnes J., eds, ASP Conf. Ser. Vol. 101, Astronomical Data Analysis Software and Systems V. Astron. Soc. Pac., San Francisco, p. 17
 Axelsson M., Hjalmarsdotter L., Done C., 2013, MNRAS, 431, 1987
 Bachetti M. et al., 2015, ApJ, 800, 109
 Bardeen J. M., Petterson J. A., 1975, ApJ, 195, L65
 Cabanac C., Henri G., Petrucci P.-O., Malzac J., Ferreira J., Belloni T. M., 2010, MNRAS, 404, 738
 Chakrabarti S. K., Molteni D., 1993, ApJ, 417, 671
 De Marco B., Ponti G., 2016, preprint ([arXiv:1605.00765](https://arxiv.org/abs/1605.00765))
 Done C., Diaz Trigo M., 2010, MNRAS, 407, 2287
 Done C., Gierlinski M., Kubota A., 2007, A&A, 15, 1
 Esin A. A., McClintock J. E., Narayan R., 1997, ApJ, 489, 865
 Fabian A. C., Rees M. J., Stella L., White N. E., 1989, MNRAS, 238, 729
 Foreman-Mackey D. et al., 2013, Astrophysics Source Code Library, record ascl:1303.002

Fragile P. C., Blaes O. M., Anninos P., Salmonson J. D., 2007, ApJ, 668, 417
 Fragile P. C., Straub O., Blaes O., 2016, preprint ([arXiv:1602.08082](https://arxiv.org/abs/1602.08082))
 Fragos T., Tremmel M., Rantsiou E., Belczynski K., 2010, ApJ, 719, L79
 Frank J., King A., Raine D. J., 2002, Accretion Power in Astrophysics, 3rd edn. Cambridge Univ. Press, Cambridge
 Galeev A. A., Rosner R., Vaiana G. S., 1979, ApJ, 229, 318
 García J., Dauser T., Reynolds C. S., Kallman T. R., McClintock J. E., Wilms J., Eikmann W., 2013, ApJ, 768, 146
 Haardt F., Maraschi L., 1991, ApJ, 380, L51
 Harrison F. A. et al., 2013, ApJ, 770, 103
 Heil L. M., Uttley P., Klein-Wolt M., 2015, MNRAS, 448, 3348
 Ichimaru S., 1977, ApJ, 214, 840
 Ingram A., Done C., 2012, MNRAS, 427, 934
 Ingram A., Motta S., 2014, MNRAS, 444, 2065
 Ingram A., van der Klis M., 2015, MNRAS, 446, 3516 (IK15)
 Ingram A., Done C., Fragile P. C., 2009, MNRAS, 397, L101
 Ingram A., Maccarone T. J., Poutanen J., Krawczynski H., 2015, ApJ, 807, 53
 Jansen F. et al., 2001, A&A, 365, L1
 Kalamkar M., Casella P., Uttley P., O’Brien K., Russell D., Maccarone T., van der Klis M., Vincentelli F., 2016, MNRAS, preprint ([arXiv:1510.08907](https://arxiv.org/abs/1510.08907))
 Kolehmainen M., Done C., 2010, MNRAS, 406, 2206
 Krolik J. H., Hawley J. F., 2015, ApJ, 806, 141
 Lense J., Thirring H., 1918, Phys. Z., 19, 156
 Li Z. B., Zhang S., Qu J. L., Gao H. Q., Zhao H. H., Huang C. P., Song L. M., 2013, MNRAS, 433, 412
 Markoff S., Nowak M. A., Wilms J., 2005, ApJ, 635, 1203
 Matt G., Fabian A. C., Ross R. R., 1993, MNRAS, 262, 179
 Miller J. M., Homan J., 2005, ApJ, 618, L107
 Morales Teixeira D., Fragile P. C., Zhuravlev V. V., Ivanov P. B., 2014, ApJ, 796, 103
 Motta S. E., Belloni T. M., Stella L., Muñoz-Darias T., Fender R., 2014, MNRAS, 437, 2554
 Motta S. E., Casella P., Henze M., Muñoz-Darias T., Sanna A., Fender R., Belloni T., 2015, MNRAS, 447, 2059
 Nealon R., Price D. J., Nixon C. J., 2015, MNRAS, 448, 1526
 Nixon C. J., King A. R., 2012, MNRAS, 421, 1201
 Nixon C., King A., Price D., Frank J., 2012, ApJ, 757, L24
 Novikov I. D., Thorne K. S., 1973, in Dewitt C., Dewitt B. S., eds, Black Holes (Les Astres Occlus). Gordon and Breach, New York, p. 343
 Poutanen J., Krolik J. H., Ryde F., 1997, MNRAS, 292, L21
 Pringle J. E., 1996, MNRAS, 281, 357
 Qu J. L., Lu F. J., Lu Y., Song L. M., Zhang S., Ding G. Q., Wang J. M., 2010, ApJ, 710, 836
 Ross R. R., Fabian A. C., 2005, MNRAS, 358, 211
 Schnittman J. D., Homan J., Miller J. M., 2006, ApJ, 642, 420
 Shakura N. I., Sunyaev R. A., 1973, A&A, 24, 337
 Shirakawa A., Lai D., 2002, ApJ, 564, 361
 Steiner J. F. et al., 2012, MNRAS, 427, 2552
 Stella L., Vietri M., 1998, ApJ, 492, L59
 Stella L., Vietri M., Morsink S. M., 1999, ApJ, 524, L63
 Stevens A. L., Uttley P., 2016, MNRAS, preprint ([arXiv:1605.01753](https://arxiv.org/abs/1605.01753))
 Stiele H., Yu W., 2016, MNRAS, preprint ([arXiv:1602.01550](https://arxiv.org/abs/1602.01550))
 Sunyaev R. A., Truemper J., 1979, Nature, 279, 506
 Tagger M., Pellat R., 1999, A&A, 349, 1003
 Thorne K. S., Price R. H., 1975, ApJ, 195, L101
 Tomsick J. A. et al., 2014, ApJ, 780, 78
 Tsang D., Butsky I., 2013, MNRAS, 435, 749
 Uttley P., Cackett E. M., Fabian A. C., Kara E., Wilkins D. R., 2014, A&AR, 22, 72
 van den Eijnden J., Ingram A., Uttley P., 2016, MNRAS, 458, 3655
 van der Klis M., 1989, ARA&A, 27, 517
 van der Klis M., 2006, Adv. Space Res., 38, 2675
 Vaughan S., Uttley P., 2005, MNRAS, 362, 235
 Veledina A., Poutanen J., Ingram A., 2013, ApJ, 778, 165

- Wagoner R. V., Silbergleit A. S., Ortega-Rodríguez M., 2001, *ApJ*, 559, L25
- Wijnands R., Homan J., van der Klis M., 1999, *ApJ*, 526, L33
- Wilkinson T., Uttley P., 2009, *MNRAS*, 397, 666
- Wilms J., Allen A., McCray R., 2000, *ApJ*, 542, 914
- Yan S.-P. et al., 2012, *Ap&SS*, 337, 137
- Zhuravlev V. V., Ivanov P. B., Fragile P. C., Morales Teixeira D., 2014, *ApJ*, 796, 104

APPENDIX A: DATA VISUALIZATION

In order to create the probability maps shown in Figs 7 and 10, we run an MCMC in *XSPEC* after finding a best-fitting model in the Fourier domain. *XSPEC* uses the *emcee* algorithm (the MCMC hammer; Foreman-Mackey et al. 2013). We use the Goodman–Weare algorithm with a chain length of 3×10^5 steps and 10^3 walkers. The starting point of the chain is a randomized realization of the best-fitting parameters. Visual inspection of the χ^2 implies that the chain takes $\sim 2 \times 10^4$ steps to converge, so we burn 2.5×10^4 steps. For the rest of the chain, the autocorrelation function of the parameters of interest is centrally peaked, indicating reasonable convergence. Even so, we note that none of our significances or error estimates use these chains, we use them purely for data visualization.

For the probability maps in Figs 7 and 10, we calculate the $E_{\text{line}}(\gamma)$ function for each step of the chain, for 400 values of γ . That is, for each step of the chain, we read in the parameters A_{1E} , A_{2E} , ϕ_{1E} , ϕ_{2E} and E_0 for that step and calculate $E_{\text{line}}(\gamma)$ from equation (6). For each γ value, we thus have 2.75×10^5 values of E_{line} , which we bin into an 800-bin histogram. Fig. A1 shows these histograms for the chain corresponding to our joint fit, for two selected QPO phases




Figure A1. Histograms for the line energy for two QPO phases (red: phase=7/16 cycles, blue: phase=11/16 cycles) created using an MCMC. The black lines are smoothed versions of these histograms (see the text for details).

(red: phase=7/16 cycles, blue: phase=11/16 cycles). We normalize each histogram to peak at unity. For plotting purposes, we smooth these histograms by averaging each of the 800 bins with the ± 10 bins either side. The black lines show the smoothed versions of the histograms. We use the smoothed versions for our probability maps.

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