Investigating hot-Jupiter inflated radii with hierarchical Bayesian modelling

Marko Sestovic^{1,2}, Brice-Olivier Demory¹, and Didier Queloz^{2,3}

¹ Center for Space and Habitability, University of Bern, Gesellschaftsstrasse 6, Bern 3012, Switzerland
e-mail: marko.sestovic@csh.unibe.ch

² Astrophysics Group, Cavendish Laboratory, J. J. Thomson Avenue, Cambridge CB3 0HE, UK

³ Observatoire de Genève, Université de Genève, 51 chemin des Maillettes, 1290 Sauverny, Switzerland

Received 27 June 2017 / Accepted 27 March 2018

ABSTRACT

Context. As of today, hundreds of hot Jupiters have been found, yet the inflated radii of a large fraction of them remain unexplained. A number of mechanisms have been proposed to explain these anomalous radii, however most of these can only work under certain conditions and may not be sufficient to explain the most extreme cases. It is still unclear whether a single mechanism can sufficiently explain the entire distribution of radii, or whether a combination of these mechanisms is needed.

Aims. We seek to understand the relationship of radius with stellar irradiation and mass and to find the range of masses over which hot Jupiters are inflated. We also aim to find the intrinsic physical scatter in their radii, caused by unobservable parameters, and to constrain the fraction of hot Jupiters that exhibit inflation.

Methods. By constructing a hierarchical Bayesian model, we inferred the probabilistic relation between planet radius, mass, and incident flux for a sample of 286 gas giants. We separately incorporated the observational uncertainties of the data and the intrinsic physical scatter in the population. This allowed us to treat the intrinsic physical scatter in radii, due to latent parameters such as the heavy element fraction, as a parameter to be inferred.

Results. We find that the planetary mass plays a key role in the inflation extent and that planets in the range $\sim 0.37\text{--}0.98 M_J$ show the most inflated radii. At higher masses, the radius response to incident flux begins to decrease. Below a threshold of $0.37 \pm 0.03 M_J$ we find that giant exoplanets as a population are unable to maintain inflated radii $\gtrsim 1.4 R_J$ but instead exhibit smaller sizes as the incident flux is increased beyond 10^6 W m^{-2} . We also find that below $1 M_J$, there is a cut-off point at high incident flux beyond which we find no more inflated planets, and that this cut-off point decreases as the mass decreases. At incident fluxes higher than $\sim 1.6 \times 10^6 \text{ W m}^{-2}$ and in a mass range $0.37\text{--}0.98 M_J$, we find no evidence for a population of non-inflated hot Jupiters. Our study sheds a fresh light on one of the key questions in the field and demonstrates the importance of population-level analysis to grasp the underlying properties of exoplanets.

Key words. planets and satellites: fundamental parameters – planets and satellites: atmospheres – methods: statistical

1. Introduction

Understanding the internal state and energy processes of giant planets is a key goal in exoplanet science. Gas giant interiors are generally modelled by assuming an adiabatic temperature profile and internal structure and calculating the evolution of the planet radius as it contracts (for a review see e.g. Fortney & Nettelmann 2010; Fortney et al. 2010; Baraffe et al. 2014).

Fortney et al. (2007) and Baraffe et al. (2008) modelled the evolution of planetary radii by including the effects of stellar irradiation and heavy element cores with state-of-the-art internal equations of state. Such model predictions can be tested against the currently observed population of exoplanets with known radii, masses, and orbital distances. At low stellar irradiation, these models are thought to agree with observations, but the high irradiation regime still presents a challenge. At fluxes greater than $\sim 2 \times 10^5 \text{ W m}^{-2}$ (Miller & Fortney 2011; Demory & Seager 2011), a sizeable fraction of gas giants are found to be anomalously larger than predicted; these gas giants include hot Jupiters WASP-17 b, WASP-121 b, and Kepler-435 b, which all have measured radii $R > 1.8 R_J$ (Anderson et al. 2011; Almenara et al. 2015; Delrez et al. 2016).

There are several candidate explanations for this inflation process, including tidal dissipation (Bodenheimer et al. 2001, 2003; Arras & Socrates 2010; Jermyn et al. 2017), kinetic heating (Guillot & Showman 2002), enhanced atmospheric opacities (Burrows et al. 2007), double diffusive convection (Chabrier & Baraffe 2007; Kurokawa & Inutsuka 2015), vertical advection of potential temperature (Youdin & Mitchell 2010; Tremblin et al. 2017), and Ohmic heating through magnetohydrodynamic effects (Batygin & Stevenson 2010; Perna et al. 2010; Wu & Lithwick 2013; Ginzburg & Sari 2016).

In most cases to date, theoretical predictions are tested against individual planets or by producing radius predictions for a narrow set of parameters. However, the radii may depend on many parameters, including some that are non-observable or poorly constrained (latent parameters) such as the core mass, system age, internal composition, and atmospheric opacity. The degeneracy between these parameters and the strength of the inflationary process limits the information we can obtain from a single planet. Thanks to the sheer count of exoplanets discovered so far, we are now in a position to study the issue of inflated radii using population-level analyses.

In such an attempt, the choice of forward model is central. The candidate model should not only reproduce a range of individual planet observations, but also predict the full distribution of the observable planet radii and their dependence on planet mass and stellar incident flux within reasonable assumptions on the latent parameters. Laughlin et al. (2011) already used the observed radii of 90 well-characterised transiting exoplanets to assess the incidence of Ohmic heating on the inflated radius anomaly (first introduced by Guillot et al. 2006). Using a population-level approach, one can also extract model-independent conclusions from the data. For example, Enoch et al. (2012) used multivariate regression models to investigate the factors that affect the radii of 119 extrasolar gas planets such as planetary equilibrium temperature, planet mass, stellar metallicity, and tidal heating rate.

To that end, it is important to separate the real physical scatter caused by the latent parameters and the scatter caused by measurement uncertainties (see Wolfgang et al. 2016 in which this was addressed for super-Earths and sub-Neptunes). For example, hot-Jupiter radii span from $\sim 1.99 R_J$ (Kepler-435 b; Delrez et al. 2016) to $\sim 0.84 R_J$ (Kepler-41 b; Santerne et al. 2011). We seek to understand how much of this spread can be accounted for by the measurement uncertainty in the radius and thus constrain the true range of radii that is driven by the currently unidentified physical process.

In this work, we infer the flux-mass-radius (FMR) relation using a population of 286 transiting gas giants with measured masses and radii. We use a very similar approach to Wolfgang et al. (2016) and Shabram et al. (2016), by employing a hierarchical Bayesian model (HBM) in which we separate the effects of intrinsic scatter and measurement uncertainty (see also Kelly 2007; Hogg et al. 2010; Demory 2014; Rogers 2015). This also allows us to include the uncertainty in both the dependent parameter (planet radius) and the independent parameters (planet mass and incident flux), which cannot be carried out with basic least-squares regression. We thus constrain the FMR relation as a distribution rather than a deterministic function, and characterise the physical scatter of the gas giant population.

It is as of yet unknown whether all hot Jupiters are inflated or if some may have radii that agree with non-inflationary models (Fortney et al. 2007). Many of the candidate inflation mechanisms, for example, tidal heating for circularised orbits (Gu et al. 2003), would not apply for all planets. Similarly, Heng (2012) and Perna et al. (2012) showed that the depth and degree of Ohmic heating are sensitive to numerous atmospheric parameters such as opacity or thermal inversions and may be inhibited in some cases. Measuring the FMR distribution and its intrinsic scatter thus provides a way to determine whether such non-anomalous planets exist.

We present the planet data used in this study in Sect. 2 and present our parameterisation of the FMR distribution in Sect. 3. The methods for fitting the parameters and selecting our model are explained in Sect. 4, and the results of our fit for the FMR distribution are shown in Sect. 5. We test the agreement of our model with the data in Sect. 6, and we finally discuss the implications of our results in Sect. 7.

2. Data

As we require both the mass and radius of the planets, we include only transiting giants that have radial velocity (RV) or transit timing variation (TTV) mass measurements. Of those, nine planets have masses derived from TTVs, which are biased towards

lower fluxes and masses, and have uncertainties that are larger on average than for the whole sample¹. We include these planets because they populate a relatively sparse region of the parameter space, however we note that they may have a different selection bias than the rest of the sample (see e.g. Wolfgang et al. (2016) in which such a bias was found for sub-Neptune planets). We also use the stellar temperature, stellar radius, and the semi-major axis of the planetary orbit to calculate a distribution for the incident flux on the planet (see Sect. 4).

The data was taken from the Exoplanet Orbit Database (EOD; Han et al. 2014) and supplemented by the NASA Exoplanet Archive (Akeson et al. 2013), both last accessed on 10 October 2016. We checked the archive and the source paper in cases in which the EOD data is incomplete. Planets that are missing required values in both databases were discarded. We also took information about the planet ages and metallicities from the exoplanet.eu database (Schneider et al. 2011). Where the raw data have asymmetric error bars, we averaged these data and took the uncertainty distribution to be Gaussian. This follows the treatment of Weiss et al. (2013) and Wolfgang et al. (2016), since we did not have access to the full likelihood distributions for all our parameters. While this may introduce a bias into our results, a short test (see Sect. 7.1) suggested the effect on our hyperparameters is well within our posterior distribution uncertainties.

Our main criterion for selecting our sample is the planet mass. We set $13 M_J$ as the upper cut-off to exclude planets where deuterium burning is thought to happen (Chabrier et al. 2005; Spiegel et al. 2011). Choosing the lower bound is related to the wider problem of whether there are any strong boundaries in gas-giant characteristics, which is one of the questions we seek to answer. Enoch et al. (2012) chose 0.1 – $0.5 M_J$ to define Saturn-mass planets and 0.5 – $2.0 M_J$ for Jupiter-like planets, finding that Saturn-mass planets do not exhibit any inflation with temperature. Laughlin et al. (2011) also used $0.1 M_J$ as a lower bound for their sample.

We chose $0.1 M_J$ as an absolute lower bound with the understanding that this introduces planets into our model that do not have the same properties as the rest of our sample, i.e. Saturn and Neptune-like planets that may not exhibit inflation, as shown by Enoch et al. (2012). Finding mass boundaries between the hot-Jupiter planets that inflate and those that do not is part of our HBM implementation (see Sects. 3 and 4).

3. Forward model

In this section, we review the known trends and present our main parametrisation for the FMR distribution. We refer to the following as the baseline model. Our current understanding of giant-planet structure and evolution is that at low incident flux, thermal evolution models can reproduce the range of planet radii present in the exoplanet population (Miller & Fortney 2011; Thorngren et al. 2016). Although previous models have predicted that radii depend on stellar irradiation, particularly at fluxes of $\sim 1 \times 10^5 \text{ W m}^{-2}$ and greater ($\sim 0.1 \text{ AU}$), it is also clear that a large percentage of gas giants in short-period orbits are much more inflated than predicted (see Fortney et al. 2007, 2010; Baraffe et al. 2008, 2014; Fortney & Nettelmann 2010). Demory & Seager (2011) and Miller & Fortney (2011) showed that these inflated hot Jupiters start to occur at incident fluxes of greater than

¹ The average fractional uncertainty is 8.8% for the RV masses and 15.0% for the TTV masses.

Fig. 1. *Top panel:* distribution of the transit radii in our sample vs. the incident flux. The mass is colour coded. The radii begin to show a correlation with flux at $\sim 2 \times 10^5 \text{ W m}^{-2}$, corresponding to where inflated hot Jupiters begin to appear (Demory & Seager 2011). *Bottom panel:* the transit radii plotted against mass, colour coded by incident flux.

$\sim 2 \times 10^5 \text{ W m}^{-2}$. Enoch et al. (2012) and Laughlin et al. (2011) also found that the radii of these hot gas giants can depend strongly on irradiation temperature. These trends can be seen in the top of Fig. 1.

Taking this into account when choosing the parametrisation of our empirical flux-mass-radius relation (represented as the forward model for radius, R , given mass and flux, M and F), we modelled the effect of stellar irradiation by considering two regimes. For planets below a threshold F_s in incident flux, the radii are taken to be constant at size C , and independent of flux. Above F_s the radii increase proportionally with $\log F$. This relation (Eq. (1)) qualitatively agrees with the distribution of radii and fluxes in Fig. 1 and is written as

$$\frac{R}{R_J} = \begin{cases} C, & F < F_s, \\ C + A \cdot (\log F - \log F_s), & F \geq F_s. \end{cases} \quad (1)$$

While planetary radii may depend on incident flux even at $F < F_s$, this dependence is very weak, and variations in age, core mass, and envelope composition are expected to play a significantly greater role; see for example Fortney et al. (2007); Baraffe et al. (2008); and Miller & Fortney (2011). Additionally, because of a lack of transiting planets with known masses found at large orbital distances (there are only 22 planets further out than 0.1 AU in our sample), combined with the observational uncertainty of the measurements, it is unlikely that we could extract

meaningful constraints about the dependence of planetary radius with incident flux for gas giants at low irradiation.

Equation (1) provides a deterministic way of modelling the planetary radius based on the incident flux and model parameters, which we must infer (A , C , and F_s). However, owing to non-observable parameters such as core mass, composition, and age, we expect to see a physical scatter in the radii at a given flux level. We include this uncertainty by attaching to the planet radius a Gaussian distribution of standard deviation σ_R , around a mean value μ_R , which is given the following same form as Eq. (1):

$$R \sim N(\mu = \mu_R, \sigma = \sigma_R), \quad (2)$$

$$\mu_R(F, M) = \begin{cases} C(M), & F < F_s(M) \\ C(M) + A(M) \cdot \log_{10} \frac{F}{F_s(M)}, & F \geq F_s(M) \end{cases}, \quad (3)$$

where $\sim N(\mu, \sigma)$ means drawn from a normal distribution with mean μ and variance σ^2 . Enoch et al. (2012) found that there were large differences in the radius-temperature relations between planets of different masses. We thus include the effect of planet mass in our model by splitting our total mass range into J separate bins, each with a different radius-flux relation of the same form as Eq. (2), but with independent model parameters (C , F_s , and A) as follows:

$$A(M) = \begin{cases} A_1, & \frac{M}{M_J} < m_1. \\ \dots \\ A_j, & m_{j-1} \leq \frac{M}{M_J} < m_j, \\ \dots \\ A_J, & m_{J-1} \leq \frac{M}{M_J}. \end{cases} \quad (4)$$

where $C(M)$ and $F_s(M)$ follow the same form and $\{m_j\}_{j=1}^{J-1}$ represent the inner boundaries separating the mass bins. Unlike previous studies, we do not fix the mass boundaries to be constant, as their placement would be arbitrary and could bias the FMR relation. Instead, the mass bin boundaries are also model parameters that we infer.

Our final baseline model uses a mass- and flux-dependent scatter, $\sigma_R(M, F)$, with three possible values for three classes of planets: planets in the lowest mass bin ($M < m_1$), and inflated and weakly irradiated planets ($F \geq F_s$ and $F < F_s$ respectively) shared for all higher mass bins ($M \geq m_1$), written as:

$$\sigma_R(F, M) = \begin{cases} \sigma_{R,1}, & \frac{M}{M_J} < m_1, \\ \sigma_{R,2}, & \frac{M}{M_J} \geq m_1 \text{ and } F < F_s, \\ \sigma_{R,3}, & \frac{M}{M_J} \geq m_1 \text{ and } F \geq F_s. \end{cases} \quad (5)$$

The baseline model, with 4 mass bins and 17 parameters, was chosen after a process of model selection, in which we considered several alternative parametrisations and compared their posterior results. We varied the number of mass bins, choosing between 3, 4, and 5 bins, as well as the form of σ_R and F_s . For the flux threshold, we compared with the case in which the threshold was the same for all mass bins. For the intrinsic scatter, we also set it to depend on just the mass instead, or on both mass and incident flux, and we tried models in which the intrinsic scatter was the same for all fluxes and masses (as a constant single parameter). We did this by considering which models converged to unimodal posteriors and with tests on our posterior predictive fit, which are described in Sect. 6.

Fig. 2. Our model shown as a Bayesian network. The forward model hyperparameters are indicated in light orange (the third level of the HBM), the true planet parameters in white (the second level of the HBM), and the fixed observed values in dark blue. The arrows denote conditional dependencies; the probability distribution of a child parameter is dependent on the values of its parent parameters (arrows go from parent to child).

4. Hierarchical Bayesian model

Our aim is to infer the values of the model parameters² of the distribution for R given F and M , or $R|F, M$. We also refer to the parameters by shorthand, for example $\mathbf{A} = \{A_j\}_{j=0}^J$, or all of the parameters together as $\alpha = (A, C, F_s, m, \sigma_R)$. These are not necessarily vectors however.

The probabilistic model described in Sect. 3 defines the relationship between the true parameters R , M , and F of a planet. However when an observation is made, for example measuring the transit radius $\tilde{R}$, the result is subject to noise. It can be taken as being drawn from a probability distribution that depends on the true radius R and the noise characteristics of the observation. This is represented as $\tilde{R} \sim p(\tilde{R}|R, \alpha_{\text{obs}})$, where α_{obs} parametrises the noise of a particular observation. We take care to differentiate between the observed value $\tilde{R}$ and the true value R . In the most simple case, $\tilde{R}$ is drawn from a Gaussian distribution with mean R and a standard deviation, which is the observational uncertainty, σ_{obs} . As shorthand, we refer to the combined observational data as $\tilde{\mathbf{x}} = \{\tilde{R}_i, \tilde{M}_i, \tilde{F}_i\}_{i=1}^N$.

The distribution of our observables $(\tilde{R}, \tilde{M}, \tilde{F})$ for a given planet is the distribution of the true parameters (R, M, F) convolved with the observational uncertainty distribution. Thus the true scatter of the radii, σ_R , is lower than the scatter of our observed data, and the posterior values for our parameters (R, M, F) are less spread out than the data (see Figs. 6 and 7). This is a general consequence in cases in which quantities drawn from the same distribution (the true values) are subject to an additional scatter; see Stein (1956) and Good (1965).

Treating this problem in a HBM framework allows us to separately incorporate both the measurement errors in all the parameters and the intrinsic scatter in the radius (which is a parameter to be inferred). Wolfgang et al. (2016), for example provided a brief overview of the advantages of HBMs for a

very similar problem. By defining the probabilistic relations and placing priors on all the parameters, we can calculate the joint posterior distribution for R , F , M , and α , given the data. The forward model relation for $R_i|M_i, F_i, \alpha$ has already been defined in Sect. 3; in this section, we define the following other probabilistic relations and priors (with their dependence relationships shown in Fig. 2):

$$2 \log_e \sigma_{R,j} \sim U(-10, 4)$$

$$m_j \sim U(m_{j-i}, m_{j+1})$$

$$\log_{10} F_{s,j} \sim N(5.3, 1)$$

$$C_j \sim U(0, 2)$$

$$\arctan A_j \sim U(-\frac{\pi}{2}, \frac{\pi}{2})$$

$$M_i \sim U(0, \infty)$$

$$\tilde{R}_i \sim N(R_i, \sigma_{R,\text{obs},i})$$

$$\tilde{M}_i \sim N(M_i, \sigma_{M,\text{obs},i})$$

$$F_i \sim SN(\mu_{F,i}, \sigma_{F,\text{obs},i}, \alpha_{F,i})$$

$$R_i|M_i, F_i, \alpha \sim N(\mu_R(M_i, F_i, \alpha), \sigma_R),$$

where $U(a, b)$ represents a uniform distribution with upper and lower bounds a and b , respectively, and $SN(\mu, \sigma, \alpha)$ represents a skew-normal with the skewness parameter α (Azzalini 1985).

We use an uninformative prior for A_j (a Jeffreys prior for the slope of a line, see VanderPlas 2014) and essentially uninformative priors for $\sigma_{R,j}$ and C_j . The cut-offs on $\sigma_{R,j}$ and C_j are placed at values beyond which we can assume to have negligible likelihood, although this is arbitrary. An average cold radius of greater than $2R_J$ or less than zero is clearly not supported by the data, and a scatter of less than $0.007R_J$ is much less than expected based on possible variations in core mass and planet age. The posteriors are not found to have any significant probability density near these limits, so our choice is justified.

The mass bin boundaries are given a uniform prior that is constrained to prevent these bin boundaries from crossing each other. There are no planets with mass less than $0.1M_J$ in our sample (see Sect. 2), thus we place a minimum cut-off for the lowest bound at $0.1M_J$. Among the model parametrisations that are tried, we find that in models with four bins (i.e. 3 boundaries), the highest boundary generally has a very wide distribution and would occasionally get stuck outside our mass range. However, we also find that when we fix the highest mass bin boundary to an arbitrary value larger than $2.0M_J$, we still obtain clearly distinct behaviour in each of the four mass bins (in terms of the radius response to inflation). In the interest of extracting the maximum amount of information from the data, we thus use four mass bins, but with the highest boundary (m_3) fixed at $2.5M_J$.

The prior for F_s is based on the results of Demory & Seager (2011); and Miller & Fortney (2011), which found that hot-Jupiter inflation is no longer obvious below an incident flux of $\sim 2 \times 10^5 \text{ W m}^{-2}$. This constraint is only weakly informative and has a conservative 1σ interval between 2×10^4 and $2 \times 10^6 \text{ W m}^{-2}$.

One of the main benefits of Bayesian techniques is the ability to set priors based on known physical constraints (for example a maximum density based on a fully iron giant planet). However in this case, we do not insert any physical priors. According to Fortney et al. (2007), radii of $0.5\text{--}0.6R_J$ are possible for low mass giants with heavy cores in a 50–50 rock/ice mixture and even smaller radii would be possible for pure rock and iron planets (down to less than $0.3R_J$), which is significantly lower than any observed radius in our sample. Thus, it is unlikely that

² In this context, known as the hyperparameters.

adding a maximum density constraint would have affected our results. We also prefer not to insert constraints from interior models into our HBM to keep it purely observation driven, so we simply constrain the true mass, radii, and fluxes to be positive. Apart from the constraints on the mass bin boundaries, we do not find that our choice of priors significantly affects the posterior distributions for any of the other parameters.

The incident flux depends on the stellar effective temperature, orbital distance, and stellar radius as

$$F = \frac{R_\star^2}{a^2} \sigma T_\star^4, \quad (6)$$

where σ is the Stefan–Boltzmann constant, a is the orbital distance (taken as the semi-major axis), and $R_\star$ and $T_\star$ are the stellar radius and temperature, respectively. For each planet, we calculate the distribution of flux based on those parameters, before solving the HBM. We find that the skew-normal distribution, described in [Azzalini \(1985\)](#), is a very good fit to the resulting asymmetric distributions of flux. We therefore fit a skew-normal distribution for each planet’s incident flux, and extract the best-fit parameters μ , σ , and the skewness α . This reduces our number of free parameters by nearly a factor of two, aiding in convergence time.

Our data come from a large range of programmes and sources, and we do not include the statistical biases introduced during observation. Observations currently favour planets with large radii and masses, however the nature of this bias depends on the instrument, and how close a parameter is to the detection threshold. The statistical biases in the flux parameter are also complicated; higher fluxes are favoured because of closer orbital distances, but disfavoured because of a smaller transit depth for larger stars. We must take this into consideration when forming our conclusions. There are also selection effects in the ground-based RV follow up that we are unable to model.

We find the posterior distributions of our parameters by drawing samples from the joint posterior distribution $p(\mathbf{R}, \mathbf{F}, \mathbf{M}, \boldsymbol{\alpha} | \tilde{\mathbf{x}})$, using a Markov chain Monte Carlo (MCMC) sampler, with a Metropolis–Hastings (M–H) algorithm ([Hastings 1970](#); [Chib & Greenberg 1995](#)) implemented in the PyMC2 package in Python ([Patil et al. 2010](#)). Hamiltonian Monte Carlo samplers such the No-U-Turn Sampler ([Hoffman & Gelman 2011](#)), implemented in PyMC3, cannot be implemented as a consequence of the discontinuous nature of our forward model. Other alternatives, such as the affine invariant ensemble sampler, implemented in emcee, are unsuitable owing to the nearly 900 dimensions in our model ([Goodman & Weare 2010](#); [Foreman-Mackey et al. 2013](#)).

While increasing the number of hyperparameters can lead to a better fit to the data, it can also lead to overfitting. Increasing the model complexity gives our model more freedom, which can lead to multi-modal posterior distributions that our M–H sampler cannot accurately sample, leading to un-converged posteriors. We fit multiple parametrisations and based our final choice firstly on whether it could converge and whether it had well-defined uni-modal posteriors. To check for convergence, we run multiple MCMC chains and calculate the Gelman–Rubin convergence metric for our posterior samples, where values close to 1 indicate convergence and good mixing of the chains ([Gelman & Rubin 1992](#)).

To choose between converged models, we use two tools. Firstly, we calculate the Bayesian information criterion (BIC;

[Schwarz 1978](#)), from

$$BIC = k \ln N - 2 \ln \hat{L}, \quad (7)$$

where k is the number of hyperparameters, N is the number of samples (the number of planets), and the highest likelihood $\hat{L}$ is defined as

$$\hat{L} = p(\tilde{\mathbf{x}} | \hat{\mathbf{M}}, \hat{\mathbf{F}}, \hat{\boldsymbol{\alpha}}, \mathcal{M}) \quad (8)$$

$$= \int p(\tilde{\mathbf{x}} | \mathbf{R}, \hat{\mathbf{M}}, \hat{\mathbf{F}}) p(\mathbf{R} | \hat{\mathbf{M}}, \hat{\mathbf{F}}, \hat{\boldsymbol{\alpha}}, \mathcal{M}) d\mathbf{R} \quad (9)$$

$$= \prod_{i=1}^N \int p(\tilde{\mathbf{x}}_i | R_i, \hat{M}_i, \hat{F}_i) p(R_i | \hat{M}_i, \hat{F}_i, \hat{\boldsymbol{\alpha}}, \mathcal{M}) dR_i, \quad (10)$$

where by marginalising out the second level of our model, we obtain a two-level model, and Eq. (10) relies on the independence between observations. This is an assumption that we have to make as we cannot model the correlation between observations and between parameters with the publicly available information. The symbol $\mathcal{M}$ refers to the model in question, and $\hat{M}_i$, $\hat{F}_i$, $\hat{\boldsymbol{\alpha}}$ represent the model parameters set to their maximum likelihood values. A lower BIC indicates a better model, where differences of ~ 10 are required for strong evidence in favour of a model (see [Schwarz 1978](#)). Because of the Gaussian approximation required for the BIC, which is not satisfied for all our posterior parameter distributions, the BIC only plays a minor role in our model selection. We only consider very large differences (more than 20 in the BIC) to be significant.

A better tool for model testing is to perform detailed statistical tests on our posterior predictive fit. We simulate the data that would be produced by our model and compare our model data to the real observed data. This essentially tests the question of if our fully marginalised model was correct, how likely would we be to produce the observed data. This model checking follows the procedure in [Wolfgang et al. \(2016\)](#) and is described in Sect. 6.

For the final baseline model, we decided to fix the highest mass boundary at $2.5 M_J$. In models where the highest mass boundary was treated as an unknown parameter, its posterior distribution was very wide and poorly constrained, generally varying between 2 and 10 Jupiter masses. This is likely due in part to the sparsity of such high mass planets in the database. Fixing it in place greatly sped up convergence times, and allowed us to focus on the lower masses where there is greater variation.

5. Results

5.1. Model posteriors

We ran 10 chains for 2 000 000 iterations each with a burn-in of 1 000 000 iterations and a thin factor of 200. The Rubin–Gelman convergence metric was evaluated for the posterior distributions, and we achieved $\hat{R} < 1.003$ for all the model parameters, indicating a high level of convergence ([Gelman & Rubin 1992](#)).

The posterior distributions for our mass bin boundaries, m_1 and m_2 (see Fig. 4), and the fixed boundary at $m_3 = 2.5$, split our sample into four mass regimes: planets with sub-Saturn masses ($< 0.37 M_J$), planets with sub-Jupiter masses ($0.37–0.98 M_J$), and planets heavier than Jupiter ($0.98–2.5 M_J$ and $2.5 + M_J$). Above the inflation threshold, the radii have a clear dependence on incident flux for all but the sub-Saturn mass planets; the

Fig. 3. Posterior distributions of the A , C , and $\log F_s$ parameters for each mass regime.

flux-radius relation for each mass bin is distinct in the sense that the joint posterior distributions for the hyperparameters that govern this relation, A , C , and F_s , have little overlap (see Fig. 3).

We plot the posteriors for C , A , and F_s in Fig. 3 for each mass bin and report the best-fit values (which we take to refer to the median sample values from here on) and error bars on the 16th and 84th percentiles for all the hyperparameters in Table 1. We find that radius inflation, $\Delta R = \mu_R - C$, governed primarily by A , is highly dependent on mass. The highest degree of inflation is found for planets with masses $0.37 - 0.98 M_J$, and it decreases with mass for the $0.98 - 2.5 M_J$ and $2.5 + M_J$ regimes. From the best-fit values of our parameters, the radius dependence at

$F > F_s$ is

$$\Delta R = \begin{cases} -0.33 \cdot (\log_{10} F - 6.09), & \frac{M}{M_J} < 0.37, \\ 0.70 \cdot (\log_{10} F - 5.5), & 0.37 \leq \frac{M}{M_J} < 0.98, \\ 0.52 \cdot (\log_{10} F - 5.8), & 0.98 \leq \frac{M}{M_J} < 2.50, \\ 0.22 \cdot (\log_{10} F - 5.2), & 2.50 \leq \frac{M}{M_J}. \end{cases} \quad (11)$$

For the lowest mass bin, $M < 0.37 M_J$, the posteriors are much broader and the scatter is much larger, as shown in Fig. 5. In Figs. 6 and 7, we see that the radius-flux distribution is more complicated than for the other mass bins, featuring a change in

Table 1. Best-fit values for the hyperparameters with error bars for the 16th and 84th percentiles.

Mass range	A	C	$\log_{10} F_s$
0.10–0.37	$-0.33^{+0.51}_{-0.91}$	$0.95^{+0.05}_{-0.08}$	$6.1^{+0.7}_{-1.6}$
0.37–0.98	$0.70^{+0.07}_{-0.06}$	0.98 ± 0.04	$5.52^{+0.07}_{-0.09}$
0.98–2.50	$0.52^{+0.07}_{-0.07}$	1.06 ± 0.03	$5.82^{+0.09}_{-0.11}$
2.50+	$0.22^{+0.08}_{-0.06}$	$0.99^{+0.05}_{-0.14}$	$5.2^{+0.3}_{-1.1}$
	$\sigma_{R,0}$	$\sigma_{R,1}$	$\sigma_{R,3}$
	$0.21^{+0.04}_{-0.03}$	$0.10^{+0.02}_{-0.01}$	0.12 ± 0.01

Notes. The mass boundaries were 0.37 ± 0.02 , $0.98^{+0.04}_{-0.05}$, and 2.5, with the latter fixed at that value.

Fig. 4. Posterior distribution of the two variable mass bin boundaries, m_1 and m_2 . We kept m_3 fixed at $2.5 M_J$.

the behaviour of the radii at $\sim 10^6 \text{ W m}^{-2}$. A number of planets appear not to follow the trend of increasing radii with flux³, and we also note a lack of inflated planets at fluxes higher than $\sim 10^6 \text{ W m}^{-2}$ in this mass range. Our chosen model, which assumes that the planets in a mass range behave uniformly with regards to their radius relationship with flux, does not seem to be a good fit for the least-massive planets in our sample. As a consequence of its large uncertainties however the model fit below $0.37 M_J$ may not be reliable. The result is still informative as it places a lower bound on the range of masses over which our model works. For a further discussion see Sect. 7.2.

We find that the mean low irradiation radius C is highest for the class of planets more massive than Jupiter ($0.98\text{--}2.50 M_J$), and decreases slightly for both heavier and lighter planets in line with what is expected from model predictions (Fortney et al. 2007; Baraffe et al. 2008). However, it would also be affected by any trends in the heavy element masses of the planets (see Miller & Fortney 2011; Thorngren et al. 2016).

The inflation threshold is only well constrained in $0.37\text{--}2.5 M_J$, and the inflation mechanism requires higher fluxes to be activated when mass is increased, as F_s is about 1.3

Fig. 5. Posterior distributions for the intrinsic physical scatter of the 3 regimes considered (see Eq. (5)).

times greater for the $0.98\text{--}2.50 M_J$ than for the $0.37\text{--}0.98 M_J$ gas giants. For $2.5 + M_J$ the distribution of F_s provides a poor constraint with an extended tail at low incident fluxes in Fig. 3. This is likely due to the lack of heavy planets at low irradiation in our sample. Similarly for $0.1\text{--}0.37 M_J$ there is also a poor constraint, with a long tail in Fig. 3, likely as a result of the poor fit of our model in this mass range.

Figure 5 shows the physical scatter posteriors, σ_R . Low mass planets ($M < 0.37 M_J$) have by far the largest scatter, at $0.21 R_J$. For masses above $0.37 M_J$, we find the scatter is slightly higher in the inflated planets above the flux threshold than for weakly irradiated planets ($0.12 R_J$ vs. $0.10 R_J$). Most of the proposed mechanisms for inflation would introduce their own latent parameters, which control the degree of inflation, for example opacity, tidal heating rate, and Ohmic heating depth. Variations in these parameters may add to the scatter already present from heavy element content (see e.g. Heng 2012).

5.2. Marginalised posterior distribution

The functional form of the FMR relations, taking the best-fit parameter values, is shown in Fig. 6, with scatter limits as dotted lines. The points represent the raw observational dataset $\tilde{\mathbf{x}} = \{\tilde{F}_i, \tilde{M}_i, \tilde{R}_i\}_{i=1}^N$ and its observational uncertainties.

Figure 7 represents a more complete picture of our posterior FMR relation. The points are the posterior true parameters and their uncertainties, which can be seen to be different from the observed values. The shaded regions show the posterior predictive distribution for the true radii and represent the parameter space within which we expect the central 68% of planets to be contained for a given incident flux and planet mass. The boundaries of the region represent the 1σ contour lines of the distribution $p(R|M_j, F, \tilde{\mathbf{x}})$ as a function of F for specific masses⁴. $p(R|M_j, F, \tilde{\mathbf{x}})$ is our forward model marginalised over the posterior distribution of the hyperparameters,

$$p(R|M, F, \tilde{\mathbf{x}}) = \int p(R|M, F, \alpha) p(\alpha|\tilde{\mathbf{x}}) d\alpha, \quad (12)$$

⁴ The masses of each plot are fixed and chosen to be in the middle of each mass bin (away from the boundaries). Near the boundaries, the scatter would be larger (and the shaded region wider) as it transitions from one bin to the next. This is because marginalising over the hyperparameters, including the mass bin boundaries, means that a particular mass could lie in either of the two adjacent mass bins with some probability. In essence, the R/F relations plotted above vary smoothly with mass (except at the $2.5 M_J$ boundary), although only cross sections are shown in Fig. 7.

³ The five planets are HD 149026 b (Butler et al. 2006; Torres et al. 2008), K2-39 b (Van Eylen et al. 2016), Kepler-101 b (Bonomo et al. 2014), Kepler-41 b (Santerne et al. 2011; Butler et al. 2006), and WASP-126 b (Maxted et al. 2016).

Fig. 6. Relations of our model plotted for parameters set at their best-fit values. The solid line is $\mu_R(F, M)$, and the dotted lines are $\mu_R(F, M) \pm \sigma_R(F, M)$; in the best-fit model, 68% of true values of the planets should lie within. The points represent the data values of the planet parameters of our sample $(\bar{F}_i, \bar{M}_i, \bar{R}_i)$ with observational uncertainty error bars.

where $p(\alpha|\tilde{x})$ is the posterior distribution of our hyperparameters. The dashed line represents the median point of the distribution in Eq. (12), and the dotted line is the central 95% coverage interval. In the Bayesian sense, Eq. (12) represents the most accurate statement of our knowledge of the FMR relation, having inferred the posterior distributions of all our model parameters using the available data, as it includes the uncertainties we have in those parameters. However, the best-fit distribution, $p(R|M, F, \hat{\alpha})$, with $\hat{\alpha}$ from Table 1 is often a good enough approximation.

The $\bar{R}_i$ values in Fig. 6 are more scattered than the true R_i in Fig. 7. This is expected (see in Sect. 4) since the scatter from observational noise is removed from the posterior distribution for R_i . While the form of model we chose affects these posteriors, they tend to always shrink towards a mean line. We justify this result in the model checking in Sect. 6. The posterior distribution should be seen as the distribution of true planet parameters most likely to reproduce the data when observational noise is applied to the values.

6. Model checking

We performed detailed tests to justify our model and checked its consistency with the data using a posterior-predictive test. Our aim was to test the ability of our model to reproduce datasets similar to the observed dataset and to quantify the discrepancy. Owing to the similar nature of our investigation, we followed the procedure of Wolfgang et al. (2016), which provides a more in-depth explanation. For a detailed overview of hierarchical model checking, see also Bayarri & Castellanos (2007).

Fig. 7. Marginalised posterior distribution of radii given flux, $p(R|F, M, \tilde{x})$, plotted for different planet masses; see Eq. (12). The shaded region is the central 68% interval within which 68% of planet true radii should lie for a given incident flux and planet mass. It represents the 1σ points of $p(R|M, F, \tilde{x})$; the central dashed line is the median and the dotted lines are the 95% interval. The red points represent the posterior true values of the planet parameters of our sample with posterior uncertainty error bars.

We calculate the posterior predictive distribution of model $\mathcal{M}$, defined as the probability of observing a new planet with parameters $x_{\text{new}} = (\bar{R}_{\text{new}}, \bar{M}_{\text{new}}, \bar{F}_{\text{new}})$, given our original observed dataset,

$$p(x_{\text{new}}|\tilde{x}, \mathcal{M}) = \int p(x_{\text{new}}|\theta, \mathcal{M})p(\theta|\tilde{x}, \mathcal{M})d\theta \quad (13)$$

$$= \mathbb{E}_{\theta|\tilde{x}, \mathcal{M}}[p(x_{\text{new}}|\theta, \mathcal{M})], \quad (14)$$

where θ refers to the combined model parameters (R, M, F, α) . Equation (14) is to be read as the expectation value of $p(x_{\text{new}}|\theta, \mathcal{M})$, for values of θ drawn from their posterior distribution, $p(\theta|x_{\text{new}}, \mathcal{M})$. This follows from Eq. (13) by the definition of the expectation value (mean in this case), and is of course trivial to calculate with the samples produced by our MCMC. From here on, we drop the $\mathcal{M}$ notation and treat it as implicit.

To perform the posterior-predictive test, we draw points from the posterior predictive distribution in Eq. (13) to create a mock dataset, $\tilde{x}_{\text{new}}$, of the same size N as our planet sample. By drawing thousands of such mock datasets, we can compare these statistically to our observed dataset.

The comparison in this case involves quantifying certain aspects of the distributions in some derived statistics. We use the same two statistics as Wolfgang et al. (2016): i.e. $f_{1\sigma}$, the fraction of the simulated radii of a given dataset that fall within the 68% coverage interval (the shaded region in Fig. 7, described in Sect. 5.2), and f_{μ} , the fraction of data points that have radius 1σ error bars that cross the median of the posterior distribution

(the dashed line in Fig. 7). While we use the posterior true distributions to aid in performing this comparison, it is important to remember that we are comparing the observed data, i.e. $\bar{R}$, produced from the true values (e.g. R), based on observational uncertainties.

The two statistics test two key attributes of our model. The value $f_{1\sigma}$ is a measure of the scatter of data produced by our model (controlled mainly by σ_R), while f_μ is a measure of how centred the data is about the median line. The value f_μ is thus a proxy for the shape of the dataset distribution, and checks whether the normal distribution chosen in Eq. (2) is justified. As an example, if the scatter of our model proved to be accurate (in reproducing $f_{1\sigma}$), but f_μ was too small, it could suggest that we should have used a distribution with heavier tails and less clustering near the mean.

We must also address how to draw a posterior predictive dataset, $\mathbf{x}_{\text{new}}$, from Eq. (13), where we also follow the treatment of Wolfgang et al. (2016). There are two methods, each tied to one level of our HBM, where the second level represents the true values ($\mathbf{R}$, $\mathbf{F}$, $\mathbf{M}$) and the third level represents the hyperparameters α (see Fig. 2).

For the first method (drawing from the second level), we draw each mock dataset from the posterior true value distributions of the real planet sample. Thus for the i th planet of our mock dataset, we draw $x_{\text{new},i}$ from

$$x_{\text{new},i} \sim \int p(x_{\text{new},i} | R_i, M_i, F_i) p(R_i, M_i, F_i | \mathbf{x}) dR_i dM_i dF_i \quad (15)$$

$$\sim \mathbb{E}_{R_i, M_i, F_i | \mathbf{x}} [p(x_{\text{new},i} | R_i, M_i, F_i)], \quad (16)$$

where

$$p(x_{\text{new},i} | R_i, M_i, F_i) = p(\tilde{R}_{\text{new},i} | R_i) p(\tilde{F}_{\text{new},i} | F_i) p(\tilde{M}_{\text{new},i} | M_i), \quad (17)$$

represents the published observational uncertainty distributions of that particular planet. Performing the test on the datasets produced above checks the posterior distributions of our true parameters for the real planet population (i.e. the points in Fig. 7).

The second method generates an entirely new population of hypothetical planet true values from the marginalised model and is written as

$$R_i \sim \int p(R_i | M_i, F_i, \alpha, \mathbf{x}) p(\alpha | \mathbf{x}) d\alpha \quad (18)$$

$$\sim \mathbb{E}_{\alpha | \mathbf{x}} [p(R | M, F, \alpha)]. \quad (19)$$

From that hypothetical planet sample, we generate a new dataset as in the first method, i.e.

$$x_{\text{new},i} \sim \mathbb{E}_{R_i, M_i, F_i | \mathbf{x}} [p(x_{\text{new},i} | R_i, M_i, F_i)]. \quad (20)$$

In this case, we draw the experimental uncertainties from distributions consistent with our data, conserving the trends between the parameter uncertainties and parameter values⁵.

⁵ The values $\sigma_{M,\text{obs}}$ and $\sigma_{F,\text{obs}}$ are strongly correlated with M and F , respectively. We thus fit a linear regressive model (including a dispersion term) for $\log \sigma_{M,\text{obs}} \sim N(f_1(\log M), \sigma)$ and $\log \sigma_{F,\text{obs}} \sim N(f_2(\log F), \sigma)$, and we draw $\sigma_{M,\text{obs}}$ and $\sigma_{F,\text{obs}}$ from the resulting highest likelihood distribution. We also find that $\sigma_{R,\text{obs}}$ is correlated with F . To assign $\sigma_{R,\text{obs}}$, we split the range of fluxes into four bins, and draw $\sigma_{R,\text{obs}}$ from a random planet that is in the same F bin.

Fig. 8. Distribution of $f_{1\sigma}$ (left) and f_μ (right) for 10 000 generated datasets from the second level of our HBM (the true posteriors, see Eq. (16)). The comparison to the values for the real dataset is represented by the vertical line. The real data falls within the 73rd and 50th percentiles of the generated distributions for $f_{1\sigma}$ and f_μ , respectively.

Our hyperparameters do not parametrise the distributions of F and M for our population, thus we must first draw M and F from the same distributions as our population to use them in Eqs. (19) and (20). Performing the model checking statistics on the hypothetical datasets produced by this method is thus a test of our posteriors for the third level of our HBM, i.e. the hyperparameters.

We calculate $f_{1\sigma}$ and f_μ for a large number of datasets produced by both the above methods, and also for the single real dataset. If our model is consistent with the data, we expect that the statistics calculated from the mock datasets and from the real data agree with each other. We define agreement as meaning that the statistic calculated for the real dataset is within the 1σ interval of the statistics calculated from the mock datasets. In other words, the real dataset must have the same properties as a typical mock dataset.

Figure 8 shows the results of our model checking for the second level of our HBM, and Fig. 9 shows the results for the third level. We measure model-data consistency by computing the percentile of the generated mock datasets that the real datasets fall into (measured by the statistic of choice).

We find that the observed f_μ fall within the typical values produced by our model for datasets produced by both methods, coming in at the 50th and 82nd percentile for the second and third levels, respectively. The observed f_σ falls in the 73rd percentile of data generated from the second level of our HBM, indicating that the spread of our posterior true values (R_i , M_i , F_i) is consistent with the data, and justifying the shrinking effect seen in the points of Fig. 7. For the data generated by the third level of our HBM, f_μ falls in the 82nd percentile, while f_σ falls in the 97th percentile. In other words, only ~6% of the generated mock data were more extreme than our observed data. The result suggests that the form of our model may have room for improvement, as we overestimate the scatter of the data. However, it does not negatively affect our conclusions in Sect. 7.3.

7. Discussion

7.1. Flux-mass-radius relation

Our results show a strong trend of decreasing radius inflation with mass; the most inflated planets are in the mass range

Fig. 9. Distribution of $f_{1\sigma}$ (left) and f_{μ} (right) for 50 000 generated datasets from the third level of our HBM (the hyperparameters, see Eqs. (19) and (20)). The comparison to the values for the real dataset are represented by the vertical line. The real data falls within the 97th and 82nd percentiles of the generated distributions for $f_{1\sigma}$ and f_{μ} , respectively.

0.37–0.98 M_J (see Fig. 1). This effect is also expected from models without additional inflation mechanisms, in which mass increases past a few Jupiter masses lead to a decreasing radius response to incident flux (Fortney et al. 2007; Baraffe et al. 2008). Below the inflation threshold, we also find that planet mass has a small effect and that the largest cold giant radii have masses of 0.98–2.50 M_J .

Thorngren et al. (2016) have found that the heavy element content of non-inflated gas giants is strongly and positively correlated with planet mass. Unless there is a different formation mechanism behind the hottest Jupiters, or in cases of significant evaporation, this trend should apply in our sample as well. This would lead to a further reduction in the radii as planet mass increases and would agree with the observed trends for inflated hot Jupiters. As more detailed predictions from models with inflation become available in the future (with predictions given for wide ranges of parameters such as in Fortney et al. 2007), we will similarly be able to probe the core mass distribution of inflated hot Jupiters.

The incidence of using binning in the mass parameter must also be examined. To do this we look at the radius residuals, $\Delta R_i = R_i - \mu_R(F_i, M_i)$, and their relation with mass. We fit a simple linear regression to the relationship, with an additional scatter parameter, to determine whether a mass-radius gradient of zero is excluded to more than 1σ . We find that within the 0.10–0.37 M_J bin there is a strong positive correlation of radius with mass, and a negative correlation for the 2.5 + M_J planets (see Fig. 10), and for both a gradient of zero is more than 1σ from the mean of the distribution. This additional dependence with mass may be the reason for the large scatter in the lowest mass bin, $M < 0.37$. We find no evidence of a mass-radius correlation for the middle two mass bins. Future models could take this into account.

Extending this analysis to other parameters, we also find a weak negative correlation of radius with metallicity for all mass bins (see Fig. 10), albeit with a very large scatter. This could be caused by a correlation between heavy element content and stellar metallicity (Guillot et al. 2006; Miller & Fortney 2011). However, we find no correlation of the radius with host star age. Our results agree broadly with Enoch et al. (2012).

We also performed a short test to determine the first-order bias introduced by averaging our observational error bars where

Fig. 10. Radius residuals, $\Delta R_i = R_i - \mu(F_i, M_i)$, plotted against host star metallicity (top), and planet mass for the lowest mass bin (middle) and the highest mass bin (bottom). The shaded grey region is the marginalised 68% coverage interval of our linear regression, and the shaded blue region represents the 1σ limits of the mean line.

they are asymmetric (see Sect. 2). The asymmetry in the uncertainties is most significant for the radii, where it is biased towards larger upper error bars. To determine the bias, we fit a simple three-level hierarchical model to the residual observed radii, $\Delta \tilde{R}_i = \tilde{R}_i - \mu_R(F_i, M_i)$, treating the population as being drawn from a normal distribution with mean μ and variance σ^2 , and having the same measurement uncertainties as the data. We then take two cases: one where we average the upper and low uncertainties as before and one where we allow them to vary. For the latter, we use a discontinuous split normal distribution

$$p(\Delta \tilde{R}_i | \Delta R_i) = \begin{cases} \text{Normal}(\Delta \tilde{R}_i, \Delta R_i, \sigma_l^2), & \Delta R_i < \Delta \tilde{R}_i \\ \text{Normal}(\Delta \tilde{R}_i, \Delta R_i, \sigma_u^2), & \Delta R_i \geq \Delta \tilde{R}_i \end{cases} \quad (21)$$

to simulate asymmetric uncertainty distributions with upper and lower error bars σ_u and σ_l respectively. Comparing the results between the two cases, we find that the difference in our posteriors for σ and μ are 0.05% and 0.1%, respectively. This is well within the error bars on the posteriors in this simple HBM and within the errors bars we present for our hyperparameters in Sect. 5, which are all of the order of 1–10%.

7.2. Low-mass inflation cut-off

A key finding we present is the inference of a boundary at $\sim 0.37 M_J$, above which we see the most inflated gas giants, with a

Fig. 11. Mass-flux distribution for all gas giants that have a radius greater than $1 R_J$. The blue dashed lines show our best-fit mass boundaries. As the mass decreases below $1 M_J$, the maximum incident flux at which we find inflated planets also decreases (black line).

clear positive dependence on incident flux. Below we find a lack of highly inflated radii with $R > 1.5 R_J$ and a more complicated trend (see Figs. 6 and 7).

From visual inspection of the data in the lowest mass bin, we see the beginning of a trend of inflating radii as a function of flux $\sim 10^5 \text{ W m}^{-2}$, particularly for planets closer to $0.37 M_J$. However this stops at $\sim 10^6 \text{ W m}^{-2}$, and there are no inflated planets beyond this threshold. Instead we see a reversal, with decreasing radii and denser planets, although we have a small sample size in that regime (five planets). Giants at the top end of the mass range, which would otherwise be inflated, also show decreasing radii (e.g. HD 149026 b has a mass of $0.37 \pm 0.03 M_J$ and evolutionary models find it to be very dense, with a predicted core of $80\text{--}110 M_\oplus$; Burrows et al. 2007). Previous studies about hot-Jupiter inflation (e.g. Laughlin et al. 2011; Enoch et al. 2012) generally use manually chosen minimum mass cut-offs for their samples. However, we recommend that future studies focus on planets heavier than $\sim 0.37 M_J$, to avoid contaminating the samples with planets where other processes may start to play a role alongside inflation.

In Fig. 11 we plot the incident flux of all the planets in our sample that have radii greater than $1 R_J$ (and are thus potentially inflated) and find that there is a maximum incident flux that depends on the planet mass and decreases as we decrease the mass.

To explain the reversal of the inflation trend and the lack of such inflated low mass giants at $F > 10^6 \text{ W m}^{-2}$, an obvious cause could be bulk evaporation or Roche-lobe overflow leading to significantly denser planets if the outer gaseous layers are stripped (see e.g. Kurokawa & Nakamoto 2014). This could also lead to such evaporated giants leaving remnants too low in mass to enter our sample ($M < 0.1 M_J$). Roche-lobe overflow and asymmetric evaporation, such as described in Gu et al. (2003) and Baraffe et al. (2005), could also stop planetary migration beyond a certain point, or even lead to outward migration, which may explain the paucity of planets in this mass range closer to their star (relative to the other mass ranges).

We must also consider the observational biases that may be present. There are few data points available in this regime, however our conclusion relies not only on the trend of over-dense planets closer to the star, but also on the clear lack of $R > 1.0 R_J$ gas giants at incident fluxes greater than $\sim 10^6 \text{ W m}^{-2}$. Below this threshold, there are numerous such large giants, and we would expect that as orbital distance is decreased and incident flux increased, the probability of detecting similar objects

should increase. Similarly, larger radii should also be easier to detect, thus if observational surveys are able to find planets with $R < 1.0 R_J$ in the mass ranges that we examine, we would also expect these surveys to be able to find more inflated planets (if they exist). Thus it is unlikely that our conclusions are impacted by an observational bias.

7.3. Lack of non-inflated hot Jupiters

Of the proposed mechanisms for radius inflation, many can only work under specific circumstances. For example, tidal dissipation requires non-zero eccentricity, unless the eccentricity can be excited (Arras & Socrates 2010) and Ohmic heating requires heat deposition deep enough near the radiative-convective boundary; see e.g. Heng (2012); Perna et al. (2012); Ginzburg & Sari (2016). Thus, not only may we expect some planets not to be inflated, but the varying degrees of inflation caused by each mechanism should leave an imprint on the distribution of radii and its scatter. As an example, we make a simple comparison of our posterior FMR distribution with planetary models that do not include inflation mechanisms (Fortney et al. 2007) to determine what fraction of radii could be explained without resorting to inflation.

To compare our observation-based results with theory, we must not only account for the mass and incident flux, but also for planet age and the unknown core mass. We assume that planet ages are distributed uniformly in the range $0.316\text{--}8.0 \text{ Gyr}$ (in rough agreement with the age distribution of our planet sample). The core mass, or heavy element fraction, is not an observable quantity. However Thorngren et al. (2016) showed that there exists a strong correlation between core mass and planet mass at low insolation, and inferred the relation through a regression fit. For the purposes of this work, we assume that the inferred core mass (M_C) to planet mass (M) relation of Thorngren et al. (2016) holds equally well for highly irradiated planets. Their relation is denoted as a probability distribution $p_T(M_C|M, \hat{\theta}_T)^6$, for the highest likelihood point of their model parameters $\hat{\theta}_T$.

Fortney et al. (2007) gave a detailed table of predictions of planet radii as a function of mass, incident flux (orbital distance), core mass, and planet age. However, using the Thorngren et al. (2016) relation, planets heavier than $\sim 1 M_J$ can have core masses greater than the domain of predictions in Fortney et al. (2007). Thus, we can only carry out the following analysis for the lowest inflated mass bin, $0.37\text{--}0.98 M_J$. There is also a systematic error in this comparison, as the Fortney et al. (2007) models considered the heavy elements to be contained fully in a rock-ice core, while the Thorngren et al. (2016) models have the heavy elements also distributed in the envelope. Models with all the heavy elements contained in the core tend to produce larger radii than models with heavy elements mixed throughout the envelope, so our model radii are overestimated (see Baraffe et al. 2008; Thorngren et al. 2016).

We produce a distribution of model radii as a function of incident flux, $p_M(R|F)$, fully marginalised over the relevant distributions of core mass and planet age (t) and marginalised over the width of the mass bin, $m_i - m_{i+1}$,

$$p_M(R|F) = \int \delta(R - f(M, M_C, t, F)) p_T(M_C|M, \hat{\theta}_T) U(t) dt dM dM_C, \quad (22)$$

⁶ The fitted model in Thorngren et al. (2016) has the form $M_C/M_\oplus = k \times 57.9 \times M^{0.61}$, where the multiplicative scatter, k is drawn from $\log_{10} k \sim N(\mu = 0, \sigma = \log_{10} 1.83)$.

Fig. 12. Our marginalised posterior radius-flux distribution for $0.37 \leq \frac{M}{M_J} < 0.98$, with the central 68% interval shaded in dark grey, and the 95% interval in light grey. Plotted against a flux-radius relation $p_M(R|F)$ predicted by a non-inflationary planet model, with the central 68% interval shaded in blue.

where $f(M, M_C, t, F)$ refers to the model predictions of Fortney et al. (2007), which are deterministic, δ is the delta distribution, and $U(t)$ is the uniform distribution for the planet age. Thus we calculate the model prediction for the distribution of radius as a function of incident flux, similar to our posterior FMR distribution. We compare the model prediction against our observation-driven posterior FMR distribution in Fig. 12.

This is intended as a rough and qualitative exploration of how we could compare marginalised theoretical models to uncertainty decoupled data. We see that below the flux threshold, our theory-predicted $p_M(R|F)$ overestimates the radius, but falls within the 1σ interval. This could be because we underestimated the core masses or planet age, but may also be caused by the previously mentioned increase in radii, when all the heavy elements are placed into a pure rock-ice core, compared to being distributed in an envelope.

With that in mind, we look at the predictions above $\sim 10^6 \text{ W m}^{-2}$. The posterior FMR relation and the model prediction drastically diverge, with negligible overlap beyond $F \sim 1.6 \times 10^6 \text{ W m}^{-2}$. Only one real planet, HATS-9 b (Brahm et al. 2015), has a posterior density that has significant overlap with $p_M(R|F)$, and we cannot rule out the fact that it may be an inflated planet with an unusually large core. Such over-dense planets have been found in the low-flux regime, for example see Kepler-539 b, with total mass $1.00 M_J$ and heavy element mass $0.49 M_J$ (Thorngren et al. 2016). Furthermore, what little overlap exists is only significant above the mean line of $p_M(R|F)$ (solid blue line). Below the mean line, where we would expect to find denser planets with large cores, there is negligible overlap, even with the 95% interval of our the FMR relation of our HBM.

While a more detailed study would be warranted, our results suggest that classical planets that do not exhibit radius inflation do not exist above fluxes of $\sim 10^6 \text{ W m}^{-2}$ for this mass range. Even if some of the observed planets could be explained as being young non-inflated gas giants with light cores, if a population of such young low-density planets were to exist, we would also expect to find similar planets with medium-mass or heavy cores and older ages (which would fall below the mean line of $p_M(R|F)$).

Although we found a slight discrepancy between our hyperparameter-marginalised HBM posteriors and the data in Sect. 6, this discrepancy suggests that our HBM is overestimating the intrinsic scatter. With a lower scatter, the overlap between the observed data and the model prediction would be decreased.

Although we found that we were overestimating the radii at low irradiation, increasing the input core masses to counter this would only serve to further drive $p_M(R|F)$ away from our FMR distribution at high fluxes.

If re-inflation of hot Jupiters is impossible (see Wu & Lithwick 2013; Ginzburg & Sari 2016), this would mean that all of the current population of highly irradiated $0.37\text{--}0.98 M_J$ giants must have migrated to their present locations very early in their lifetimes. This could however be avoided with tidal re-inflation (e.g. Jermyn et al. 2017) followed by delayed contraction from Ohmic heating or another mechanism.

Furthermore, the findings of Heng (2012) showed that variations in atmospheric scattering and absorption should produce a large scatter in the degree of Ohmic dissipation. In certain cases (such as atmospheres with temperature inversions and high optical opacities, or planets with too weak a magnetic field) we may even expect no significant Ohmic dissipation at depth (Heng 2012; Perna et al. 2012). This should be reflected in the radius distribution of hot Jupiters. Thus if Ohmic heating is the dominant inflation process, the lack of non-inflated hot Jupiters must be explained by either a lack of high optical opacities and temperature inversions in this temperature regime, or would require a different inflation mechanism for the cases where Ohmic heating might fail.

We also briefly consider the observational biases that may be present in our data. While there is a detection bias towards finding larger radii and masses that we cannot fully incorporate into our model, since we use various sources of data for which we may not have well-documented completeness functions, we note that the non-inflated planets that we are comparing with have roughly the same radii as gas giants found below the inflation threshold. Thus surveys that can detect the giants below the flux threshold should detect closer-in gas giants of the same radii as well. Planets closer to their star should also have a higher chance of transiting, and thus being detected, so we should be biased towards finding planets in the high incident flux regime.

A more thorough treatment of the observational biases would be the next step, as well as a comparison with newer models that include inflation mechanisms. Within the framework of such a hierarchical model, we could also extract a more quantitative measure of the percentage of gas giants that are non-inflated using a mixture model (see McLachlan & Peel 2000; Celeux 2007). Another approach, applied recently by Thorngren & Fortney (2018), is to infer the heating required without specifying an underlying mechanism. These authors find that the functional form of the heating rate $\epsilon(F)$ required to explain the radii most closely matches the profile predicted by Ohmic heating. Given our findings that all $0.37\text{--}0.98 M_J$ hot Jupiters beyond $\sim 10^6 \text{ W m}^{-2}$ show inflation, it would be informative to further explore the scatter of $\epsilon(F)$ at a particular flux within their framework.

8. Conclusions

We have constructed a theory-independent HBM to constrain the probabilistic relation between the planet radius, mass, and incident flux for 286 gas giants. We thus find the intrinsic scatter of the gas giant population, decoupled from the observational uncertainties in both the radius parameter, and the mass and incident flux parameters. The posterior distribution we inferred may be used to test and constrain theoretical models for gas giant radius evolution, especially involving inflation mechanisms. Constraining the intrinsic scatter allows us to determine

the real range of radii that need to be reproduced by theory at each regime of incident flux and mass.

Our key results are summarised below:

- We find that planetary mass plays a significant role in the degree of radius inflation, with the most inflated hot Jupiters coming from 0.37–0.98 M_J range, and that the response of radius to incident flux decreases as mass is increased. This could be caused by either a trend of increasing core masses with increasing total mass, similar to that found for gas giants below the inflation threshold (see [Thorngren et al. 2016](#)), or by the increase in surface gravity with mass, which would make them harder to inflate.
- Below 0.37 Jupiter masses, there is an abrupt lack of heavily inflated radii, and we find that radii begin to decrease with incident fluxes near $\sim 10^6 \text{ W m}^{-2}$. Such a trend of decreasing radii is not present at higher masses. We also note that there is a cut-off point in incident flux beyond which hot Jupiters no longer exist and that for inflated planets below 1.0 M_J this cut-off point decreases with decreasing mass.
- We use our inferred posterior distribution to show that there is no evidence for non-inflated hot Jupiters at fluxes greater than $F \sim 1.6 \times 10^6 \text{ W m}^{-2}$ and masses of 0.37–0.98 M_J . In this case we define non-inflated to mean that their radii could be reproduced by an evolutionary model that only considers incident flux energy deposition in the photosphere with no additional inflationary effects ([Fortney et al. 2007](#)).

To extend this study, a mixture model could be a useful tool to provide a more quantitative measure of the fraction of gas giants that do not exhibit any inflation, and could be incorporated within a Bayesian hierarchical model (see [McLachlan & Peel 2000](#); [Celeux 2007](#)). With increasingly detailed inflated model predictions given for wide ranges of latent parameters (such as those in [Fortney et al. \(2007\)](#) for core mass and age), we will begin to extract information about hot-Jupiter inflation from the entire observed distribution and its scatter. The forward model of our HBM, $p(R|M, F, \alpha)$, may be replaced by a theoretically driven model for model selection purposes, and/or to constrain latent parameters in such models. Using a mixture model, we could also find the fraction of objects that are consistent with a particular model, and in which regions of mass, incident flux, and other parameters this may be the case. With the increasingly large number of discovered exoplanets, such statistics-based studies will continue to become more important.

Acknowledgements. We thank the referee, Mathieu Havel, for a comprehensive review that improved our paper. We thank Kevin Heng for valuable discussions and ideas. M.S acknowledges support from the Swiss National Science Foundation (PP00P2-163967). B.-O.D. acknowledges support from the Swiss National Science Foundation in the form of a SNSF Professorship (PP00P2-163967). This work has been carried out within the framework of the NCCR PlanetS supported by the Swiss National Science Foundation. Calculations were performed on UBELIX (<http://www.id.unibe.ch/hpc>), the HPC cluster at the University of Bern. This research has made use of the Exoplanet Orbit Database and the Exoplanet Data Explorer at exoplanets.org. This research has made use of the NASA Exoplanet Archive, which is operated by the California Institute of Technology, under contract with the National Aeronautics and Space Administration under the Exoplanet Exploration Program.

References

Akeson, R. L., Chen, X., Ciardi, D., et al. 2013, *PASP*, **125**, 989
 Almenara, J. M., Damiani, C., Bouchy, F., et al. 2015, *A&A*, **575**, A71
 Anderson, D. R., Smith, A. M. S., Lanotte, A. A., et al. 2011, *MNRAS*, **416**, 2108
 Arras, P., & Socrates, A. 2010, *ApJ*, **714**, 1

Azzalini, A. 1985, *Scand. J. Stat.*, **12**, 171
 Baraffe, I., Chabrier, G., Barman, T. S., et al. 2005, *A&A*, **436**, L47
 Baraffe, I., Chabrier, G., & Barman, T. 2008, *A&A*, **482**, 315
 Baraffe, I., Chabrier, G., Fortney, J., & Sotin, C. 2014, *Protostars and Planets VI* (Tucson: University of Arizona Press), 763
 Batygin, K., & Stevenson, D. J. 2010, *ApJ*, **714**, L238
 Bayarri, M. J., & Castellanos, M. E. 2007, *Statist. Sci.*, **22**, 322
 Bodenheimer, P., Lin, D. N. C., & Mardling, R. A. 2001, *ApJ*, **548**, 466
 Bodenheimer, P., Laughlin, G., & Lin, D. N. C. 2003, *ApJ*, **592**, 555
 Bonomo, A. S., Sozzetti, A., Lovis, C., et al. 2014, *A&A*, **572**, A2
 Brahm, R., Jordán, A., Hartman, J. D., et al. 2015, *AJ*, **150**, 33
 Burrows, A., Hubeny, I., Budaj, J., & Hubbard, W. B. 2007, *ApJ*, **661**, 502
 Butler, R. P., Wright, J. T., Marcy, G. W., et al. 2006, *ApJ*, **646**, 505
 Celeux, G. 2007, *Mixture Models for Classification*, eds. R. Decker, & H. J. Lenz (Berlin, Heidelberg: Springer), 3
 Chabrier, G., & Baraffe, I. 2007, *ApJ*, **661**, L81
 Chabrier, G., Baraffe, I., Allard, F., & Hauschildt, P. H. 2005, ArXiv e-prints [[arXiv:astro-ph/0509798](https://arxiv.org/abs/astro-ph/0509798)]
 Chib, S., & Greenberg, E. 1995, *Am. Stat.*, **49**, 327
 Delrez, L., Santerne, A., Almenara, J.-M., et al. 2016, *MNRAS*, **458**, 4025
 Demory, B.-O. 2014, *ApJ*, **789**, L20
 Demory, B.-O., & Seager, S. 2011, *ApJS*, **197**, 12
 Enoch, B., Collier Cameron, A., & Horne, K. 2012, *A&A*, **540**, A99
 Foreman-Mackey, D., Hogg, D. W., Lang, D., & Goodman, J. 2013, *PASP*, **125**, 306
 Fortney, J. J., & Nettelmann, N. 2010, *Space Sci. Rev.*, **152**, 423
 Fortney, J. J., Marley, M. S., & Barnes, J. W. 2007, *ApJ*, **659**, 1661
 Fortney, J. J., Baraffe, I., & Militzer, B. 2010, *Giant Planet Interior Structure and Thermal Evolution*, ed. S. Seager, 397
 Gelman, A., & Rubin, D. B. 1992, *Statist. Sci.*, **7**, 457
 Ginzburg, S., & Sari, R. 2016, *ApJ*, **819**, 116
 Good, I. J. 1965, *The Estimation of Probabilities: An Essay on Modern Bayesian Methods*, Research Monograph No. 30 (Cambridge, USA: MIT Press)
 Goodman, J., & Weare, J. 2010, *Comm. App. Math. Com. Sci.*, **5**, 65
 Gu, P.-G., Lin, D. N. C., & Bodenheimer, P. H. 2003, *ApJ*, **588**, 509
 Guillot, T., & Showman, A. P. 2002, *A&A*, **385**, 156
 Guillot, T., Santos, N. C., Pont, F., et al. 2006, *A&A*, **453**, L21
 Han, E., Wang, S. X., Wright, J. T., et al. 2014, *PASP*, **126**, 827
 Hastings, W. 1970, *Biometrika*, **57**, 97
 Heng, K. 2012, *ApJ*, **748**, L17
 Hoffman, M. D., & Gelman, A. 2011, ArXiv e-prints [[arXiv:1111.4246](https://arxiv.org/abs/1111.4246)]
 Hogg, D. W., Myers, A. D., & Bovy, J. 2010, *ApJ*, **725**, 2166
 Jermyn, A. S., Tout, C. A., & Ogilvie, G. I. 2017, *MNRAS*, **469**, 1768
 Kelly, B. C. 2007, *ApJ*, **665**, 1489
 Kurokawa, H., & Nakamoto, T. 2014, *ApJ*, **783**, 54
 Kurokawa, H., & Inutsuka, S.-I. 2015, *ApJ*, **815**, 78
 Laughlin, G., Crismani, M., & Adams, F. C. 2011, *ApJ*, **729**, L7
 Maxted, P. F. L., Anderson, D. R., Collier Cameron, A., et al. 2016, *A&A*, **591**, A55
 McLachlan, G. J., & Peel, D. 2000, *Finite mixture models*, (New York, Chichester: John Wiley & Sons, Inc.), 419
 Miller, N., & Fortney, J. J. 2011, *ApJ*, **736**, L29
 Patil, A., Huard, D., & Fonnesbeck, C. 2010, *J. Stat. Softw.*, **35**, 1
 Perna, R., Menou, K., & Rauscher, E. 2010, *ApJ*, **724**, 313
 Perna, R., Heng, K., & Pont, F. 2012, *ApJ*, **751**, 59
 Rogers, L. A. 2015, *ApJ*, **801**, 41
 Santerne, A., Bonomo, A. S., Hébrard, G., et al. 2011, *A&A*, **536**, A70
 Schneider, J., Dedieu, C., Le Sidaner, P., Savalle, R., & Zolotukhin, I. 2011, *A&A*, **532**, A79
 Schwarz, G. 1978, *Ann. Statist.*, **6**, 461
 Shabram, M., Demory, B.-O., Cisewski, J., Ford, E. B., & Rogers, L. 2016, *ApJ*, **820**, 93
 Spiegel, D. S., Burrows, A., & Milsom, J. A. 2011, *ApJ*, **727**, 57
 Stein, C. 1956, *Proc. 3rd Berkeley Sympos. Math. Stat. Prob.* (Berkeley, CA: University of California Press), 197
 Thorngren, D. P., & Fortney, J. J. 2018, *AJ*, **155**, 214
 Thorngren, D. P., Fortney, J. J., Murray-Clay, R. A., & Lopez, E. D. 2016, *ApJ*, **831**, 64
 Torres, G., Winn, J. N., & Holman, M. J. 2008, *ApJ*, **677**, 1324
 Tremblin, P., Chabrier, G., Mayne, N. J., et al. 2017, *ApJ*, **841**, 30
 Van Eylen, V., Albrecht, S., Gandolfi, D., et al. 2016, *AJ*, **152**, 143
 VanderPlas, J. 2014, ArXiv e-prints [[arXiv:1411.5018](https://arxiv.org/abs/1411.5018)]
 Weiss, L. M., Marcy, G. W., Rowe, J. F., et al. 2013, *ApJ*, **768**, 14
 Wolfgang, A., Rogers, L. A., & Ford, E. B. 2016, *ApJ*, **825**, 19
 Wu, Y., & Lithwick, Y. 2013, *ApJ*, **763**, 13
 Youdin, A. N., & Mitchell, J. L. 2010, *ApJ*, **721**, 1113