

Detecting episodes of star formation using Bayesian model selection

Andrew J. Lawler¹★ and Viviana Acquaviva^{2,3}

¹*Department of Statistics, Baylor University, One Bear Place #97140, Waco, TX 76798, USA*

²*Department of Physics, CUNY NYC College of Technology, 300 Jay Street, Brooklyn, NY 11201, USA*

³*Institut de Ciències del Cosmos, Universitat de Barcelona, Carrer Martí i Franques 1, E-08028 Barcelona, Spain*

Accepted 2020 January 13. Received 2020 January 13; in original form 2020 May 8

ABSTRACT

Bayesian model comparison frameworks can be used when fitting models to data in order to infer the appropriate model complexity in a data-driven manner. We aim to use them to detect the correct number of major episodes of star formation from the analysis of the spectral energy distributions (SEDs) of galaxies, modelled after 3D-*HST* galaxies at $z \sim 1$. Starting from the published stellar population properties of these galaxies, we use kernel density estimates to build multivariate input parameter distributions to obtain realistic simulations. We create simulated sets of spectra of varying degrees of complexity (identified by the number of parameters), and derive SED fitting results and pieces of evidence for pairs of nested models, including the correct model as well as more simplistic ones, using the BAGPIPES codebase with nested sampling algorithm MULTINEST. We then ask the question: is it true – as expected in Bayesian model comparison frameworks – that the correct model has larger evidence? Our results indicate that the ratio of pieces of evidence (the Bayes factor) is able to identify the correct underlying model in the vast majority of cases. The quality of the results improves primarily as a function of the total S/N in the SED. We also compare the Bayes factors obtained using the evidence to those obtained via the Savage–Dickey density ratio (SDDR), an analytic approximation that can be calculated using samples from regular Markov Chain Monte Carlo methods. We show that the SDDR ratio can satisfactorily replace a full evidence calculation provided that the sampling density is sufficient.

Key words: methods: statistical – galaxies: evolution – galaxies: star formation.

1 INTRODUCTION

The spectral energy distribution (SED) of a galaxy encodes a vast amount of information about the galaxy’s physical properties. Complex processes such as the formation and the evolution of stars and the chemical composition of the interstellar medium between the stars dominate the emission from different parts of the electromagnetic spectrum. These processes are often highly interdependent, making it difficult to isolate the contribution of each component separately (e.g. McKee & Ostriker 2007; Walcher et al. 2011; Conroy 2013).

Past and ongoing photometric and spectroscopic surveys, such as SDSS (York et al. 2000), COSMOS (Scoville et al. 2007), 2MASS (Skrutskie et al. 2006), UltraVISTA (McCracken et al. 2012), CANDELS (Grogin et al. 2011), and 3D-*HST* (Skelton et al. 2014) provide measurements from the ultraviolet to infrared wavelengths in a range of redshifts for millions of galaxies. To interpret these data, a variety of models to describe these physical processes and their interplay have been developed in the past two decades. With the improvement in data quality and volume, models have grown increasingly sophisticated.

Among the many components of these models there are, for example, the initial stellar mass function (IMF), the star formation history (SFH), the dust attenuation law, the chemical enrichment history, and assumptions about the underlying stellar population

models. In particular, the SFH of a galaxy has received a great deal of attention because it provides important insights on which mechanisms dominate galaxy formation and evolution.

The method of SED fitting compares libraries of models to spectro-photometric data to derive physical properties of interest. Techniques such as χ^2 fitting have historically been used to identify a best-fitting model and estimate parameters (Arnouts et al. 1999; Bolzonella, Miralles & Pelló 2000). More recently, Bayesian methods employing Markov chain Monte Carlo (MCMC) sampling algorithms have been employed to better constrain the correlated parameter spaces of galaxy models and analyse posterior probability distributions for parameter inference (e.g. Acquaviva, Gawiser & Guaita 2011).

Bayesian model comparison is a statistical technique aimed at determining the optimal model complexity that is warranted by the data. Despite being used widely in other fields of astronomy, for example, in the context of cosmological model selection (Mukherjee, Parkinson & Liddle 2006; Shaw, Bridges & Hobson 2007; Trotta 2007, 2008), it has seldom been applied to SED fitting thus far. Han & Han (2012) used MULTINEST to calculate and compare the Bayesian evidence from the SED models for hyperluminous infrared galaxies; the same authors used the evidence to compare two popular stellar population synthesis models in Han & Han (2014, 2018). Dries, Trager & Koopmans (2016) and Dries et al. (2018) used Bayesian evidence in a hierarchical setting to compare the prior spaces for several choices of IMFs. Salmon et al. (2016) used Bayesian evidence to test dust laws on a sample of galaxies in CANDELS.

* E-mail: lawler.andy.j@gmail.com

One of the crucial issues in SED fitting is that galaxies are extremely complicated objects, and modelling them with a few parameters is a necessary but a very crude simplification. Until recently, galaxies have typically been described as combinations of simple stellar population models of fixed metallicity and stellar ages. The process of stellar assembly was often described using the exponentially declining (τ) model, which assumes that galaxies form stars at a declining rate, parametrized by the half-time τ , and the attenuation due to dust was often described by a one parameter ‘screen’ model (e.g. Calzetti et al. 2000). However, more recently, it has become clear that these extremely simplified assumptions are unable to capture the complexity of galaxies’ behaviors, and also cause biases in the determination of the physical properties of galaxies. In particular, it has been shown that wrongly reconstructed SFHs introduce noticeable biases in many parameters that are usually estimated through SED fitting, such as stellar masses, stellar age indicators, dust content, and redshift (e.g. Mobasher et al. 2015; Pacifici et al. 2015; Iyer & Gawiser 2017; Leja et al. 2017). In Acquaviva, Raichoor & Gawiser (2015), we evaluated the impact of different sources of non-algorithmic systematics on the recovered SED fitting parameters and concluded that a wrong SFH is the most harmful. Similarly, Iyer & Gawiser (2017) and Carnall et al. (2018) found that fitting the SFH using single stellar populations and simple functional forms (e.g. exponentially declining or constant models) leads to a bias of up to 70 per cent in the recovered physical properties of the galaxies.

Progress in the modelling of stellar populations has been more promising, with several groups developing sophisticated algorithms used to generate and fit simulated spectra in the ultraviolet-through-infrared range with increased flexibility (da Cunha, Charlot & Elbaz 2008; Conroy, Gunn & White 2009; Chevallard & Charlot 2016; Leja et al. 2017; Carnall et al. 2018; Iyer et al. 2019). This means that we are now technically able to explore much larger parameter spaces, but are subject to the curse of dimensionality: the process of SED fitting might take a much longer time, which does not suit large upcoming surveys, and/or the data might not be high enough quality to resolve the many degeneracies inherent to models with a large number of parameters.

For these reasons, in this paper we set out to determine whether Bayesian model selection can be used to infer *from the data* what is the true model complexity. Because the choice of SFH is crucial to this issue, our first exploratory question is as follows: are we able to successfully detect the correct number of major episodes of star formation, which is an index of model complexity, by comparing the evidence values of the different models? We will use simulated spectro-photometric data to answer this question under different assumptions, and then validate the method.

While in principle the Bayesian model comparison method is applicable to models that can occupy different prior spaces, in practice the relevant calculations depend on both the volume occupied by the likelihood and the volume occupied by priors. As a result, it is difficult to disentangle the contribution of the data and the priors to the model comparison. For this reason, we use *nested* models: a hierarchical setting where we compare a more complicated $N + k$ parameter model to a simpler N parameter model. Comparison between these two models is achieved by the use of the Bayes factor, which is simply the ratio of the Bayesian pieces of evidence for each respective model. The mechanics of the Bayes factor can be interpreted as the factor by which the model space shrinks when the data arrive and will only favour a model with more parameters when the data warrant the additional complexity (e.g. MacKay 2003).

This paper is organized as follows. In Section 2, we describe the data, inferred parameters, and noise profile used to construct our mock data. This includes a novel use of kernel density estimation to derive realistic multivariate distributions for the parameter space, as well as the error structure. We then review the SED fitting code BAGPIPES, which is used to generate a realistic galaxy catalogue and to fit the models. We make a number of assumptions about the SFH, IMF, and dust law used for the galaxy population and define the prior space used. In Section 3, we review Bayesian inference methods and summarize MULTINEST, the sampling algorithm used for both parameter estimation and evidence calculation. In Section 4, we define the SDDR ratio, discuss its expected range of applicability, and compare the two approaches used to calculate the Bayes factor. In Section 5, we describe the simulated scenarios that we explore, which includes analysing combinations of various design points to assess the effectiveness of the Bayes factor as a model comparison tool, and discuss our results. Our conclusions are summarized in Section 7.

2 DATA

We apply our methodology to simulated galaxies modelled after those in the CANDELS GOODS-South field from v4.1 of the 3D-*HST* catalogue; this ensures that we employ realistic distributions for input parameters and observational errors. The main catalogue consists of ultraviolet to mid-infrared wavelength measurements of 50 507 galaxies, and it includes fitted stellar population parameters obtained through FAST (Kriek et al. 2009).

We select objects according to the following criteria:

(i) Objects with the `use_phot` flag set to 1. This is defined as an object not classified as a star or close to a bright star, is well exposed on the F125W and F160W bands, has an $S/N > 3$ in F160W, and has both a reasonable photometric redshift fit and a reasonable stellar population fit according to the EAZY/FAST codes.

(ii) Objects with the `z_peak` flag between 0.9 and 1.1, which is defined as the peak of the photometric redshift distribution according to the EAZY code.

The above criteria results in a reference catalogue of 4567 galaxies; we use a total of 23 bands for each galaxy. The parameter values for this sample estimated by FAST and reported by the 3D-*HST* collaboration are used as the basis for the kernel density estimation of the input parameter space. We also use kernel density estimation methods for creation of the mock catalogue error structure, as described in Section 2.1 below.

2.1 Kernel density estimation

In order to generate a realistic parameter space for SED fitting, we choose to leverage kernel density estimates (KDEs) of the fitted stellar population parameters from FAST, which include stellar mass, stellar age, dust extinction, photometric redshift, and the time-scale of star formation (τ), as the FAST code assumes an exponentially declining model. Since these parameters were fitted on a grid of possible values, there are gaps in the fits that limit the precision of the estimated values. However, our multivariate KDE approach allows us to build a realistic population of parameter values on a continuum while correctly capturing the correlations among the parameters.

KDE is a non-parametric data smoothing technique that uses a bandwidth parameter to approximate the probability distribution function of a data sample. We use the k -nearest neighbour (kNN) density estimator in `scikit-learn` (Pedregosa et al. 2011), where

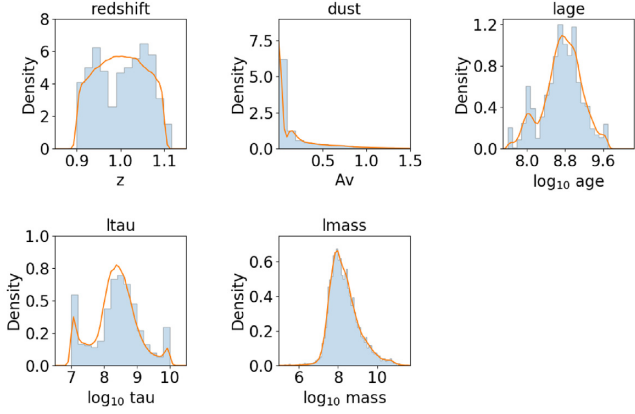


Figure 1. The distributions of the FAST parameter values from 4567 galaxies in the GOODS-S field of the 3D-*HST* catalogue are represented by the blue histograms. The marginalized KDE of each parameter distribution from the multivariate KDE are represented by the orange lines. A bandwidth of 0.0861 is used for the multivariate KDE therefore the marginalized KDEs also have a bandwidth of 0.0861. Note that both the histograms and the marginalized KDEs are normalized.

for a given point \mathbf{x} , the multivariate kNN density estimator estimates the density by (Tran, Wehrens & Buydens 2006)

$$\hat{f}_{\text{knn}}(\mathbf{x}) = \frac{1}{nV_k(\mathbf{x})} \sum_{i=1}^n K\left(\frac{\mathbf{x} - X_i}{\mathbf{H}}\right), \quad (1)$$

where $V_k(\mathbf{x}) = c_d r_k^d(\mathbf{x})$ is the volume of the d -dimensional sphere, $r_k^d(\mathbf{x})$ is the Euclidean distance from \mathbf{x} to the k th nearest neighbour (in other words, the radius of the sphere), and c_d is the constant factor applied (i.e. $c_1 = 2$, $c_2 = \pi$, $c_3 = \frac{4\pi}{3}$, etc.). Also, $K(\cdot)$ is a multivariate kernel, and \mathbf{H} is a vector of bandwidths h , which are all equal in this case. The generalized Euclidean distance metric is given by

$$d_{i,j} = \sqrt{\sum_{k=1}^p (x_{ik} - x_{jk})^2}. \quad (2)$$

We use a multivariate standard normal kernel, given by

$$K\left(\frac{\mathbf{x} - X_i}{\mathbf{H}}\right) = (2\pi)^{-(d/2)} \exp\left(-\frac{(\mathbf{x} - X_i)^T (\mathbf{x} - X_i)}{2\mathbf{H}}\right). \quad (3)$$

This kernel is chosen as it has the desirable properties of being symmetric and unimodal. Several of the parameter distributions reported by FAST are also approximately lognormal, making this kernel suitable for density estimation.

The bandwidth plays a crucial role in determining the bias-variance relationship of the resulting distribution. A larger bandwidth results in a smoother distribution with less variance and more bias, whereas a smaller bandwidth results in a less smooth distribution with more variance and less bias. The bandwidth parameter is chosen using a grid-based search with five-fold cross-validation. Outliers above the 99.5 percentile in the right tail of each FAST parameter distribution are removed so as not to overestimate the tails of the KDEs. The total log-likelihood of the data in each test set is calculated for each bandwidth, and the bandwidth is chosen according to the largest mean log-likelihood. Marginalized distributions for each parameter in the multivariate KDE are shown in Fig. 1. The agreement between the histograms of estimated parameters and the probability distributions derived from the KDE is very good; furthermore, this method preserves all the correlations between different parameters, ensuring




Figure 2. A selection of the distributions of the observational errors from 4567 galaxies in the GOODS-S field of the 3D-*HST* catalogue are represented by the blue histograms. The marginal KDE of each distribution from the multivariate KDE are represented by the orange lines. Each KDE has a bandwidth chosen by five-fold cross-validation.

that the input parameter space of our simulated galaxies is fully realistic.

Similarly, we create KDEs of the distributions of observational errors. However, these are not treated in a multivariate way since errors are assumed to be independent across observations. As with the multivariate parameter KDE, values above the 99.5 percentile are removed before each error KDE is constructed. These are shown in Fig. 2.

2.2 Generating mock data

Armed with our realistic multivariate distribution of galaxy properties, we are now ready to generate our mock catalogues. We use version 0.7.9 of BAGPIPES (Carnall et al. 2018), a modern SED fitting code that we chose for its sophistication in modelling galaxies and its ability to infer parameters using nested sampling, which will be crucial for the evidence calculations in model comparison.

In recent years more complex parametric models, such as the delayed exponentially declining tau model with additional stochastic bursts of star formation, have emerged as a more realistic alternative to the exponentially declining tau model (see e.g. Leja et al. 2019 and references therein).

For our main mode of star formation, we use the delayed exponentially declining tau model, defined as (Gavazzi et al. 2002; Behroozi, Conroy & Wechsler 2010; Lee et al. 2010)

$$\text{SFR}_{\text{delayed}}(t, t_0, \tau) \propto (t - t_0) e^{-\frac{t-t_0}{\tau}}, \quad (4)$$

where t is the age of the universe at observation, t_0 is the time of the onset of star formation, and τ is the width of the SFH.

Random samples from the multivariate KDE based on the FAST parameter distributions (described in Section 2.1) are used as parameter inputs for a delayed tau model.

When present, a second stellar population (burst) is modelled as a simple delta function in time, with parameters burst age and mass formed in the burst. In this case, we uniformly sample a burst mass of 10–90 per cent of the total mass (effectively splitting the randomly drawn stellar mass sample from the multivariate KDE into two components), and we uniformly sample the random burst age between 0.05 and 0.90 Gyr.

We assume a Calzetti (Calzetti et al. 2000) dust law, parametrized by the attenuation in the V band (A_V), with an additional multiplicative

Table 1. BAGPIPES parameter inputs used to generate mock data catalogue. Parameters denoted by ‘KDE’ are drawn from a multivariate input distribution obtained by using kernel density estimation. For stellar mass, ‘KDE*random’ indicates that the total mass is obtained through the KDE method, and the mass of the secondary population is obtained as a uniform random number in the (0.1, 0.9) interval.

BAGPIPES parameter	Value	Range
Global components		
redshift	KDE	0.9–1.1
t_bc	0.01	
Dust components		
law	Calzetti	
A_v	KDE	0–3.0
η	2.0	
Nebular components		
$\log U$	–3.0	
SFH components		
$\log_{10}(\text{mass}_{\text{delayed}} \tau)$	KDE*random	6.5–12.5 M/M_{\odot}
$\log_{10}(\text{mass}_{\text{burst}})$	KDE*random	6.5–12.5 M/M_{\odot}
Metallicity	0.02	
$\log_{10}(\text{age}_{\text{delayed}} \tau)$	KDE	7.7–9.7 Gyr
$\text{Age}_{\text{burst}}$	Random	0.10–0.90 Gyr
$\log_{10}(\tau)$	KDE	7.0–10.0 Gyr

factor η for stars in birth clouds of age < 0.01 Gyr. We set the ionization parameter $\log_{10} U$ for the nebular component to -3.0 , and we assume a Kroupa & Boily (2002) IMF.

A summary of all the BAGPIPES parameter inputs is provided in Table 1.

2.3 Mock data versus 3D-*HST* data

We next compare the distribution of observational fluxes in several bands generated from BAGPIPES using an exponentially declining (τ) SFH model with input parameters $\log_{10}(\text{mass})$, age, τ , A_v , and z sampled from the multivariate KDE in Section 2.1 with that of the observational flux values in several bands of the 3D-*HST* catalogue from redshift range $0.9 < z < 1.1$. The τ SFH model in BAGPIPES was chosen because it was also the model used to fit the observations by FAST, and since the multivariate KDE of the parameter space was based on these FAST values, comparison of these distributions acts as an important validation step. The flux values in the 3D-*HST* catalogue are scaled to an AB magnitude zero-point of 25.0, so a scaling factor of $\frac{1}{10^{0.44}} = 0.3631$ was employed before comparing it to the BAGPIPES generated flux values, which are in microjanskys (μJy). Fig. 3 shows that the BAGPIPES forward model is able to generate generally realistic 3D-*HST* observational flux values at the population level. Table 2 provides summary statistics on the flux distribution in each band; some differences can be attributed to the different IMF used by the two approaches (Kroupa and Chabrier, respectively).

3 BAYESIAN PARAMETER ESTIMATION AND MODEL EVIDENCE CALCULATION

3.1 Bayesian inference

The goal of our analysis, which is carried out within the realm of Bayesian statistics, is to find the right balance between keeping the model simple and providing a good fit to the data, a concept often known as Occam’s razor. Quantitatively, this translates into applying the model comparison formalism. The evidence (the denominator

in Bayes theorem, sometimes also called marginal likelihood) is sensitive to not only the complexity of the parameter and prior space, but also the quality of the fit. Therefore, it can be used as a model comparison tool: models that have larger evidence are preferable, not because they necessarily fit the data better, but because they strike the right balance between complexity and accuracy.

The hypothesis we test in the paper is whether the evidence (or more accurately, the ratio of the evidence of two models, called the Bayes factor) can be successfully used to ‘recognize’, in a purely data-driven manner, the correct model. To this end, we generate simulated data for a five-parameter or six-parameter model (the ‘correct model’ in the remainder of the paper), and then we ask: if we were to fit the data with the correct model (five or six parameters, respectively) or a more simplistic model (four or five parameters), would the Bayes factor favour the more complex model? Would we be able to recognize the ‘true’ complexity using the Bayes factor? If our tests are successful, this indicates a possible path to derive the ‘true’ model complexity from the data, giving us a way to discriminate between useful, meaningful parameters, and ‘noisy’ ones that add complexity without sufficiently improving the quality of the fit.

We start with estimation of the posterior distribution given a model \mathbf{M} , its parameters Θ , and a set of data \mathbf{D} . We follow the notational convention of Speagle (2020). We use Bayes Rule:

$$P(\Theta|\mathbf{D}, \mathbf{M}) = \frac{P(\mathbf{D}|\Theta, \mathbf{M})P(\Theta|\mathbf{M})}{P(\mathbf{D}|\mathbf{M})}, \quad (5)$$

where $P(\mathbf{D}|\Theta, \mathbf{M})$ is the likelihood of the data given the parameters in the model, $P(\Theta|\mathbf{M})$ is the prior for the parameters, and

$$P(\mathbf{D}|\mathbf{M}) = \int_{\Theta_{\Omega}} P(\mathbf{D}|\Theta, \mathbf{M})P(\Theta|\mathbf{M}) d\Theta \quad (6)$$

is the marginal likelihood, also known as the *evidence*, integrated over the entire parameter space Ω . Rewriting Bayes Rule using shorthand notation, we have

$$P(\Theta_{\mathbf{M}}) = \frac{L(\Theta_{\mathbf{M}})\pi(\Theta_{\mathbf{M}})}{Z_{\mathbf{M}}} \quad (7)$$

for a given model \mathbf{M} , where

$$Z_{\mathbf{M}} = \int_{\Theta_{\Omega}} L(\Theta_{\mathbf{M}})\pi(\Theta_{\mathbf{M}}) d\Theta. \quad (8)$$

The evidence $Z_{\mathbf{M}}$ can therefore be understood simply as the likelihood space multiplied by the prior space integrated over the entire parameter space.

3.2 Nested sampling

Numerous MCMC methods have been used to successfully generate samples and their associated weights for posterior estimation. However, when the posterior is multimodal or there are significant parameter degeneracies, traditional MCMC methods often struggle due to algorithmic limitations, computational expense, or both. Moreover, calculating evidence to a sufficient degree of accuracy using MCMC methods is often challenging if not impossible (e.g. Feroz & Hobson 2008).

Nested sampling is an alternative to MCMC sampling developed by Skilling (2006) that can perform both accurate evidence calculation and effective posterior estimation while avoiding the issues noted above. It harnesses the relationship between the likelihood and prior, and samples from the prior space subject to ever-increasing lower bounds on the likelihood. Each successful iteration of the nested sampling algorithm results in an estimate of the remaining prior ‘volume’ as well as a likelihood evaluation, which are used to calculate

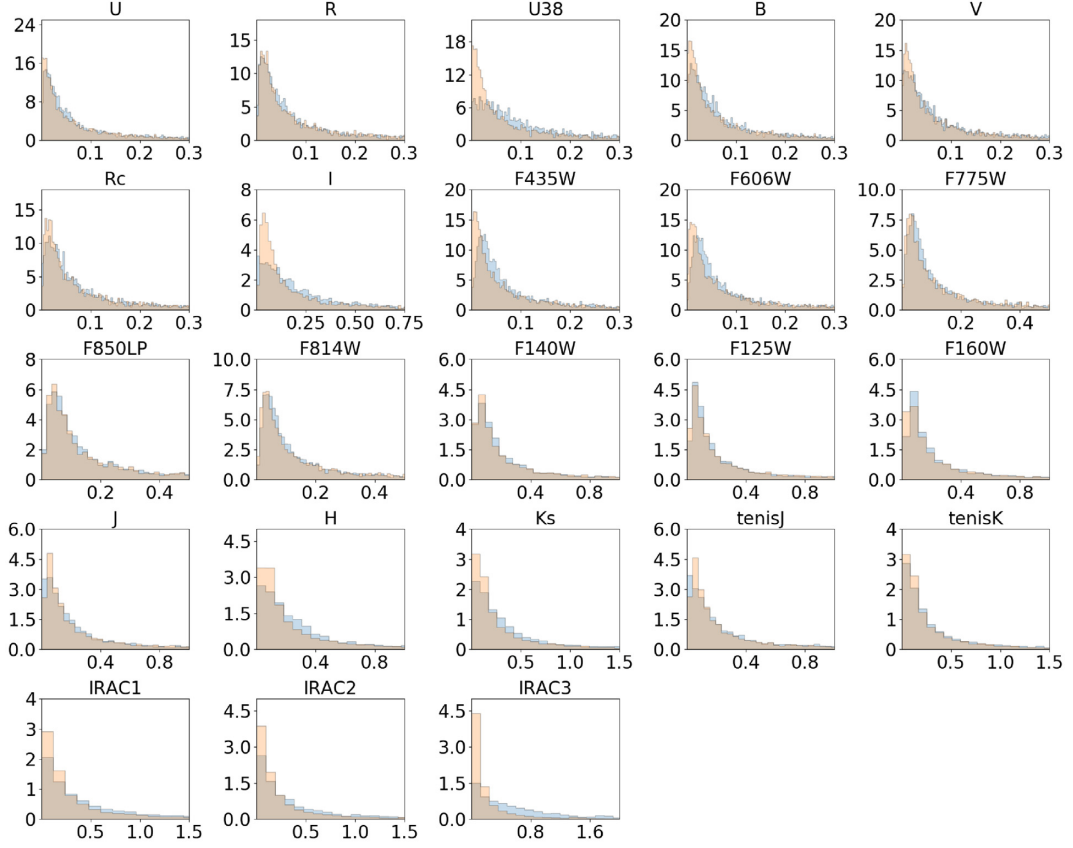


Figure 3. The blue histogram represents the distribution of observed flux values for each band in the reference 3D-*HST* sample described in the text. The orange histogram represents the distribution of flux values for each band in our simulated catalogue, obtained using BAGPIPES to generate the photometry.

a fraction of the total evidence, and by extension a posterior sample. This process creates nested ‘shells’ of iso-likelihood contours, hence the name. BAGPIPES use the PYTHON wrapper PYMULTINEST (Buchner 2016) to utilize the sampling routines in MULTINEST.

To start, we can rewrite equation (8) as

$$Z = \int_{\Theta_{\Omega}} P(\mathbf{D}|\Theta)P(\Theta) d\Theta = \int_{\Theta_{\Omega}} L(\Theta)\pi(\Theta) d\Theta. \quad (9)$$

The estimate of equation (9) is simply its expected value. Note that in mathematical statistics, for a positive random variable, the area above a cumulative distribution function (CDF) and below 1 is the expected value. In other words, if the probability density function (PDF) of X is f and the CDF of X is F , then

$$E(X) \equiv \int_0^{\infty} (1 - F(x))dx = \int_0^{\infty} xf(x)dx. \quad (10)$$

Furthermore, let the CDF of $L(\Theta) = \lambda$ above be defined as

$$F(\lambda) \equiv \int_{L(\Theta) < \lambda} \pi(\Theta)d\Theta, \quad (11)$$

so estimating equation (9) is equivalent to estimating $\int_0^{\infty} (1 - F(x))dx$, where F is defined in equation (11).

Then, if we define the prior volume X as $dX = \pi(\Theta)d\Theta$, such that

$$X(\lambda) = 1 - F(\lambda) = \int_{L(\Theta) > \lambda} \pi(\Theta)d\Theta \quad (12)$$

(noting the change in equality in the integration bounds) then equation (9) is equal to $\int_0^{\infty} X(\lambda)d\lambda$. Lastly, inverting $X(\lambda)$ and

evaluating the integral of the inverse from 0 to 1, we have

$$\int_0^1 X^{-1}(a) = \int_{\Theta_{\Omega}} L(\Theta)\pi(\Theta) d\Theta = \int_0^1 L(X)dx, \quad (13)$$

where $L(X)$, the inverse of equation (12), is a monotonically decreasing function of X . Thus, the evidence integral can be written as

$$Z = \int_0^1 L(X)dx, \quad (14)$$

and its calculation is recast from a difficult multidimensional integration over Θ to an easier 1D integration over X .

The next step in evidence calculation involves computing the likelihoods $L_i = L(X_i)$, where X_i decreases as $i = 0, 1, 2, \dots, M$ so that

$$0 < X_M < \dots < X_2 < X_1 < X_0 = 1. \quad (15)$$

The evidence is then found using quadrature as

$$Z = \sum_{i=1}^M L(\Theta_i)w_i, \quad (16)$$

where $w_i = \frac{1}{2}(X_{i-1} - X_{i+1})$ (Feroz, Hobson & Bridges 2009).

3.3 Calculating the evidence

The process of calculating the evidence involves several steps. First, N ‘live’ points are sampled from the prior space and corresponding likelihoods are calculated for each live point. The likelihoods are then

Table 2. Median values and representative percentiles, in each band, for the reference 3D-*HST* sample at $z \sim 1$ described in the text and the simulated galaxies generated by BAGPIPES using the KDE distributions from Section 2.1. We find good agreement between the observed and simulated catalogue in all bands. Some residual differences can be attributed to the different IMF used by the two approaches (Kroupa and Chabrier, respectively).

Band	Median	16%	84%
U _{3D-HST}	0.0488	0.0127	0.2200
U _{BP}	0.0434	0.0097	0.2070
R _{3D-HST}	0.0572	0.0155	0.2929
R _{BP}	0.0557	0.0147	0.3005
U38 _{3D-HST}	0.0842	0.0230	0.2872
U38 _{BP}	0.0438	0.0097	0.2081
B _{3D-HST}	0.0539	0.0140	0.2320
B _{BP}	0.0440	0.0100	0.2163
V _{3D-HST}	0.0571	0.0149	0.2569
V _{BP}	0.0463	0.0111	0.2330
RC _{3D-HST}	0.0643	0.0174	0.3104
RC _{BP}	0.0540	0.0144	0.2879
I _{3D-HST}	0.1894	0.0494	0.8176
I _{BP}	0.1132	0.0349	0.6181
F435W _{3D-HST}	0.0552	0.0194	0.2222
F435W _{BP}	0.0445	0.0100	0.2156
F775W _{3D-HST}	0.0582	0.0207	0.2589
F775W _{BP}	0.0489	0.0126	0.2546
F606W _{3D-HST}	0.0920	0.0325	0.4542
F606W _{BP}	0.0914	0.0260	0.4962
F814W _{3D-HST}	0.1231	0.0411	0.6331
F814W _{BP}	0.1189	0.0369	0.6470
F850LP _{3D-HST}	0.1031	0.0354	0.5044
F850LP _{BP}	0.0993	0.0293	0.5400
F140W _{3D-HST}	0.1668	0.0551	0.8863
F140W _{BP}	0.1580	0.0491	0.9222
F125W _{3D-HST}	0.1606	0.0553	0.8934
F125W _{BP}	0.1527	0.0472	0.8806
F160W _{3D-HST}	0.1701	0.0622	0.9500
F160W _{BP}	0.1512	0.0485	0.8965
J _{3D-HST}	0.1757	0.0456	0.9135
J _{BP}	0.1523	0.0465	0.8834
H _{3D-HST}	0.2362	0.0609	1.0381
H _{BP}	0.1554	0.0499	0.9346
Ks _{3D-HST}	0.2822	0.0715	1.4908
Ks _{BP}	0.1740	0.0562	1.0789
tenisJ _{3D-HST}	0.1907	0.0439	1.0315
tenisJ _{BP}	0.1539	0.0473	0.9146
tenisK _{3D-HST}	0.2265	0.0537	1.4288
tenisK _{BP}	0.1739	0.0562	1.0794
IRAC1 _{3D-HST}	0.3669	0.0708	2.1866
IRAC1 _{BP}	0.1886	0.0590	1.1943
IRAC2 _{3D-HST}	0.2908	0.0518	1.7301
IRAC2 _{BP}	0.1426	0.0441	0.9038
IRAC3 _{3D-HST}	0.5486	0.1055	2.5733
IRAC3 _{BP}	0.1003	0.0306	0.6496

ordered L_0, L_1, \dots, L_N , where L_0 is the smallest likelihood value. We will assume henceforth that the prior volume at this point in the sampling is $X_0 = \mathbf{1}$ (i.e. the unit hypercube). Next, L_0 is removed from the set of live points and becomes a ‘dead’ point. A new point is then drawn from the prior space, and if its likelihood value is greater than new smallest likelihood value (i.e. $L_{\text{new}} > L_{0_{\text{new}}}$), it is added to the pool of live points. If, on the other hand, the new point’s likelihood value is less than the new smallest likelihood value (i.e. $L_{\text{new}} < L_{0_{\text{new}}}$), then it is discarded and added to the pool of ‘dead’ points. At each new live point (each ‘successful’ iteration), a new

iso-likelihood contour is defined with which to constrain future live points, the prior space X shrinks, and the evidence value increases.

Nested sampling is therefore a form of rejection sampling, and indeed the biggest challenge in nested sampling is sampling a new point from the prior space to be added to the pool of live points subject to the constraint $L_{\text{new}} > L_{0_{\text{new}}}$. The trick is to find a way to both efficiently sample from the prior volume defined in equation (18) and also to fully explore the likelihood space while satisfying this constraint.

3.4 Posterior estimation

If we are interested in posterior samples, we use the dead points from the nested sampling process (the points with the lowest likelihood values). This is written as

$$p(\Theta_i) = \frac{L(\Theta_i)w_i}{Z}, \quad (17)$$

where Z and w_i are given in equation (16). One can then construct relevant marginal distributions and descriptive statistics such as means and percentile values.

3.5 Parameter estimation in SED fitting using MULTINEST

We now return to our simulated data, where we used samples from the KDEs of the parameter space inferred by FAST for the 3D-*HST* catalogue as inputs to our galaxy catalogues.

For the purposes of realistic Bayesian model fitting and parameter estimation, we include noise, modelled after the 3D-*HST* catalogue, in the simulated observations created in Section 2. This is achieved by sampling from the KDE of the observational error in each band and randomly adding or subtracting the sampled amount to/from the corresponding observational flux. The minimum KDE error sample values are restricted to each of the minimum observational error values of the 23 bands from the selection of 4567 galaxies, and the maximum KDE error sample values are restricted to the 95th percentile of the KDE errors.

The original prior configuration is shown in Table 7. We note that age prior is adaptive, in the sense that the upper bound is restricted to be no larger than the age of the universe at the relevant redshift.

4 BAYES FACTORS IN MODEL COMPARISON

4.1 Calculating evidence

We can use the evidence values of a given pair of models, along with prior beliefs about the models, to calculate the Bayes factor. We have

$$Z_{\mathbf{M}} \equiv P(\mathbf{D}|\mathbf{M}) = \int_{\Theta_{\mathbf{M}}} P(\mathbf{D}|\Theta, \mathbf{M})P(\Theta|\mathbf{M}) d\Theta, \quad (18)$$

where $Z_{\mathbf{M}_i}$ is the evidence and $\pi(M_i)$ is the prior belief for M_i , $i = (0, 1)$. Henceforth, we will assume $\pi(M_0) = \pi(M_1) = \frac{1}{2}$, so the Bayes factor is

$$B_{01} = \frac{P(\mathbf{D}|\mathbf{M}_0)}{P(\mathbf{D}|\mathbf{M}_1)} = \frac{Z_{\mathbf{M}_0}}{Z_{\mathbf{M}_1}}. \quad (19)$$

Interpretation of the Bayes factor can be done through the Jeffreys’ scale, which is an empirically calibrated scale (adapted from Jeffreys 1998 and Trotta 2017). Small absolute values of B_{01} between model \mathbf{M}_0 and model \mathbf{M}_1 indicate that the data and prior structure do not prefer either model. Large negative values express a preference for the \mathbf{M}_0 model (in our formalism, the simpler model), large positive

Table 3. This table shows the parameter fits for Scenario 1, which includes fits and summary statistics for both the five parameter values (the first value in a given column) and four-parameter model (the second value in a given column).

	Parameter	% of objects in 68% region	% of objects in 95% region	Median bias	Scatter	OLF
0×	Massburst	63	88	0.01	0.06	0.0
Noise reduction	Age _{delayed} τ	63/45	89/69	0.00/−0.01	0.05/0.06	0.00/0.00
	Mass _{delayed} τ	56/30	86/53	−0.01/0.02	0.06/0.04	0.00/0.01
	τ	66/46	90/69	0.04/0.01	0.09/0.11	0.09/0.11
	A _v	62/60	89/85	0.00/0.01	0.13/0.14	0.12/0.17
5×	Massburst	47	74	0.00	0.05	0.00
Noise reduction	Age _{delayed} τ	56/20	83/36	0.00/−0.03	0.04/0.05	0.01/0.00
	Mass _{delayed} τ	50/12	77/24	0.00/0.02	0.04/0.02	0.00/0.00
	τ	50/17	81/36	0.02/−0.05	0.08/0.09	0.07/0.04
	A _v	63/48	87/70	0.00/0.00	0.10/0.12	0.08/0.13
10×	Massburst	44	73	0.00	0.04	0.00
Noise reduction	Age _{delayed} τ	50/12	76/24	0.00/−0.03	0.03/0.05	0.01/0.00
	Mass _{delayed} τ	46/7	76/14	0.00/0.01	0.03/0.01	0.00/0.00
	τ	48/12	79/25	0.01/−0.06	0.07/0.08	0.05/0.01
	A _v	62/42	86/59	0.00/0.00	0.08/0.11	0.06/0.14
20×	Massburst	45	71	0.00	0.03	0.00
Noise reduction	Age _{delayed} τ	49/9	74/15	0.00/−0.03	0.03/0.04	0.01/0.00
	Mass _{delayed} τ	46/3	74/8	0.00/0.01	0.02/0.01	0.00/0.00
	τ	45/9	76/18	0.00/−0.06	0.06/0.07	0.03/0.00
	A _v	62/32	87/47	0.00/0.01	0.07/0.11	0.04/0.15

Table 4. This table shows the parameter fits for Scenario 2, which includes fits and summary statistics for both the six-parameter values (the first value in a given column) and five-parameter model (the second value in a given column).

	Parameter	% of objects in 68% region	% of objects in 95% region	Median bias	Scatter	OLF
0×	Age _{burst}	66	95	0.02	0.03	0.02
Noise reduction	Massburst	73/55	94/78	0.03/−0.01	0.06/0.07	0.00/0.00
	Age _{delayed} τ	69/61	93/85	−0.01/0.01	0.05/0.05	0.00/0.01
	Mass _{delayed} τ	69/56	92/79	−0.02/0.00	0.05/0.06	0.00/0.00
	τ	65/62	90/84	0.05/0.03	0.09/0.09	0.12/0.08
5×	A _v	64/61	87/86	0.00/0.01	0.14/0.15	0.13/0.17
	Age _{burst}	56	88	0.01	0.03	0.00
	Massburst	62/34	87/53	0.02/−0.05	0.05/0.08	0.00/0.00
	Age _{delayed} τ	59/36	87/62	0.00/0.00	0.04/0.05	0.00/0.01
Noise reduction	Mass _{delayed} τ	62/35	86/55	−0.01/0.01	0.04/0.03	0.00/0.00
	τ	60/39	87/64	0.03/0.00	0.08/0.09	0.06/0.06
	A _v	56/54	84/78	0.00/0.00	0.11/0.12	0.08/0.13
	Age _{burst}	57	87	0.01	0.03	0.00
10×	Massburst	55/22	84/37	0.01/−0.07	0.05/0.08	0.00/0.00
Noise reduction	Age _{delayed} τ	55/25	85/50	0.00/0.00	0.04/0.04	0.00/0.01
	Mass _{delayed} τ	56/26	85/44	0.00/0.01	0.04/0.02	0.00/0.00
	τ	53/27	84/52	0.02/−0.01	0.07/0.08	0.05/0.05
	A _v	56/50	84/74	0.00/0.00	0.10/0.11	0.08/0.11
20×	Age _{burst}	59	87	0.01	0.03	0.00
	Massburst	50/14	82/24	0.00/−0.10	0.05/0.08	0.00/0.00
	Age _{delayed} τ	53/21	81/37	0.00/0.00	0.04/0.04	0.01/0.01
	Mass _{delayed} τ	53/16	83/27	0.00/0.01	0.03/0.02	0.00/0.00
Noise reduction	τ	49/19	82/35	0.01/−0.01	0.07/0.07	0.04/0.03
	A _v	55/45	80/67	0.00/0.00	0.09/0.10	0.06/0.10

values denote a preference for the \mathbf{M}_1 model (in our formalism, the more complex model). The reference threshold values we use are listed in Table 8.

4.2 Bayes factors in Savage–Dickey density ratio

The Savage–Dickey density ratio (SDDR) is an analytical approach to calculating Bayes factors for model selection (Dickey 1971,

Verdinelli & Wasserman 1995) and was first introduced for use in a cosmological context by Trotta (2007). Given an N parameter model M_1 (where $N \geq 2$) with free parameters $\theta = \theta_1, \dots, \theta_N$ and a nested $N - 1$ parameter model M_0 with free parameters $\theta_1, \dots, \theta_{N-1}$ and fixed $\theta_N = \theta_*$, the SDDR can be calculated using equation (19) by

$$B_{01} = \frac{P(\theta | \mathbf{D}, \mathbf{M}_1)}{\pi(\theta | \mathbf{M}_1)} \Big|_{\theta_N = \theta_*} \quad (20)$$

Table 5. This table shows the parameter fits for Scenario 3, which includes fits and summary statistics for both the six-parameter values (the first value in a given column) and four-parameter model (the second value in a given column).

	Parameter	% of objects in 68% region	% of objects in 95% region	Median bias	Scatter	OLF
0× Noise reduction	Age _{burst}	66	95	0.02	0.03	0.02
	Mass _{burst}	73	94	0.03	0.06	0.00
	Age _{delayed} τ	69/46	93/73	−0.01/0.00	0.05/0.06	0.00/0.01
	Mass _{delayed} τ	69/56	92/79	−0.02/0.02	0.05/0.04	0.00/0.01
	τ	65/49	90/74	0.05/0.03	0.09/0.10	0.12/0.11
5× Noise reduction	A_V	64/63	87/87	0.00/0.01	0.14/0.14	0.13/0.16
	Age _{burst}	56	88	0.01	0.03	0.00
	Mass _{burst}	62	87	0.02	0.05	0.00
	Age _{delayed} τ	59/28	87/49	0.00/0.00	0.04/0.05	0.00/0.00
	Mass _{delayed} τ	62/15	86/32	−0.01/0.02	0.04/0.01	0.00/0.00
10× Noise reduction	τ	60/28	87/50	0.03/0.01	0.08/0.09	0.06/0.05
	A_V	56/51	84/75	0.00/0.00	0.11/0.12	0.08/0.13
	Age _{burst}	57	87	0.01	0.03	0.00
	Mass _{burst}	55	84	0.01	0.05	0.00
	Age _{delayed} τ	55/19	85/39	0.00/0.00	0.04/0.04	0.00/0.00
20× Noise reduction	Mass _{delayed} τ	56/9	85/18	0.00/0.02	0.04/0.01	0.00/0.00
	τ	53/21	84/41	0.02/−0.01	0.07/0.07	0.05/0.04
	A_V	56/49	84/70	0.00/0.00	0.10/0.11	0.08/0.12
	Age _{burst}	59	87	0.01	0.03	0.00
	Mass _{burst}	50	82	0.00	0.05	0.00
Noise reduction	Age _{delayed} τ	53/16	81/30	0.00/0.00	0.04/0.03	0.01/0.00
	Mass _{delayed} τ	53/6	83/11	0.00/0.02	0.03/0.01	0.00/0.00
	τ	49/15	82/28	0.01/−0.02	0.07/0.07	0.04/0.02
	A_V	55/43	80/65	0.00/0.00	0.09/0.10	0.06/0.12

Evaluation of the SDDR therefore requires estimation of the PDF for both the marginalized posterior and the prior of θ_N to calculate the height of the densities at θ_* . Log spline density estimation (Stone & Koo 1986), in which the logarithm of a PDF is modelled using polynomial spline, has been found to perform well for estimation of the SDDR under general conditions (Wagenmakers et al. 2010) and tends to outperform kernel-based density estimation (Wetzels et al. 2009). We will use the R package `logspline` (Stone et al. 1997) to calculate the log spline density estimates in order to derive the SDDR.

The SDDR requires that the two competing models M_0 and M_1 are nested and that the priors are separable, i.e.

$$\pi(\theta|M_1) = \pi(\theta_1|M_1) \pi(\theta_2|M_1) \dots \pi(\theta_N|M_0). \quad (21)$$

In addition, the prior for the fitted parameters of M_1 should equal the prior for the fitted parameters under M_0 in order for M_1 to reduce to M_0 . Importantly, in order to effectively calculate the SDDR, a sufficient number of posterior samples are needed in the neighbourhood of θ_* . We emphasize that, depending on the dimensionality of the SED model being fit, the number of samples and coverage of the posterior space scales with the number of live points used in MULTINEST. To obtain sufficient coverage of posterior samples, we choose 1500 live points. The sampling efficiency in MULTINEST is inversely proportional to the enlargement factor of the volume of the ellipsoid region. A greater enlargement factor therefore explores more completely the likelihood space. As we are interested in not only accurate evidence values but also accurate posterior estimates, we use a sampling efficiency of 0.25 instead of the default 0.3 in MULTINEST.

We note that the Bayes factor is effectively the height of the posterior density of θ_N divided by the height of the prior density of θ_N evaluated at θ_* , the choice of prior has a great impact upon the derivation of the SDDR. This caveat notwithstanding, the

computational resources needed to calculate the SDDR are minimal, so it is an appealing alternative to a nested sampling approach to calculating Bayes factors.

5 RESULTS

We analyse different simulation scenarios in order to assess, on one hand, the impact of incorrect assumptions on the derived physical parameters, and on the other, the reliability of the Bayes factor as a tool to empirically discern the model complexity from the data. While the latter is inherently new to this paper, the former is a necessary validation step for the SED-fitting process.

Our simulations involve the following approaches. In all cases, we begin by assuming a noise profile matching the one of the 3D-*HST* galaxies described earlier in paper. Since we found in our preliminary studies that S/N was the main factor driving the effectiveness of the Bayes factor, we then reduce the noise (maintaining the same profile) of the mock galaxy catalogue by factors of 5, 10, and 20. The input parameters for the simulations are the same throughout this process to ensure a fair comparison in the presence of modelling systematics.

5.1 How many stellar populations?

In choosing our set of simulations, we are motivated by the question ‘can we detect multiple major episodes of star formation?’. In fact, many physically motivated models of star formation, including semi-analytic models and hydrodynamical simulations (e.g. Genel et al. 2014; Somerville & Davé 2015; Pillepich et al. 2018), show that the stellar assembly of galaxies might be well described by a ‘smooth’ component, describing the process of star formation in isolation, and a stochastic component, driven by interactive events such as mergers and winds. We describe each scenario in detail below.

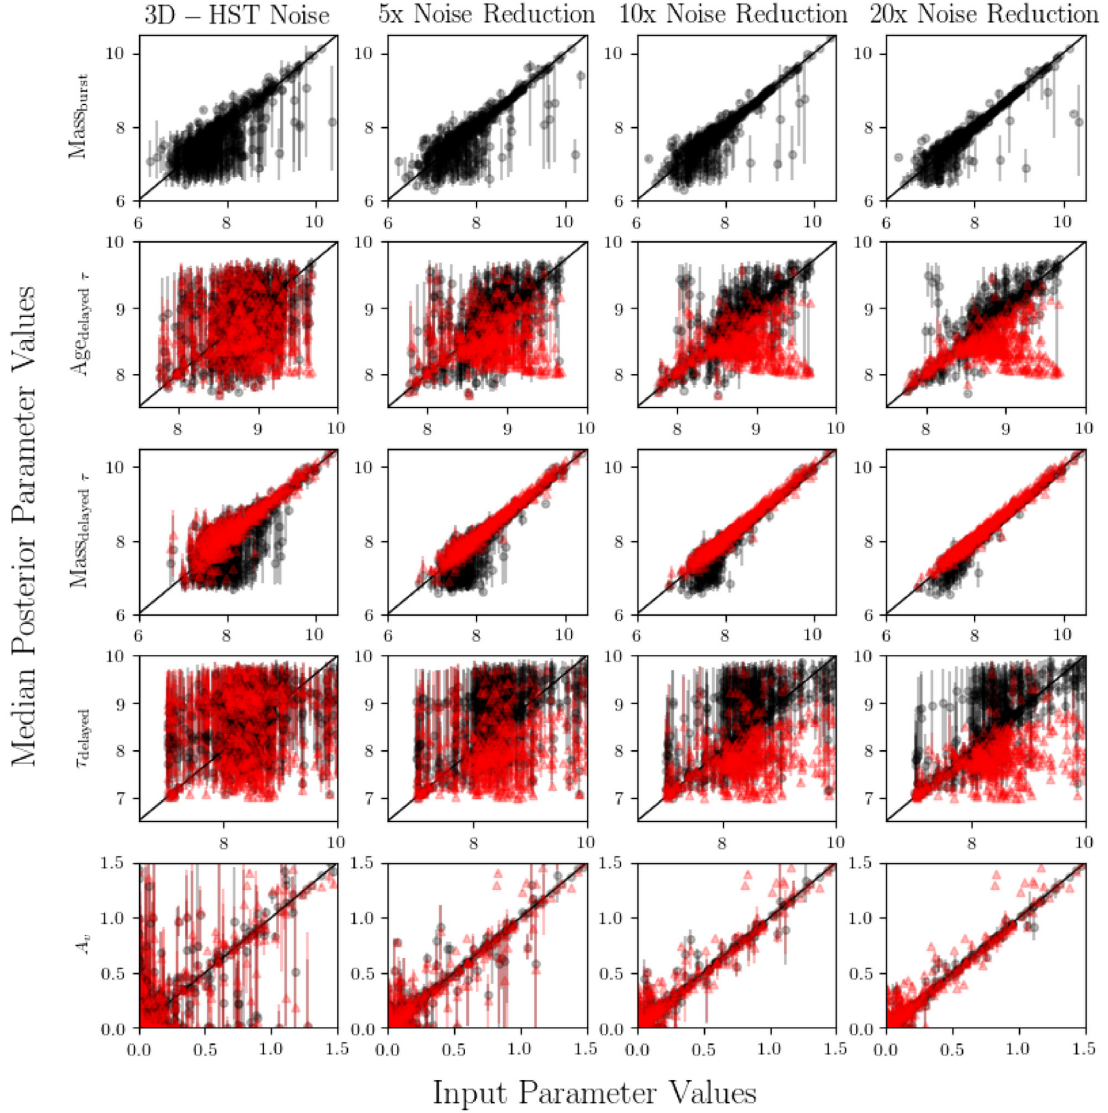


Figure 4. The black circles represent the input versus median posterior parameter values for the five-parameter model in Scenario 1. The black vertical lines represent the 16–84 per cent credible intervals the parameters for the same model. The red triangles, on the other hand, are the input versus median posterior parameter values for the four-parameter model in Scenario 1. The red vertical lines represent the 16–84 per cent credible intervals the parameters for the same model.

(i) **Scenario 1:** In our first test, we generate models using a five-parameter model, with a ‘main’ stellar population component described by a delayed τ model, with parameters stellar age, stellar mass, τ , and dust attenuation, and an additional recent burst component of star formation with age = 0.1 Gyr and variable stellar mass. We then perform parameter fits and evidence calculation both with a five-parameter model (\mathbf{M}_1 for the purposes of Bayesian model comparison) with the parameters listed above, and a four-parameter model where the burst mass is fixed at $\log_{10}M/M_{\odot} = 6.5$. In practice, the latter corresponds to a one-component stellar population, because the median mass for the entire sample is $\log_{10}M/M_{\odot} \simeq 8.4$ and so the burst mass contribution to the total mass is at the per cent level, but we can not simply set the mass in the burst to zero because that would not be part of the prior parameter space in the higher model, which is a necessary condition of nested models. These conditions aim to answer the question: ‘Can we detect a recent burst of star formation?’

(ii) **Scenario 2:** In our second test, we generate models using a six parameter model, where the ‘main’ stellar population component is still described by a delayed τ model, with parameters stellar age, stellar mass, τ , and dust attenuation, and an additional burst component of star formation with variable age (between 0.05 and 0.9 Gyr) and variable stellar mass. We then perform parameter fits and evidence calculation both with a six parameter model (\mathbf{M}_1 for the purposes of Bayesian model comparison) with the parameters listed above, and a five-parameter model where we vary the stellar mass, but fix the burst age at 0.1 Gyr. These conditions aim to answer the question: ‘What happens if we incorrectly estimate the age of a secondary burst of star formation?’

(iii) **Scenario 3:** Our third test involves generating models using the same six parameter model as in Scenario 2, and a four-parameter model where the burst age is fixed at 0.1 Gyr and the burst mass is fixed at $\log_{10}M/M_{\odot} = 6.75$. These conditions aim to answer

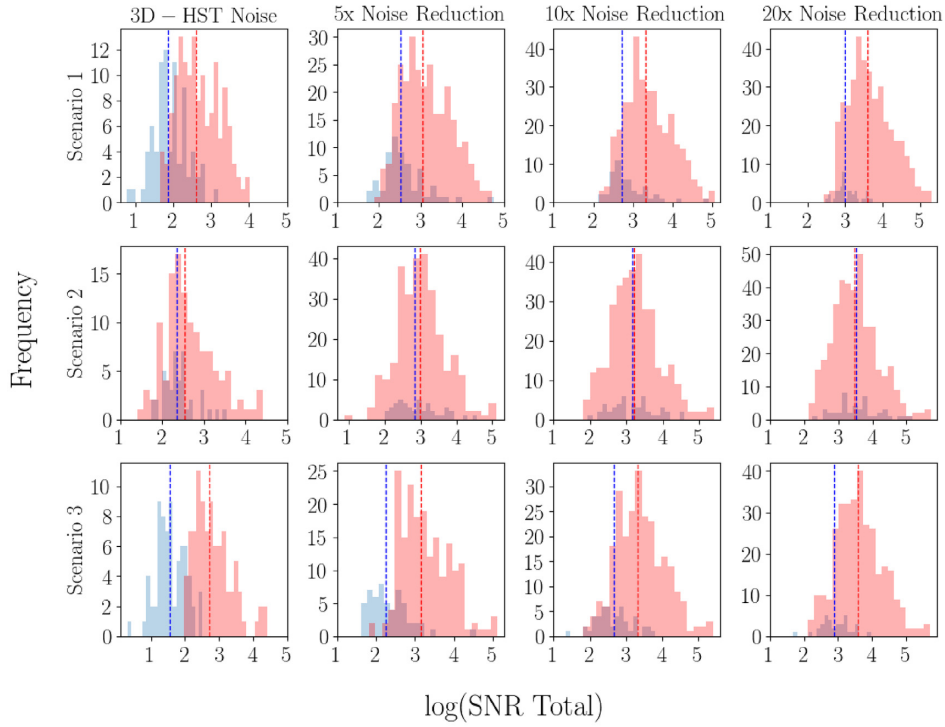


Figure 5. The blue histograms represent the distribution of cases in which the Bayes factor incorrectly prefers the simpler model ($\ln B_{01} < -1$), and the red histograms represent the distribution of cases in which the Bayes factor correctly prefers the more complex model ($\ln B_{01} > 1$). We show results for Scenarios 1–3 at each noise level, as a function of the total S/N of each SED. The total S/N is the sum of every band’s photometric flux measurement divided by its 1σ error measurement. The dotted vertical line represents the median value for each distribution. The S/N of each SED is an important factor in determining the method’s success in Scenarios 1 and 3.




Figure 6. The blue histograms represent the distribution of cases in which the Bayes factor incorrectly prefers the simpler model ($\ln B_{01} < -1$), and the red histograms represent the distribution of cases in which the Bayes factor correctly prefers the more complex model ($\ln B_{01} > 1$), for each scenario and S/N level. The horizontal axis of each subplot is the age separation, which is defined as the time between the onset of star formation in the delayed τ SFH component and the burst SFH component (in Gyr). The dotted vertical line represents the median value for each distribution in each subplot.




Figure 7. The black circles represent the input versus median posterior parameter values for the five-parameter model in Scenario 1. The black vertical lines represent the 16–84 per cent credible intervals the parameters for the same model. The red triangles, on the other hand, are the input versus median posterior parameter values for the five-parameter model in Scenario 4. The red vertical lines represent the 16–84 per cent credible intervals the parameters for the same model.

the question: ‘Can we detect a secondary burst component of star formation, independent of age?’

(iv) **Scenario 4:** Our fourth and last test attempts to test the impact of priors, in which we employ the same model as in Scenario 1 but with linear rather than \log_{10} priors on stellar age, τ , and dust attenuation parameters.

5.2 Posterior fits

For scenarios 1–3, which involve different input parameter and fit parameters structure, we report the results of the parameter estimation procedure through a series of summary statistics for the parameter fits in Tables 3, 4, and 5 following the conventions of Acquaviva et al. (2015). For each parameter, we report the mean bias and scatter, as well as the fraction of outliers (defined as the percentage of objects for which the input and estimated values differ at more than 15 per cent level). Furthermore, in order to assess the accuracy of the reported

posteriors, we report the fraction of objects for which the input values lie within the reported 68 per cent and 95 per cent intervals.

For Scenario 1, we also show a set of scatter plots, displaying the input (true) values versus estimated values; results from the more complicated model are shown in black, and from the simpler model in red, in Fig. 4. The behaviour of Scenarios 2 and 3 follow a very similar pattern.

Our results indicate that in all scenarios, the posterior width is underestimated when the simpler model is used. This becomes more stringent as the noise level is reduced: in the four-parameter model (and at a modest level, even in the five-parameter model) of Scenario 1, the size of the uncertainties, i.e. width of the posteriors, are very severely underestimated at the greatest amount of noise reduction.

In Scenarios 1–3, the bias, scatter, and outlier fraction for the simpler and more complicated models are comparable. They are generally present at the few per cent level. Scenario 4 parameter fits are noteworthy for the significant outlier fraction for both τ and A_v

Table 6. This table shows the parameter fits for Scenario 1 versus Scenario 4, which includes fits and summary statistics for both the five-parameter values with \log_{10} priors (the first value in a given column) and five-parameter model with linear priors (the second value in a given column).

	Parameter	% of objects in 68% region	% of objects in 95% region	Median bias	Scatter	OLF
0×	Mass _{burst}	63/54	88/78	0.01/0.01	0.06/0.06	0.00/0.01
Noise reduction	Age _{delayed} τ	63/36	89/70	0.00/0.04	0.05/0.06	0.00/0.06
	Mass _{delayed} τ	56/56	86/79	−0.01/0.00	0.06/0.06	0.00/0.01
	τ	66/7	90/20	0.04/0.16	0.09/0.10	0.09/0.55
	A_v	62/28	89/48	0.00/0.05	0.13/0.14	0.12/0.37
	Mass _{burst}	47/27	74/53	0.00/0.01	0.05/0.04	0.00/0.00
5×	Mass _{burst}	47/27	74/53	0.00/0.01	0.05/0.04	0.00/0.00
Noise reduction	Age _{delayed} τ	56/26	83/52	0.00/0.03	0.04/0.05	0.01/0.05
	Mass _{delayed} τ	50/37	77/64	0.00/−0.01	0.04/0.04	0.00/0.00
	τ	50/9	81/20	0.02/0.15	0.08/0.1	0.07/0.49
	A_v	63/24	87/52	0.00/0.03	0.10/0.11	0.08/0.30
	Mass _{burst}	44/22	73/38	0.00/0.01	0.04/0.04	0.00/0.00
10×	Mass _{burst}	44/22	73/38	0.00/0.01	0.04/0.04	0.00/0.00
Noise reduction	Age _{delayed} τ	50/21	76/38	0.00/0.03	0.03/0.05	0.01/0.05
	Mass _{delayed} τ	46/32	76/56	0.00/0.00	0.03/0.04	0.00/0.00
	τ	48/11	79/21	0.01/0.14	0.07/0.11	0.05/0.44
	A_v	62/24	86/51	0.00/0.02	0.08/0.10	0.06/0.25
	Mass _{burst}	45/22	71/38	0.00/0.01	0.03/0.04	0.00/0.00
20×	Mass _{burst}	45/22	71/38	0.00/0.01	0.03/0.04	0.00/0.00
Noise reduction	Age _{delayed} τ	49/21	74/39	0.00/0.03	0.03/0.05	0.01/0.05
	Mass _{delayed} τ	46/33	74/56	0.00/0.00	0.02/0.04	0.00/0.00
	τ	45/9	76/21	0.00/0.14	0.06/0.11	0.03/0.44
	A_v	62/32	87/50	0.00/0.02	0.07/0.10	0.04/0.25

parameters, and significant bias in τ compared to those parameters Scenarios 1–3. Due to this, evidence values derived from Scenario 4 should not be considered reliable.

5.3 Bayes factors and model comparison

We can now turn to the analysis of the Bayes factors, i.e. the natural log ratio of pieces of evidence of the more complex (correct) versus simpler (incorrect) model for the three scenarios; our results are shown in Fig. 8. Each set of multihistograms (column) shows the Bayes factor; the columns are organized according to the Jeffrey scale described in Table 8. The leftmost column displays cases in which the simpler model is preferred with moderate confidence ($-2.5 < \ln B_{01} < -1$), and moving to the right, we find inconclusive cases, and cases in which the more complex model is preferred with weak, moderate, and high confidence, respectively.

In an ideal case, the Bayes factor would always prefer the more complex model since we know that this is the correct underlying model; however, the most crucial issue is whether the Bayes factor can actually be misleading, preferring the simpler (incorrect) model with high confidence.

In Scenario 1, at the 3D-*HST* noise level, we see that the Bayes factor is indecisive in the majority of cases (56 per cent of the total), and is able to ‘pick’ the more complex model with high confidence ($\ln B_{01} > 2.5$) only in 15 per cent of the cases. Importantly, though, there are only a handful of failing cases (~ 6 per cent of the total, not shown) in which the Bayes factor strongly prefers the simpler model, and a fraction of cases (16 per cent of the total) in which the Bayes factors indicates a moderate preference for the simpler model. As the noise level is reduced, however, the predictive power of the evidence increases, and the Bayes factor indicates a strong preference for the more complex model in the majority of cases (57 per cent for a noise reduction of 10, and 69 per cent for a noise reduction of 20). We note that the latter constraints could also easily be recast as a condition on stellar mass of galaxies within the 3D-*HST* survey at $z \sim 1$; roughly

speaking, the Bayes factor favours the more complex model with high confidence for galaxies with a stellar mass $= 2.5 \times 10^9 M_{\odot}$.

In Scenario 2, we observe that the Bayes factor remains a fair estimator of the preferred model, with no cases where the Bayes factor strongly favours for the simpler model, and a small fraction (12 per cent of the total) of failures, in which the simpler model is weakly favoured. Similar to scenario 1, at the 3D-*HST* noise level the Bayes factor is largely indecisive, but becomes more and more able to select the six-parameter model with confidence as the noise is reduced, with a strong preference for the more complex model for 27 per cent and 36 per cent of galaxies in the case of a 10× and 20× noise reduction, respectively.

In Scenario 3, the failures tend to disappear more quickly as S/N increases, but otherwise we observe a similar trend as in Scenario 2 in the Bayes factor’s ability to select the more complex six-parameter model as the noise is reduced. At the 3D-*HST* noise level, the Bayes factor fails only 14 per cent of the time, in which the simpler model is weakly favoured. At the 10× and 20× noise reduction, the Bayes factor strongly prefers the more complex model 22 per cent and 35 per cent, respectively.

Our results indicate that even in the presence of modelling systematics, the Bayes factor is a fairly reliable indicator of model complexity, and is able to detect a secondary episode of star formation in the vast majority of cases, especially at high S/N.

5.4 Investigating failures

We now attempt to investigate of the ‘failures’ of the Bayes factor (i.e. when the Bayes factor incorrectly prefers the simpler model, or $\ln B_{01} < -1$) as a tool to discern model complexity. We would like to understand, if possible, what factors may contribute to skew the evidence toward the simpler (incorrect) model, resulting in the simpler model being strongly or moderately preferred. We consider two possible explanations: the S/N ratio of each individual spectrum, which we expect to be important given the improvement seen at the population level with increasing S/N, and the age separation between




Figure 8. The top bar chart (Scenario 1) shows the Bayes factor for data simulated using a five-parameter model (delayed τ component + recent burst) and fitted with a four-parameter model (delayed τ component), as described in the text. The middle bar chart (Scenario 2) shows the Bayes factor for data simulated using a six-parameter model (delayed τ component + burst of varying mass/age) and fitted with a five-parameter model (delayed τ component + recent burst). The bottom bar chart (Scenario 3) shows the Bayes factor for data simulated using a six-parameter model (delayed τ component + burst of varying mass/age) and fitted with a four-parameter model (delayed τ component). The horizontal axis shows the distribution of the Bayes Factor for different S/N ratios, grouped according to the criteria of the Jeffrey Scale. From left to right, the ‘columns’ indicate that the BF moderately prefers the simpler model, is indecisive, moderately prefers the more complex model, decisively prefers the more complex model. In all scenarios, the Bayes Factor ‘fails’ (prefers the simpler model) in a minority of cases, especially at high S/N.

the main and the secondary population. The rationale for the latter is that if the two population have similar ages, they effectively act as a single population, and in this sense the single-population model is still – in practical if not formal terms – the correct choice.

In Fig. 5, we plot the distribution of cases in which the Bayes factor incorrectly prefers the simpler model ($\ln B_{01} < -1$) and those in which the Bayes factor correctly prefers the more complex model ($\ln B_{01} > 1$), for each scenarios and noise level, as a function of the total S/N of each SED.

Scenarios 1 and 3 both show a significantly higher median total S/N for the succeeding cases in comparison to the failing cases, indicating that S/N is probably the main driver of the expected performance of the method. This is encouraging because Scenarios 1 and 3 correspond to a practical situation in which we are trying to detect a secondary population (recent, in Scenario 1, and of any age, for Scenario 2) by comparing a two-population fit with a single-

population fit. From this plot, as well as the results presented in Section 5.3, we can conclude that for sufficiently deep photometry, the Bayes factor is a reliable indicator of model complexity in the vast majority of cases.

The situation is not the same for Scenario 2, in which we attempt to fit two-population SEDs by fixing the age of the secondary population, and varying its mass. In this case, we see that the SED’s total S/N is less correlated with the distribution of failing cases. This happens possibly because of increased degeneracies between the simpler (five-parameter) model and the more complex (six-parameter) model for this scenario, which make it harder to isolate one single discriminating factor. We note that this Scenario is the least akin to a physical situation in which we would want to use the Bayes factor as a discriminating tool for model complexity.

A similar trend can be seen when we plot the same histograms of failing and succeeding cases as a function of age separation between




Figure 9. The top bar chart (Scenario 1) shows the Bayes factor for data simulated using a five-parameter model (delayed τ component + recent burst) and fitted with a four-parameter model (delayed τ component). The bottom bar chart (Scenario 4) shows the Bayes factor for the same parameter space, but when uniform priors in A_v , $\text{Age}_{\text{delayed } \tau}$, and τ are used, instead of uniform priors in the logarithm of the same variables, like in Scenario 1. Notation is the same as Fig. 8. Despite the strong decay in the quality of parameter estimates observed in Scenario 4 (see Table 6), the Bayes factor remains reliable as a model complexity selection tool.

the main (delayed τ SFH) component and the secondary (burst) component. Age separation here is defined as the time between the onset of star formation in the delayed τ SFH and the burst component. In Fig. 6, Scenarios 1 and 3 show a significant difference in the median age separation for the failing versus succeeding cases at each noise level; as expected, the Bayes factor is less able to pick the correct model if the age separation between the two components is small because they effectively ‘become one’. Conversely, in Scenario 2 there is little difference in the median age separation for these two cases.

5.5 Design space and impact of priors

One often neglected issue in Bayesian parameter estimation is the role that the choice of priors plays in determining the final results. We are interested in the recovery of physical parameters, as well as in how the choice of priors correlates with the reliability of the Bayes factor as a model selection tool.

We show a preliminary example (Scenario 4). We begin by comparing two five-parameter models: one analogue to the one described in Scenario 1, and one in which we employ a uniform prior in the age, τ parameter, and dust parameter A_v of the main stellar population. We note that since the multivariate input parameter distributions are derived directly from the data (as opposed to being generated from a prior distribution), this test should be effective at determining which priors are preferable at a general level.

Fig. 7 and Table 6 show that the set of priors from Scenario 1 are highly preferable to the ones used here. The decay in the quality of the fits is evident, with the emergence of significant biases in some of the affected parameters, such as A_v and τ , which also present very large outlier fractions, of the order of 20–30 and 50 per cent, respectively.

Our recommendation from this simple test is that uniform priors in $\log(A_v)$ and $\log(\tau)$ should be used.

In Fig. 9, we show the Bayes factor from the comparison of a five-parameter model, described in Scenario 1 above, and a five-parameter model in which the priors are uniform in age, τ , and A_v , as opposed as in \log_{10} .

The main result from this test is that at the 3D-*HST* noise level, priors play an important role in determining the Bayes factor, and the results obtained by using the uniform priors from Scenario 4 are not trustworthy in a higher number of cases. On the other hand, as the noise is reduced, the contribution of the data to the marginal likelihood increases, and the Bayes factor remains a reliable estimator of model complexity.

6 EVIDENCE VERSUS SAVAGE-DICKEY DENSITY RATIO

Our last novel result comes from the comparison of the Bayes factor evaluated through the SDDR versus through the evidence. As mentioned above, the SDDR can be easily obtained by ordinary MCMC sampling, as opposed to the evidence calculation that requires expensive ad hoc methods such as nested sampling.

Table 9 provides a measure of the linear correlation between the natural log of the Bayes factor derived from the SDDR and the natural log of the Bayes factor derived from nested sampling in Scenarios 1–2 at each noise level. Scenario 3 is not included because it involves comparison of a six- and four-parameter model, and the SDDR method is valid only with N versus $N - 1$ model scenarios. Scenario 4 is also not included because the poor posterior fits in Section 3.4 indicate that the evidence calculations are not reliable for model comparison. The first column (ρ_{total}) refers to the Pearson

Table 7. Parameter priors used for SED fitting in Scenarios 1, 2, and 3. We note that the prior used for the mass variables is effectively a uniform prior in $\log_{10} M/M_{\odot}$.

Parameter	Prior	Units	Sampling
A_v	(0.0001, 3.0)	Magnitudes	Uniform in \log_{10}
age	(0.05, age_{univ})	Gyr	Uniform in \log_{10}
burst age	(0.05, 0.90)	Gyr	Uniform
τ	(0.01, 10.0)	Gyr	Uniform in \log_{10}
\log_{10} , Mass	(6.5, 12.5)	M/M_{\odot}	Uniform
\log_{10} , Burst mass	(6.5, 12.5)	M/M_{\odot}	Uniform

Table 8. Bayes factors are used to compare the evidence of two competing models M_0 and M_1 . The table includes rule-of-thumb threshold values known as Jeffreys’ Scale, and the probability column is in reference to posterior probabilities of the model that is favoured.

$ \ln B_{01} $	Odds	Probability	Strength of evidence
<1.0	<3: 1	<0.750	Inconclusive
1.0	~3: 1	0.750	Weak evidence
2.5	~12: 1	0.923	Moderate evidence
5.0	~150: 1	0.993	Strong evidence

Table 9. ρ_{total} refers to the Pearson correlation coefficient derived from the Bayes factor calculated using either nested sampling or the SDDR, from the full set of 500 models. ρ_{sub} refers to the Pearson correlation coefficient calculated using only a subset of models that have at least 1% of their samples within 0.25×10^9 solar masses of θ_* . Numbers with * indicate that a single outlier was removed to calculate ρ_{sub} .

	ρ_{total}	ρ_{sub}	Models (ρ_{sub})	Median post samps (ρ_{sub})
Scenario 1	0.770	0.963	392	953
Noise levels	0.739	0.944	255	953
	0.779	0.952	186	1069
	0.757	0.932	139	1000
	0.761	0.928	472	1138
Scenario 2	0.838	0.889*	409	1109
Noise levels	0.717	0.892*	373	1267
	0.783	0.907	330	1321

correlation coefficient calculated before any requirement is made on the number of samples near θ_{star} , so it includes all 500 models for each scenario at each noise level. The second column (ρ_{sub}) lists the Pearson correlation coefficient calculated when considering on only models that have at least 1 per cent of their samples within 0.25×10^9 solar masses of θ_* . The third column lists the resulting number of models considered due to this requirement, and the fourth column lists the median number of posterior samples used to calculate ρ_{sub} .

These results show a high degree of correlation between the SDDR and nested sampling methods of Bayes factor computation, as long as the region of the posteriors where θ_* is evaluated are sufficiently well sampled. They suggest that with dense sampling of the tails of the posterior and when the other conditions for the applicability of the SDDR approximation hold – most notably nested models – the SDDR can be successfully used as a proxy for the full calculation stemming from the evidence.

7 CONCLUSIONS AND FUTURE WORK

In this paper, we have explored several simulated scenarios in order to establish whether Bayesian model comparison can be relied upon as a method to estimate the ‘true’ model complexity afforded by the data. Our main conclusions are the following:

(i) The Bayes factor is a promising tool to evaluate model complexity. In all the scenarios we have explored, the Bayes factor is able to select the correct model with high confidence in the majority of cases, particularly at high S/N, while it ‘fails’, i.e. selects the wrong model with high confidence, in only a few per cent of the cases. This is relevant in view of data from upcoming telescopes, for example, the Nancy Grace Roman Space Telescope (NRST, formerly *WFIRST*) and the James Webb Space Telescope (JWST). When combined, data from these two observatories will provide deep coverage over the entire range of wavelengths (with the exception of the bluest channels, but extending to mid-infrared) considered here. While obtaining a direct comparison is difficult without making reference to a specific observing proposal, it can generally be expected that data from JWST and NRST will be 2–3 magnitudes deeper than those currently available from *HST* and *Spitzer*, for deep fields of area 0.25 deg^2 or larger (A. Koekemoer, in proceedings of NRST conference, Space Telescope, 2020 October, and Koekemoer et al. 2019). A uniform improvement of 2.5 magnitude deeper data across the range of the spectrum would correspond to the ‘ $10\times$ noise reduction’ case we have described in this paper.

(ii) In the two scenarios that most closely resemble a comparison between a single component of star formation and a two-component model with an additional (recent or otherwise) burst, the Bayes factor performs increasingly well as a function of increasing S/N and age separation between the two components.

(iii) Priors play a large role in parameter estimation and Bayes factors estimation, especially at low S/N, as is to be expected. Our recommendation is that uniform priors in the logarithms of stellar age, A_v , and τ , are preferred to uniform priors in the same variables, although we also observe that the Bayes factor is less affected by incorrect priors than the estimates of physical parameters.

(iv) The SDDR can be a good proxy for the Bayes factor when the region of the posteriors where θ_* is evaluated is sufficiently well sampled. This enables the usage of Bayesian model comparison even without the computationally intensive need for nested sampling algorithms, broadening the range of possible scientific applications of this method.

Finally, we believe that our KDE method to extract realistic multivariate distributions for input parameters from data, keeping into account the covariances between different physical quantities, will be useful to the broader astrophysical community.

In a future paper, we aim to consider more carefully the sources of systematic uncertainty of the Bayes factor, and to run multiple

simulations on each scenario to account for the stochasticity of the evidence calculation in the MULTINEST nested sampling algorithm.

ACKNOWLEDGEMENTS

We are grateful to Adam Carnall, Kartheik Iyer, and Eric Gawiser for many useful discussions, as well as to the anonymous referee for many useful comments that helped us improve the paper in content and form. This work was partially supported by a PSC–CUNY Type A award (62101-0050) and a Google Cloud Platform Research for Education grant. VA thanks the Institute of Cosmic Sciences of the University of Barcelona for hospitality during the completion of this work.

DATA AVAILABILITY

This work is based on observations taken by the 3D-*HST* Treasury Program (GO 12177 and 12328) with the NASA/ESA *HST*, which is operated by the Association of Universities for Research in Astronomy, Inc., under NASA contract NAS5-26555. The data underlying this article are available in 3D-*HST* at http://monoceros.astro.yale.edu/RELEASE_V4.0/Photometry/GOODS-S/goodss_3dhst.v4.1.cats.tar.gz.

REFERENCES

- Acquaviva V., Gawiser E., Guaita L., 2011, *ApJ*, 737, 47
 Acquaviva V., Raichoor A., Gawiser E., 2015, *ApJ*, 804, 8
 Arnouts S., Cristiani S., Moscardini L., Matarrese S., Lucchin F., Fontana A., Giallongo E., 1999, *MNRAS*, 310, 540
 Behroozi P. S., Conroy C., Wechsler R. H., 2010, *ApJ*, 717, 379
 Bolzonella M., Miralles J.-M., Pelló R., 2000, *A&A*, 363, 476
 Buchner J., 2016, *Astrophysics Source Code Library*, record ascl:1606.005
 Calzetti D., Armus L., Bohlin R. C., Kinney A. L., Koornneef J., Storchi-Bergmann T., 2000, *ApJ*, 533, 682
 Carnall A. C., McLure R. J., Dunlop J. S., Davé R., 2018, *MNRAS*, 480, 4379
 Chevillard J., Charlot S., 2016, *MNRAS*, 462, 1415
 Conroy C., 2013, *ARA&A*, 51, 393
 Conroy C., Gunn J. E., White M., 2009, *ApJ*, 699, 486
 da Cunha E., Charlot S., Elbaz D., 2008, *MNRAS*, 388, 1595
 Dickey J. M., 1971, *Ann. Math. Statist.*, 42, 204
 Dries M., Trager S. C., Koopmans L. V. E., 2016, *MNRAS*, 463, 886
 Dries M., Trager S. C., Koopmans L. V. E., Popping G., Somerville R. S., 2018, *MNRAS*, 474, 3500
 Feroz F., Hobson M. P., 2008, *MNRAS*, 384, 449
 Feroz F., Hobson M. P., Bridges M., 2009, *MNRAS*, 398, 1601
 Gavazzi G., Bonfanti C., Sanvito G., Boselli A., Scodreggio M., 2002, *ApJ*, 576, 135
 Genel S. et al., 2014, *MNRAS*, 445, 175
 Grogin N. A. et al., 2011, *ApJS*, 197, 35
 Han Y., Han Z., 2012, *ApJ*, 749, 123
 Han Y., Han Z., 2014, *ApJS*, 215, 2
 Han Y., Han Z., 2018, *ApJS*, 240, 3
 Iyer K., Gawiser E., 2017, *ApJ*, 838, 127
 Iyer K. G., Gawiser E., Faber S. M., Ferguson H. C., Kartaltepe J., Koekemoer A. M., Pacifici C., Somerville R. S., 2019, *ApJ*, 879, 116
 Jeffreys H., 1998, *The Theory of Probability*. Oxford Univ. Press, Oxford
 Koekemoer A. M. et al., 2019, *Bull. Am. Astron. Soc.*, 51, 550
 Kriek M., van Dokkum P. G., Labbé I., Franx M., Illingworth G. D., Marchesini D., Quadri R. F., 2009, *ApJ*, 700, 221
 Kroupa P., Boily C. M., 2002, *MNRAS*, 336, 1188
 Lee S.-K., Ferguson H. C., Somerville R. S., Wiklind T., Giallisco M., 2010, *ApJ*, 725, 1644
 Leja J. et al., 2019, *ApJ*, 877, 140
 Leja J., Johnson B. D., Conroy C., van Dokkum P. G., Byler N., 2017, *ApJ*, 837, 170
 MacKay D., 2003, *Information Theory, Inference, and Learning Algorithms*. Cambridge Univ. Press, Cambridge
 McCracken H. J. et al., 2012, *A&A*, 544, A156
 McKee C. F., Ostriker E. C., 2007, *ARA&A*, 45, 565
 Mobasher B. et al., 2015, *ApJ*, 808, 101
 Mukherjee P., Parkinson D., Liddle A. R., 2006, *ApJ*, 638, L51
 Pacifici C. et al., 2015, *MNRAS*, 447, 786
 Pedregosa F. et al., 2011, *J. Mach. Learn. Res.*, 12, 2825
 Pillepich A. et al., 2018, *MNRAS*, 473, 4077
 Salmon B. et al., 2016, *ApJ*, 827, 20
 Scoville N. et al., 2007, *ApJS*, 172, 1
 Shaw J. R., Bridges M., Hobson M. P., 2007, *MNRAS*, 378, 1365
 Skelton R. E. et al., 2014, *ApJS*, 214, 24
 Skilling J., 2006, *Bayesian Anal.*, 1, 833
 Skrutskie M. F. et al., 2006, *AJ*, 131, 1163
 Somerville R. S., Davé R., 2015, *ARA&A*, 53, 51
 Speagle J. S., 2020, *MNRAS*, 493, 3132
 Stone C. J., Koo C.-Y., 1986, *Contemp. Math*, 59, 15
 Stone C. J., Hansen M. H., Kooperberg C., Truong Y. K., 1997, *Ann. Stat.*, 25, 1371
 Tran T. N., Wehrens R., Buydens L. M., 2006, *Comput. Stat. Data Anal.*, 51, 513
 Trotta R., 2007, *MNRAS*, 378, 72
 Trotta R., 2008, *Contemp. Phys.*, 49, 71
 Trotta R., 2017, preprint ([arXiv:1701.01467](https://arxiv.org/abs/1701.01467))
 Verdinelli I., Wasserman L., 1995, *J. Am. Stat. Assoc.*, 90, 614
 Wagenmakers E.-J., Lodewyckx T., Kuriyal H., Grasman R., 2010, *Cogn. Psychol.*, 60, 158
 Walcher J., Groves B., Budavári T., Dale D., 2011, *Ap&SS*, 331, 1
 Wetzels R., Raaijmakers J. G., Jakob E., Wagenmakers E.-J., 2009, *Psychonomic Bull. Rev.*, 16, 752
 York D. G. et al., 2000, *AJ*, 120, 1579

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