



# Bayesian Inference of the Rates of Surface Reactions in Icy Mantles

J. Holdship , N. Jeffrey , A. Makrymallis, S. Viti , and J. Yates 

Department of Physics and Astronomy, University College London, Gower Street, London WC1E 6BT, UK; [jrh@star.ucl.ac.uk](mailto:jrh@star.ucl.ac.uk)

Received 2015 November 5; revised 2018 September 13; accepted 2018 September 13; published 2018 October 19

## Abstract

Grain surface chemistry and its treatment in gas-grain chemical models is an area of large uncertainty. While laboratory experiments are making progress, there is still much that is unknown about grain surface chemistry. Further, the results and parameters produced by experiments are often not easily translated to the rate equation approach most commonly used in astrochemical modeling. It is possible that statistical methods can reduce the uncertainty in grain surface chemical networks. In this work, a simple model of grain surface chemistry in a molecular cloud is developed and a Bayesian inference of the reactions rates is performed through Markov Chain Monte Carlo sampling. Using observational data of the solid state abundances of major chemical species in molecular clouds, the posterior distributions for the rates of seven reactions producing CO, CO<sub>2</sub>, CH<sub>3</sub>OH, and H<sub>2</sub>O are calculated in a form that is suitable for rate equation models. This represents a vital first step in the development of a method to infer reaction rates from observations of chemical abundances in astrophysical environments.

*Key words:* astrochemistry – dust, extinction – ISM: molecules – methods: numerical – methods: statistical

## 1. Introduction

Dust grain chemistry plays an important role in the physical processes happening deep inside dark molecular clouds during star formation (Draine 2003; Williams & Cecchi-Pestellini 2015). These dust grains are vital to every part of the star formation process and ultimately contribute to the basic matter from which icy planetesimals are formed (van Dishoeck 2004). It is in fact evident that molecules such as water and methanol in dust grain ice mantles are primarily formed through solid state chemistry rather than accreted directly from the gas phase (Parise et al. 2005; Ceccarelli et al. 2007). In recent years, even more complex molecules have been observed in both prestellar cores and star-forming regions (see reviews by Herbst & Van Dishoeck 2009; Caselli & Ceccarelli 2012), some of which cannot currently be explained by pure gas phase chemistry. Therefore, chemical reactions leading to simple as well as complex molecules must occur on the surface of icy dust grains.

Experimentally, it has been known since the work of Hagen et al. (1979) and Pirronello et al. (1982) that grains can be chemical nanofactories on which surface reactions, UV photons, and cosmic-ray radiation can synthesize complex molecules and even prebiotic species, starting from simple atoms or molecules such as H, C, O, N, CO. Therefore, understanding the surface chemistry that takes place on dust grains is key to understanding not only the origins of stars, but also how rocky and gaseous planets are formed.

Initially, surface reaction networks in chemical models were based on chemical intuition and gas phase chemistry analogs. However, over the past two decades, laboratory astrochemists have been using experimental techniques to test and evaluate surface reactions. As a result, the efficiencies of reaction routes are being properly explored and important information on how molecules form on grain surfaces is being revealed (see Williams & Cecchi-Pestellini 2015 for a review). The first experimental work on the dust surfaces studied the formation of molecular hydrogen (Pirronello et al. 1997). Several more experiments followed studying either the formation of more complex molecules (e.g., Watanabe et al. 2005; Ioppolo et al. 2009; Minissale et al. 2015)

or the ice morphology and ice mantle mechanisms (e.g., Fraser et al. 2004; Collings & McCoustra 2006). Surface reactions can be experimentally investigated within a constrained range of laboratory conditions. Typically, these conditions include different atomic fluxes, ice temperatures, ice morphologies, and mixture ratios, as well different energetic processes. The aims of the experiments are to investigate surface molecule formation, desorption, and diffusion. However, experimental data for interstellar ices are limited, the main reason being that the experimentation process is neither simple nor fast. In order to make the best use of experimental resources, the chemical data that models require need to be prioritized according to what will have the most impact.

Bayesian methods are widely used in astronomy as a means of deriving posterior probability distributions (PPDs) for model parameters from observations (e.g., Palau et al. 2014; Schmalzl et al. 2014; Bevan 2018). It is the de facto standard in the field of cosmology, but is becoming more and more widely used in other areas of astronomical research (e.g., Lomax et al. 2016; Testi et al. 2016). In astrochemistry, Bayesian inference has been used to derive parameters such as the gas density and cosmic-ray ionization rates within a dark molecular cloud from observations of species in the gas and ices using chemical models (Makrymallis & Viti 2014).

In this work, we investigate the chemistry itself, studying the rates of reactions on the dust grain surfaces in an attempt to infer their rates and provide a list of reactions for which an accurate rate is particularly important. This is the first such work in an astrochemical context but Bayesian methods have been used to determine rate coefficients for combustion chemistry on Earth (Prager et al. 2013). This work represents a necessary first step in which we determine whether we can use a reduced chemical model and very simple observational constraints to learn more about the parameters in a grain surface chemical model.

The grain chemistry model used in this work is described in Section 2. The inference process including the choice of Markov Chain Monte Carlo (MCMC) sampler is presented in Section 3. The results of our analysis are presented in Section 4 along with an additional discussion in Section 5. Finally, our conclusions are discussed in Section 6.

**Table 1**  
Species Included in the Chemical Model

Species
CH <sub>3</sub> OH, CO, CO <sub>2</sub> , CS, CS <sub>2</sub> , H, H <sub>2</sub> CO, H <sub>2</sub> CS, H <sub>2</sub> O, H <sub>2</sub> S, H <sub>2</sub> S <sub>2</sub> , HCO, HCS, HOCS, HS, HSO, O, OCS, OH, S, SO, SO <sub>2</sub>

## 2. The Chemical Model

A simple chemical model was developed that considers only the solid state chemistry in the ice mantles of dust grains in a dark molecular cloud. The simplified model is a time-dependent single-point model that generates a time series of solid phase molecular abundances as a function of the physical conditions of the molecular cloud and the chemical parameters of the defined chemical network. The chemical network consists of 23 species and 24 surface reactions that are listed in Tables 1 and 2, respectively.

To model the surface chemistry of a dark cloud the abundance of each solid species is derived by solving rate equations for grain surface chemistry. The formation and destruction mechanisms for a species  $i$  are given by the following kinetic equation:

$$\frac{dn_i}{dt} = \sum_{l,m} k_{lm}^i n_l n_m - n_i \sum_{i=r} k_r n_r - k_i^{\text{des}} n_i + k_i^{\text{ads}} n_{i,\text{gas}}, \quad (1)$$

where  $k_{lm}^i$  is the reaction rate of all the reactions between species  $l$  and  $m$  that produce  $i$ ,  $n_i$  is the concentration of species  $i$  (with the subscript gas indicating the concentration of the species in the gas phase),  $k_r$  represents the reaction rates of all the reactions where species  $i$  participates as reactant, while  $k_i^{\text{des}}$  and  $k_i^{\text{ads}}$  are the desorption and adsorption rates, respectively.

The reactions in Table 2 are mainly hydrogenation reactions of common gas phase species and reactions between species that are likely to be abundant on the grains. Where possible, reactions that have been found to be efficient, or even dominant, routes to forming a species have been chosen. For example, the hydrogenation of CO to form CH<sub>3</sub>OH is well studied (Fuchs et al. 2009; Chuang et al. 2016), and so reactions 21–24 are the only considered route to form CH<sub>3</sub>OH. Similarly, the formation of CO<sub>2</sub> via reaction 3 is known to be efficient (Ioppolo et al. 2011), and other routes suffer from large energetic barriers. Note that in cases where H or H<sub>2</sub> is a product of a reaction such as in reaction 11, it is ignored and the total H abundance is not conserved. This is for simplicity and the lost H represents too small a fraction of the H abundance to affect the model.

There is no gas phase chemistry in the model, and so the freezeout of species from the gas phase must be parameterized. The adsorption rate is assumed to be zero for all but the following six species: CO, CS, O, H, OH, and S. To derive the adsorption rate of these species, the gas-grain chemical code UCLCHEM (Holdship et al. 2017) was utilized. UCLCHEM was run with a network of 220 species with gas phase reactions from UMIST12 (McElroy et al. 2013), freezeout of gas phase species, and the nonthermal desorption of grain surface species. A single-point model of this full gas-grain chemistry was run in which the gas increased in density under freefall from 10<sup>2</sup> cm<sup>-3</sup> to 2 × 10<sup>4</sup> cm<sup>-3</sup>, which is appropriate for a dark molecular

**Table 2**  
Reaction Network Used in the Chemical Model

No.	Reactions				
1.	O	+	H	→	OH
2.	OH	+	H	→	H <sub>2</sub> O
3.	CO	+	OH	→	CO <sub>2</sub>
4.	S	+	H	→	HS
5.	HS	+	H	→	H <sub>2</sub> S
6.	H <sub>2</sub> S	+	S	→	H <sub>2</sub> S <sub>2</sub>
7.	CS	+	H	→	HCS
8.	HCS	+	H	→	H <sub>2</sub> CS
9.	CO	+	S	→	OCS
10.	OCS	+	H	→	HOCS
11.	H <sub>2</sub> S	+	CO	→	OCS
12.	H <sub>2</sub> S	+	H <sub>2</sub> S	→	H <sub>2</sub> S <sub>2</sub>
13.	H <sub>2</sub> S <sub>2</sub>	+	CO	→	CS <sub>2</sub> + O
14.	H <sub>2</sub> S	+	O	→	SO <sub>2</sub>
15.	CS <sub>2</sub>	+	O	→	OCS + S
16.	CO	+	HS	→	OCS
17.	S	+	O	→	SO
18.	SO	+	O	→	SO <sub>2</sub>
19.	SO	+	H	→	HSO
20.	HSO	+	H	→	SO
21.	CO	+	H	→	HCO
22.	HCO	+	H	→	H <sub>2</sub> CO
23.	H <sub>2</sub> CO	+	H	→	H <sub>3</sub> CO
24.	H <sub>3</sub> CO	+	H	→	CH <sub>3</sub> OH

**Note.** The rates of these reactions are the parameters of interest in this work.

cloud. The chemistry progresses over 10 Myr at 10 K and the freezeout rates for the six species above were extracted from this model.

The freezeout rates from UCLCHEM were inserted as source terms in the ODEs for those species in the simple grain surface model. The grain surface models starts with an abundance of zero for all species, representing bare grains. The model then progresses for 10 Myr considering only the 24 reactions in Table 2, the freezeout rates, and the nonthermal desorption of each species. In this way, the grain surfaces in a dark molecular cloud are effectively modeled while the computation time is low as the gas phase treatment is reduced to the incoming (freezeout) and outgoing (desorption) flux of molecules. Note that the cloud age is arbitrary, the model reaches the molecular cloud density at 6 Myr, and the chemistry is then allowed to progress until a total age of 10 Myr. The exact choice of final time has only a small affect on the abundances in the model.

Whether one parameterizes the rate of surface reactions in a similar way to the Kooji–Arrhenius equation used for gas phase chemistry (Occhiogrosso et al. 2012) or considers the diffusion and reaction of species across the ice surface (Hasegawa et al. 1992; Chang et al. 2007), the rate of a reaction is constant for a given temperature and dust composition. Therefore,  $k$  in this model is treated as a constant rate of reaction in units of cm<sup>3</sup> s<sup>-1</sup> as the temperature in the model is constant at 10 K. This reduces the number of parameters in the model and reflects the available data, i.e., ice phase abundances in quiescent, approximately isothermal clouds.

The result of these approximations and modifications is a model of the grain surface chemistry under the conditions of a dark molecular cloud at a constant temperature of 10 K. The freezeout rates and gas density are specific to a dark cloud, and

so the model is not of applicable to arbitrary ices. This model has a run time that is approximately 1000 times shorter than an equivalent run of UCLCHEM. This reduction in run time is vital due to the number of model runs required for an MCMC inference procedure.

### 3. Bayesian Inference

#### 3.1. Inference Procedure

The aim of this work is to obtain information about the set of reaction rates  $\mathbf{k} = (k_1, k_2, \dots, k_{23})$  of the surface chemical network, where  $k_j$  is the reaction rate of reaction  $j$ . For a given set of rates, the model produces simulated molecular abundances  $\mathcal{Y} = (\mathcal{Y}_1, \mathcal{Y}_2, \dots, \mathcal{Y}_{22})$ , where  $\mathcal{Y}_i$  is the abundance of species  $i$ . These quantities are related through the chemical code  $\mathcal{C}(\cdot)$  so that  $\mathcal{Y} = \mathcal{C}(\mathbf{k})$ .

For any set of simulated abundances, the probability of the corresponding parameter values can be evaluated through the use of Bayes's rule:

$$P(\mathbf{k}|\mathbf{d}) = \frac{P(\mathbf{d}|\mathbf{k})P(\mathbf{k})}{P(\mathbf{d})} \propto P(\mathbf{d}|\mathbf{k})P(\mathbf{k}), \quad (2)$$

where  $\mathbf{d}$  is the data, representing a set of observational constraints on  $\mathcal{Y}$ .  $P(\mathbf{k}|\mathbf{d})$  is the PPD of  $\mathbf{k}$ , which expresses the level of certainty about the reaction rates after considering the data and any prior information. The denominator is known as the Bayesian evidence, but for the purposes of this study can simply be treated as a normalization factor. The prior probability distribution ( $P(\mathbf{k})$ ) adopted for the reaction rates is a logarithmically uniform distribution that is nonzero when the reaction rates are between  $10^{-5}$  and  $10^{-30}$  and zero elsewhere. The limits of the prior distributions represent a larger range than that of rates typical of gas phase reactions. This is a choice to reflect the exploratory nature of the work and is expected to cover all likely rates regardless of the nature of surface reactions.

The likelihood function  $P(\mathbf{d}|\mathbf{k})$  must give the likelihood of having obtained the data given the assumed set of rates. Here, the likelihood encodes measurement noise and is given as

$$P(\mathbf{d}|\mathbf{k}) = \exp\left(-\frac{1}{2} \sum_i \frac{(d_i - \mathcal{Y}_i)^2}{\sigma_i^2}\right), \quad (3)$$

where  $\mathcal{Y}_i$  are the model abundances of each species for which there are data and  $\sigma_i$  is the Gaussian uncertainty of each observed fractional abundance.

#### 3.2. Implementation and Data

In order to constrain the reaction rates, data are required in the form of species abundances in the ices. The solid state fractional abundances of species in quiescent gas illuminated by background stars were taken from a comprehensive recent review (Boogert et al. 2015, Table 2). Of the species in the grain surface model, this review gives constraints on the abundance of H<sub>2</sub>O, CO, CO<sub>2</sub>, and CH<sub>3</sub>OH. The value of the median fractional abundance is provided for each along with the upper and lower quartile values. To formulate the likelihood, it is assumed these values describe a Gaussian distribution for the abundance of each species. It is also assumed that the uncertainties on the abundances are independent, which is likely given that the abundances and statistics presented by Boogert et al. (2015) are combinations of

**Table 3**

(Upper Section) Adopted Abundances and Uncertainties of Species Observed in the Ices Used as Data in the Parameter Inference and (Lower Section) Upper Limits of the Fractional Abundance for Other Species that Are Used to Further Constrain the Reactions Rates

Species	Abundance
H <sub>2</sub> O	$(4.0 \pm 1.3) \times 10^{-5}$
CO	$(1.2 \pm 0.8) \times 10^{-5}$
CO <sub>2</sub>	$(1.3 \pm 0.7) \times 10^{-5}$
CH <sub>3</sub> OH	$(5.2 \pm 2.4) \times 10^{-6}$
H <sub>2</sub> S	$<1.6 \times 10^{-6}$
H <sub>2</sub> CO	$<3.0 \times 10^{-5}$
OCS	$<1.2 \times 10^{-7}$
SO <sub>2</sub>	$<4.0 \times 10^{-6}$

**Note.** All values are adapted from Boogert et al. (2015) as discussed in the text.

different data sets for each species. The median value is taken as the mean and the upper and lower quartile values can then be assumed to be  $0.68\sigma$  from that mean. The resulting abundances and uncertainties are listed in the upper half of Table 3. Due to the low number of observations, these distributions are not perfect representations of the data as the quartile values are not precisely symmetric about the median.

In order to evaluate the posterior distribution function for all values of  $\mathbf{k}$ , a sampler must be used. The emcee python package (Foreman-Mackey et al. 2013) was chosen for this purpose. This is an implementation of the affine-invariant Monte Carlo sampler proposed by Goodman & Weare (2010). Rates were sampled by 128 “walkers,” each producing chains of  $\sim 10^6$  samples where the frequency of the appearance of a particular rate value in the chain is proportional to its likelihood. These walkers start from random positions in rate space (i.e., all 24 rates have a random value from  $10^{-30}$  to  $10^{-5} \text{ cm}^3 \text{ s}^{-1}$ ). The sampling took approximately 100 hr using a single node on the DiRAC CSD3 platform's Skylake-Peta4 system utilizing emcee's built-in MPI tools. In the Appendix, heuristics are presented that demonstrated the chains were likely to have converged.

#### 3.3. Upper Limits

The abundances of H<sub>2</sub>O, CO, CO<sub>2</sub>, and CH<sub>3</sub>OH are the only strong constraints on the abundances of the species in this model. The reaction rates that are acquired as a result of performing a Bayesian parameter inference procedure with these abundances are presented and discussed in Section 4.1. However, weaker constraints do exist for other species. Boogert et al. (2015) provide upper limits on the abundances of OCS in dark clouds as well as upper limits on H<sub>2</sub>CO, SO<sub>2</sub>, and H<sub>2</sub>S in other objects; the upper limits used are given in Table 3. Therefore, a second parameter inference procedure was performed that was identical to the first, except that the likelihood function was modified to take into account the upper limits. In order to be conservative, when deriving upper limits on species that have not been detected toward background stars, the upper limits toward YSOs were increased by an order of magnitude to allow for larger abundances in molecular clouds. This is the case for the upper limits on H<sub>2</sub>CO and SO<sub>2</sub>.

To account for these upper limits, modifications were made to Equation (3). When including the upper limits, the likelihood

of a model was calculated as

$$P(\mathbf{d}|\mathbf{k}) = \exp\left(-\frac{1}{2}\delta_i \sum_i \frac{(d_i - \mathcal{Y}_i)^2}{\sigma_i^2}\right) (1 - S(C_i))^{1-\delta_i}, \quad (4)$$

where  $\delta_i$  is 1 for observed species and 0 for species with upper limits.  $C$  is the upper limit of a species' abundance, and  $S(C_i)$  is the survival function. This modification to the likelihood is standard for left-censored data, i.e., ones where a detection limit provides only an upper limit on a quantity (Klein & Moeschberger 2003). The survival function for a Gaussian distribution is used:

$$S(C_i) = 1 - \frac{1}{2} \left[ 1 + \operatorname{erf}\left(\frac{C_i - \mathcal{Y}_i}{\sigma_i}\right) \right], \quad (5)$$

where  $\operatorname{erf}()$  is the error function.  $\sigma_i$  is assumed to be one-third of the value of the upper limit.

Equation (4) is equivalent to Equation (3) for detected species. However, for upper limits it takes the value of  $1 - S$ . For model abundances much less than the upper limit, this likelihood is equal to 1, and for model abundances much larger than the limit, it takes a value of zero. Thus, while model abundances close to the upper limit are accepted, models with much larger predicted abundances have a likelihood of zero.

### 3.4. Testing the Method

In order to ascertain whether this method would be able to predict reactions rates from measured abundances in the case that the model was an accurate representation of reality, it was tested using abundances obtained from the model itself. First, random rates for all 24 reactions were generated, and the model was run using those rates. The abundances of the four species in Table 3 for which observed abundances are available were then stored. ‘‘Noisy’’ abundances were then generated by drawing randomly from a Gaussian distribution with a mean value of the model abundances and a standard deviation set to a 50% error. This produced four ‘‘observations’’ obtained from the model with known rates. The 50% error was chosen as it is the approximate fractional error of the real observations. The MCMC procedure was then performed to see whether the known reaction rates could be obtained.

It was found that the majority of rates could not be recovered. However, in tests where the rates for reactions that produced  $\text{H}_2\text{O}$ ,  $\text{CO}$ ,  $\text{CO}_2$ , and  $\text{CH}_3\text{OH}$  were high enough to produce observable abundances, the rates of those reactions were recovered. That is to say that, intrinsically, this method appears to only be able to give information on the rates of reactions that form the species for which observational constraints are available.

## 4. Results

### 4.1. Gaussian Abundance Constraints

The results are presented in the form of marginalized PPDs for the reaction rate coefficients. The density of each marginalized PPD reveals the areas where the corresponding reaction rate is more probable based on the imposed constraints. In Figure 1, the marginalized PPDs of selected reactions are plotted. These show a large probability density only for a much smaller range of values than the prior. As expected from the tests in Section 3.4, these are generally reactions that form the

constrained species. The PPDs are shown as histograms and Gaussian kernel density estimates, using the full MCMC chains from all 128 walkers. It is believed that these chains have converged and the relevant tests are discussed in the Appendix.

The PPD of reactions 1, 2, 3, and 21 are well constrained and involve species directly constrained by observation. Reaction 1 provides OH required to form  $\text{H}_2\text{O}$  and  $\text{CO}_2$  through reactions 2 and 3. Those are in turn constrained by the observed abundances of  $\text{CO}_2$  and  $\text{H}_2\text{O}$  and mutual competition for OH. Reaction 21 uses up CO, and so it is expected there would be an upper limit due to the observed abundance of CO and a lower limit due to competition with reaction 3. Reactions 22–24 form  $\text{CH}_3\text{OH}$  from HCO. The competition between reactions and the correlation between the rates of reactions 21–24 are explored further in Section 5.2.

The other PPDs are broadly similar to the prior distributions, and the implications of this should be stated. Essentially, the reactions where the rates have uniform probability distributions are reactions that do not impact the likelihood of the model. It should be noted, however, that changes in the abundance of species not included in the likelihood calculation are possible. PPDs that are similar to the priors indicate that when modeling only  $\text{H}_2\text{O}$ ,  $\text{CO}$ ,  $\text{CO}_2$ , and  $\text{CH}_3\text{OH}$ , the rates of those reactions are unimportant.

### 4.2. Inclusion of Upper Limits

The PPDs of each reaction are largely unchanged when upper limits are included, indicating that the upper limits may in fact be too conservative. The only major change is that the rate of reaction 10 takes a minimum value of  $10^{-17} \text{ cm}^3 \text{ s}^{-1}$ . This is required for a models to produce a lower OCS abundance than the upper limit.

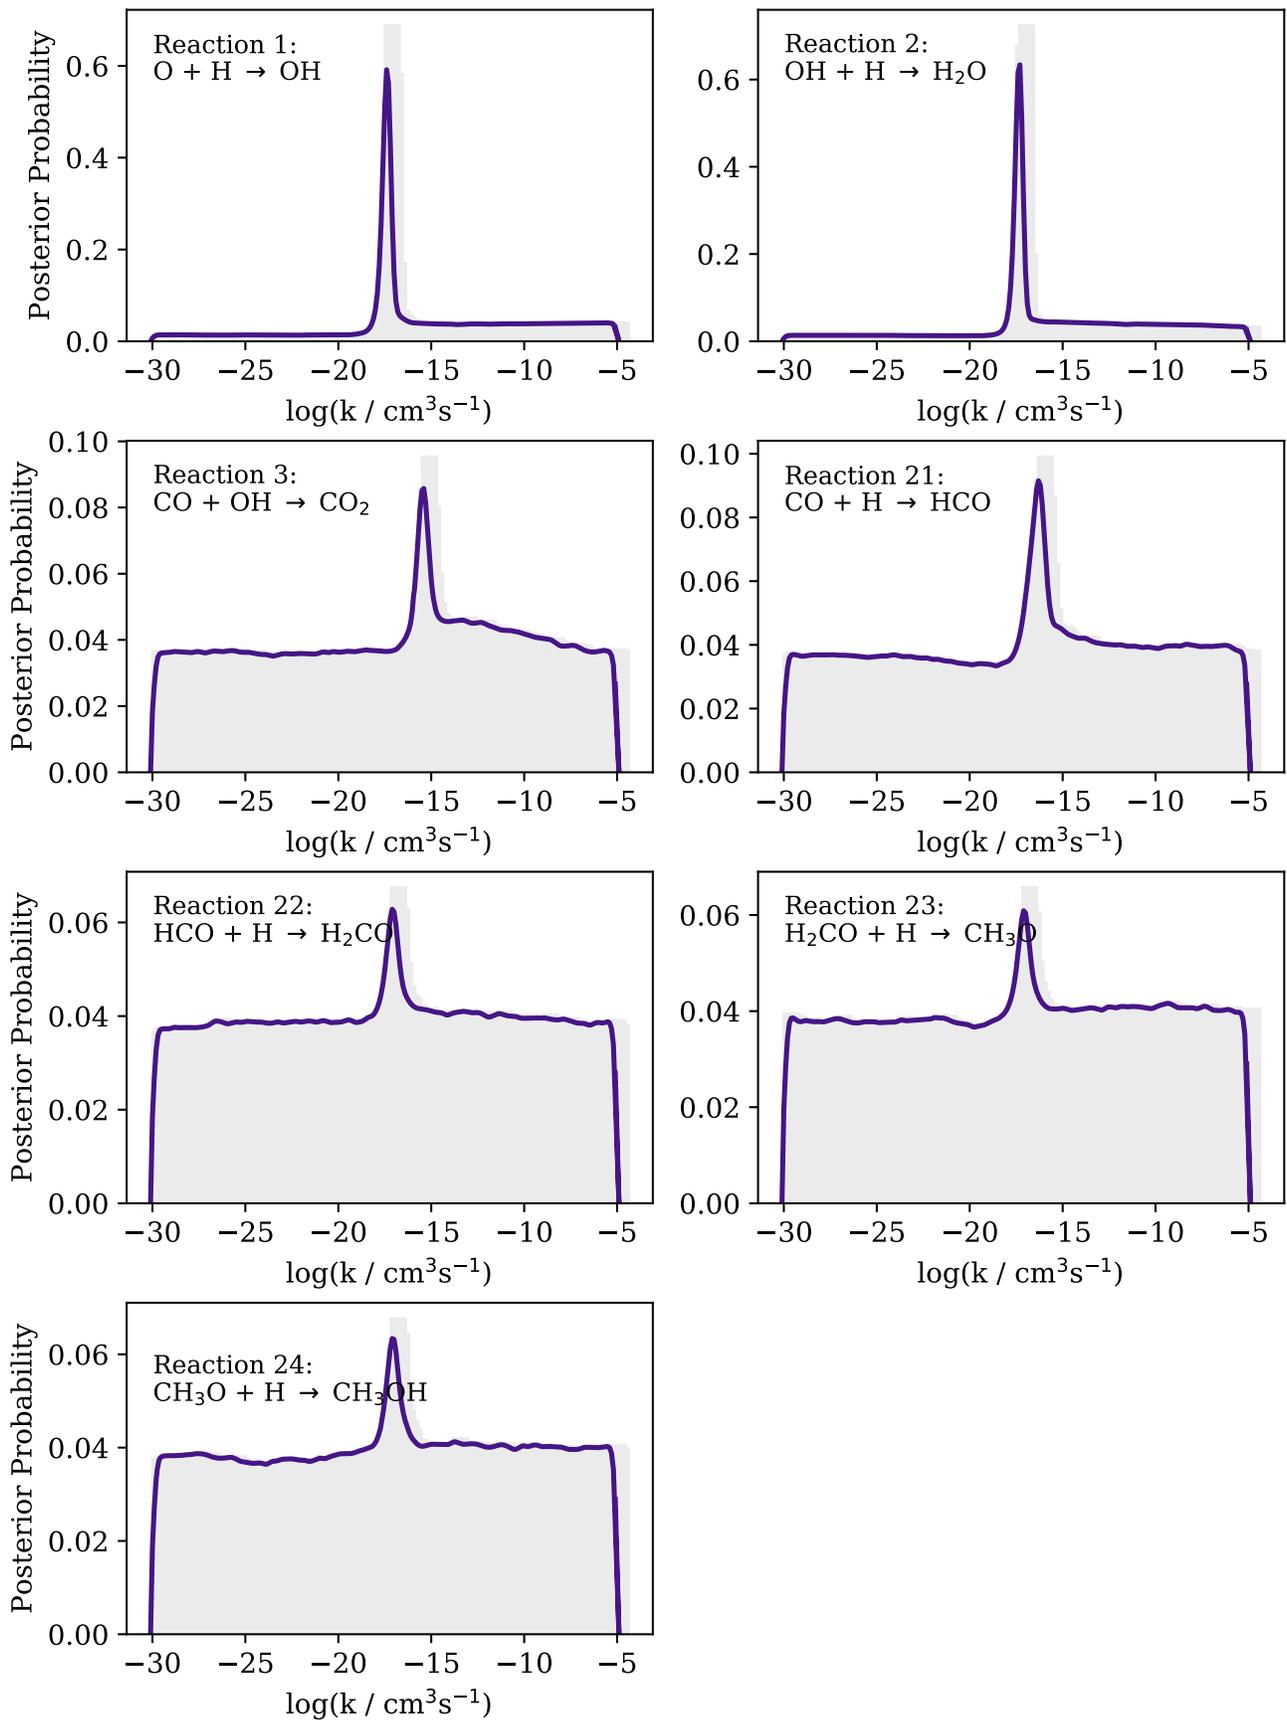
One may expect the upper limit on  $\text{H}_2\text{CO}$  to improve the level of certainty in the reaction series 21–24. However, using the conservative value in Table 3, no change is seen in the posteriors. If the value for YSOs is utilized instead, only a small change is observed. In that case, the peaks that were apparent in the marginalized posterior of each reaction in Figure 1 become more pronounced, such that the majority of the probability density lies within them. Thus, the value of  $\text{H}_2\text{CO}$  would represent an important constraint if an appropriate value for molecular clouds could be obtained.

The most likely rates for the well-constrained reactions from this MCMC analysis are presented in Table 4. These most likely rates are the same whether they are taken from the MCMC chains produced using observed abundances only or from the chains that included upper limits. It is clear that the available upper limits are not sufficiently constraining to improve the parameter inference.

## 5. Discussion

### 5.1. Model Abundances

The results of the MCMC run give the marginalized posterior distributions of the rates of each reaction. In order to understand how well these rates reproduce the observed abundances, the model must be run with rates drawn from the probability distributions. This will allow the uncertainty in the model that arises from the uncertain rates to be quantified. The model was run 1000 times, with the rates of the reactions randomly sampled from the marginalized posterior distributions derived from the upper limit MCMC procedure. By



**Figure 1.** Marginalized posterior probability distributions of the rates of the reactions with posterior distributions that are well constrained. The y-axis scaled such that the total probability density contained by the histogram is 1. All other reactions have posterior distributions that are approximately flat and similar to the prior.

**Table 4**  
Most Likely Values for the Rates of Well-constrained Reactions

Reaction	Rate ( $\text{cm}^3 \text{s}^{-1}$ )	65% Probability Range ( $\text{cm}^3 \text{s}^{-1}$ )
1	$4.0 \times 10^{-18}$	$1.0 \times 10^{-18}$ – $3.2 \times 10^{-10}$
2	$5.0 \times 10^{-18}$	$1.5 \times 10^{-18}$ – $1.5 \times 10^{-10}$
3	$4.0 \times 10^{-16}$	$6.8 \times 10^{-26}$ – $2.2 \times 10^{-10}$
21	$5.0 \times 10^{-17}$	$6.6 \times 10^{-26}$ – $3.2 \times 10^{-10}$
22	$7.9 \times 10^{-18}$	$4.6 \times 10^{-26}$ – $3.2 \times 10^{-10}$
23	$7.9 \times 10^{-18}$	$4.6 \times 10^{-26}$ – $3.2 \times 10^{-10}$
24	$7.9 \times 10^{-18}$	$4.6 \times 10^{-26}$ – $4.6 \times 10^{-26}$

**Note.** The intervals containing 65% of the probability density of the marginalized posteriors are also noted.

plotting the average abundance of each species and the standard deviation of those abundances, the uncertainty in the model can be evaluated.

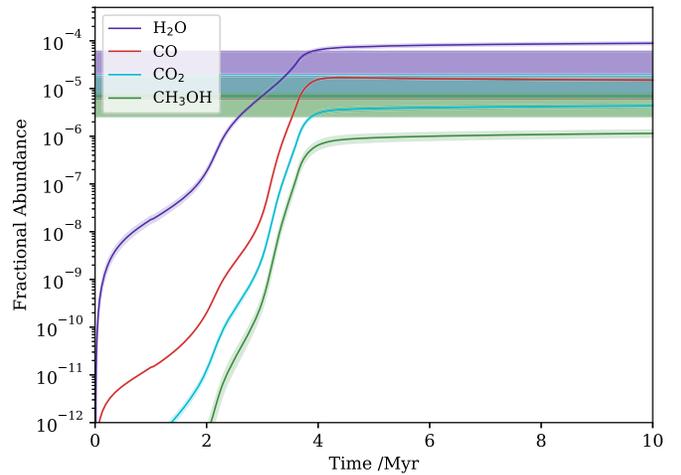
The abundances for selected species from the model runs are plotted in Figures 2 and 3. It can immediately be seen from Figure 2 that, for the species with observational constraints, the uncertainty in the rates does not lead to a large uncertainty in the model abundances. The model appears to consistently underproduce  $\text{CH}_3\text{OH}$  and overproduce  $\text{H}_2\text{O}$ . However, the difference is small if the errors on the observations are accounted for. This is a good result for such a simple model. It may be that if the reduced network was expanded, these results would be improved. Equally, it may be that the constraints are broad enough that a poor  $\text{CH}_3\text{OH}$  abundance is not affecting the overall likelihood as much as a poor  $\text{CO}$  or  $\text{CO}_2$  abundance would.

The fractional abundances of the species with upper limits were not as well constrained, and so it might be expected that they are much more varied. In Figure 3, it can be seen that this is not the case, which implies that their abundances are also strongly tied to the rates of the reactions in Figure 1. However, the average abundance of  $\text{OCS}$  is an order of magnitude higher than the observed upper limit. Examining the full abundance distribution, it appears there is a fraction of the model runs that fall outside the 67% confidence interval and give  $\text{OCS}$  abundances below the upper limits. This illustrates the problem inherent in drawing from the marginalized posteriors. Drawing from each marginalized posterior individually gives sets of reaction rates that break the upper limit used to infer the posterior distributions.

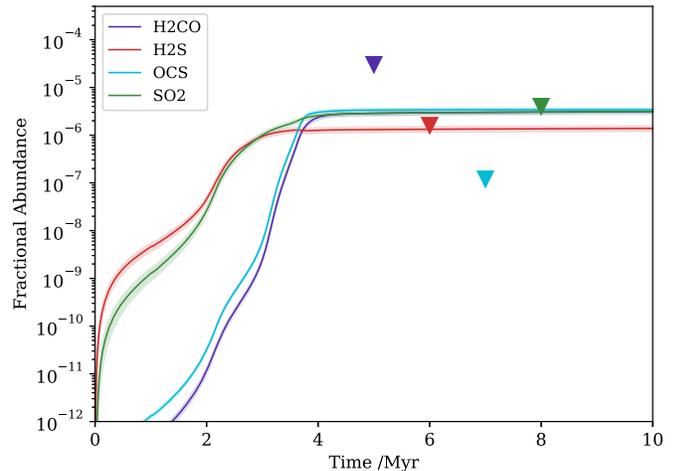
### 5.2. Network Connectivity

The PPDs presented in Figure 1 are marginalized, that is to say they represent the likelihood of a given reaction rate averaged over the values of the other rates. However, not all rates are independent, and it is possible that some areas of the rate space are only likely for one reaction when a second takes a particular value. To investigate this, the joint posteriors of pairs of reactions were examined. These give the likelihood of pairs of reaction rate values so that it can be seen whether the two reactions are in some way correlated.

For example, the joint probability distribution of reactions 1 and 2 is shown in Figure 4. It can be seen that either reaction 1 or reaction 2 can take a value much higher than their respective most likely value, but only when the other is at its most likely value. This shows that in reality there must be a certain amount of  $\text{O}$  converted to  $\text{H}_2\text{O}$  in the model, and as long as one step in



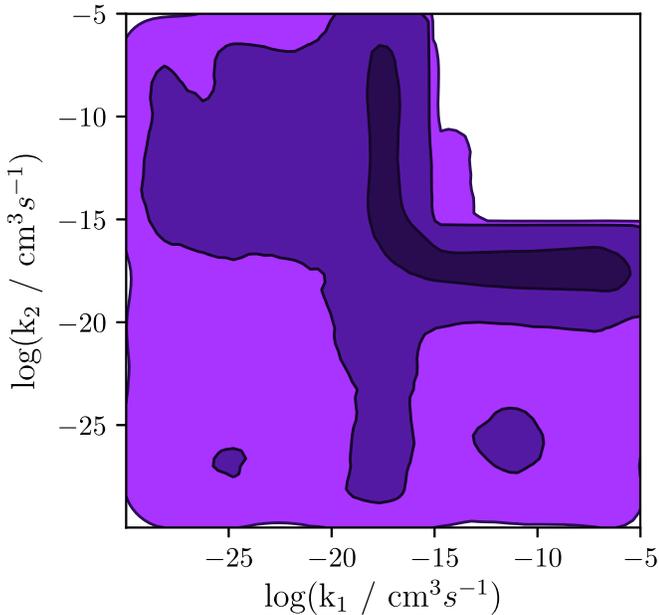
**Figure 2.** Fractional abundances for the constrained species from 1000 model runs. The rates of each reaction are randomly sampled from the marginalized posterior distributions. The average abundances produced by the models at each time are plotted along with the 67% confidence interval displayed as a shaded area.



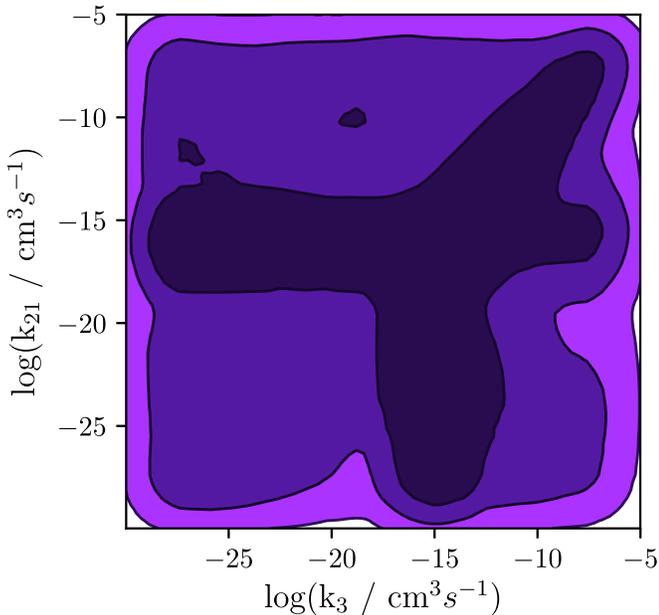
**Figure 3.** Similar to Figure 2, but for the species with upper limits. The values of the upper limits are plotted as triangles.

that process limits the rate to the correct amount then the other can freely vary. In order to break this degeneracy, limits on the  $\text{OH}$  abundance in the ice are required. Similar joint distributions are seen for reactions 22–24, as a certain amount of  $\text{CO}$  must be converted to  $\text{CH}_3\text{OH}$ .

In Figure 5, the joint probability distributions of reactions 3 and 21 are plotted. Both reactions are less well constrained than reactions 1 or 2. This can be seen from the large area taken up by the  $1\sigma$  contour. However, the high probability density areas are those where at least one reaction takes the most likely value from their respective marginalized posterior distributions. There is also a line of increased probability density where both reactions have approximately equal rates that are higher than the peak value. It is therefore likely that the reactions compete for  $\text{CO}$  and the rate of reaction 3 is poorly constrained as the availability of  $\text{CO}$  is the main factor in the amount of  $\text{CO}_2$  produced. Tighter observational ranges on the abundances of  $\text{CO}$ ,  $\text{CO}_2$ , and  $\text{CH}_3\text{OH}$  may reduce this degeneracy and allow the rate of reaction 3 to be more clearly determined.



**Figure 4.** Joint probability distributions of the rates of reactions 1 and 2; darker areas represent higher probability densities. The  $1\sigma$ ,  $2\sigma$ , and  $3\sigma$  contours are plotted. These two reactions are tightly coupled, so one must take the value of  $\sim 5.0 \times 10^{-18} \text{ cm}^3 \text{ s}^{-1}$  and act as the rate limiting step. The other must then have this value or higher in order for enough  $\text{H}_2\text{O}$  to be produced in the model.



**Figure 5.** Similar to Figure 4, but for reactions 3 and 21. While a large amount of the probability density is at the location of the peaks in the marginalized posterior distributions, there is a noticeable correlation when either increases above the most likely rate. This is likely due to the fact that both reactions destroy CO and are therefore competing to produce enough  $\text{CO}_2$  and  $\text{HCO}$ , respectively.

## 6. Conclusions

A novel way to tackle uncertainty about surface reactions and rate coefficients using Bayesian inference is presented. To prove the efficiency of Bayesian techniques in providing insight on the chemical parameters of surface reactions, the algorithm was tested with a proof-of-concept example. A simple chemical code was created by parameterizing the

freezeout of important species and neglecting all other adsorption to the grain surface. This left a model where only grain surface chemistry needed to be accounted for, greatly reducing the complexity and opening up the possibility of exploring a large parameter space.

The rates of the reactions in the chemical model were found through Bayesian inference. Using an MCMC sampling algorithm, the model was run with varying reaction rates, and the likelihood of the model was evaluated each time. This likelihood was calculated by comparing the model to observations of ices toward background stars, which are reasonable values for a molecular cloud. It was possible to strongly constrain the rates of reactions that are involved in the production or destruction of species for which measurements exist. These rates are presented in Table 4.

Future improvements should include a more complex chemical code, including the grain surface reactions directly in a gas-grain chemical code. This would allow for an improved treatment of the freezeout and nonthermal desorption among other effects. However, the added complexity would make this a vastly more computationally intensive procedure, initial tests with UCLCHEM taking approximately 1000 times longer per run. Improved rates could also be achieved by including more observational data, particularly constraining species for which there are currently only upper limits. The parameter space could also be reduced by including the PPDs from the results of this work as priors in future work.

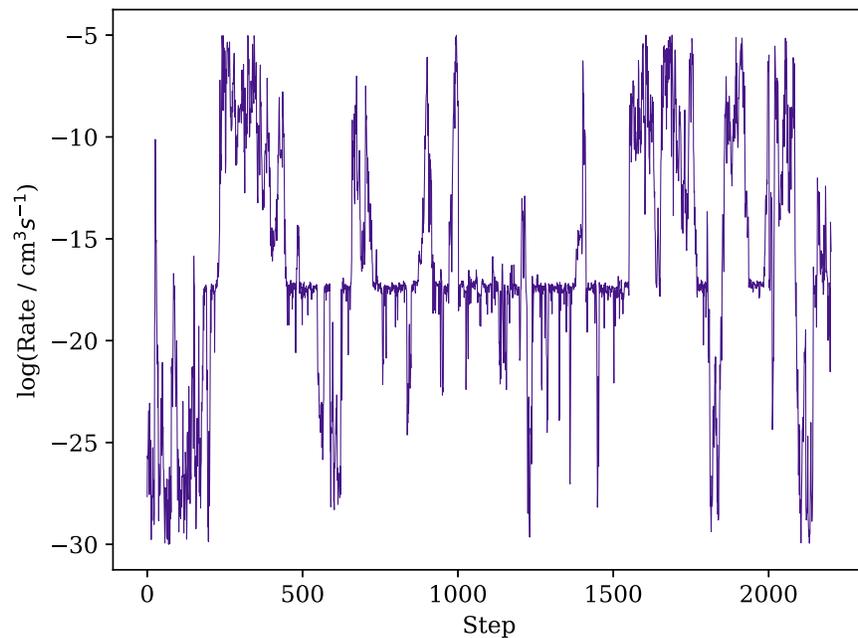
The authors thank the referees for their constructive comments that greatly improved this work. They also thank Damien De Mijolla for conversations relating to the likelihood function. J.H. is funded by the STFC grant ST/M001334/1. A.M. was partially supported by an IMPACT studentship. The authors also acknowledge DiRAC for use of their HPC system, which allowed this work to be performed.

## Appendix Convergence

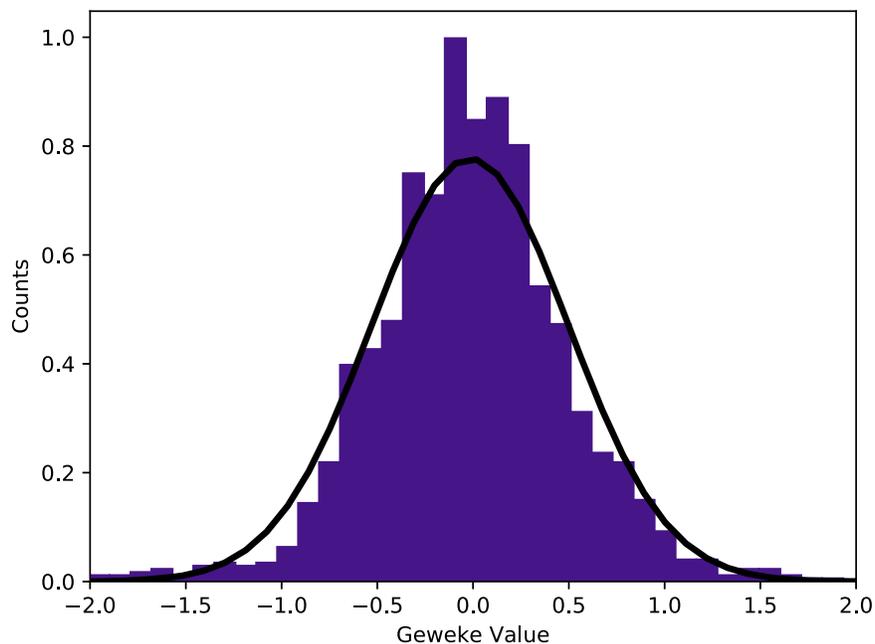
It is important to understand whether an MCMC procedure has converged to a stationary distribution. The sampling is not complete if increasing the length of that chains would considerably alter the posterior distribution of the reaction rates. It is not possible to be certain that convergence has been reached, but several heuristics are available and are considered in this appendix.

Most simply, the chains themselves can be inspected. Figure 6 shows every 500th step in an example chain for reaction 1. The walker repeatedly leaves the area of maximum likelihood and then returns. This cycle repeats a sufficient number of times for it to be unlikely that there is an undiscovered mode.

More rigorously, the Geweke diagnostic can be used (Geweke 1992). In this test, it is considered that if the chain has converged any two samples of the chain will have the same mean, within the variance of the samples. This is typically tested on the first 10% and the final 50% of a chain that is thought to have converged. In this work, a sample of chains was tested by breaking each chain into 10 subsamples and comparing each to the mean of the final 50% of the chain. In every case, the mean of the subsample was consistent with the mean of the larger sample. Since the value of the Geweke diagnostic should be zero, within the variance of the chains,



**Figure 6.** Trace plots showing the sampled rate of reaction 1 for every 500th step in an example chain. The chain repeatedly returns to the optimal rate while fully exploring the range of possible rates.



**Figure 7.** Geweke diagnostic for subsamples of the MCMC chains run for this work. The distribution of the values of the diagnostic should follow a normal distribution if the chains have converged. A normal distribution with the same standard deviation as the Geweke diagnostic samples has been plotted in black for comparison.

multiple values for converged chains should follow a normal distribution. In Figure 7, the values of the diagnostic for many subsamples of the chains from this work are plotted as a histogram with a normal distribution plotted for comparison.

The autocorrelation time is another diagnostic that can be calculated, and the tutorial provided in the emcee documentation<sup>1</sup> is used for this. This provides two heuristics. First, once the chains reach a sufficient length that the autocorrelation time

can be reliably calculated, it is likely that the chain has converged. Second, one use of the quantity is to calculate the sampling error in an MCMC chain. An autocorrelation time of  $10^4$  steps was calculated from the chains, effectively giving approximately 100 independent samples per chain. If the mean value of the chain is considered, then the variance on this mean is given through the equation

$$\sigma^2 = \frac{\tau_f}{N} \text{Var}[f(\theta)], \quad (6)$$

<sup>1</sup> <https://emcee.readthedocs.io/en/latest/tutorials/autocorr/>

where  $N$  is the number of sample and  $\text{Var}[f(\theta)]$  is the variance of the chain. In the case of an average chain in this work, the sampling error on the value of the mean is  $\sim 1\%$ .

Finally, the posterior distribution was also evaluated using the code `pyMultinest` (Buchner et al. 2014), and the marginalized posteriors were consistent with those found using `emcee`. The consistency between these two different methods of sampling the posterior is good evidence for convergence. Ultimately, given the above heuristics and the fact that the initial MCMC runs produced approximately the same results for chains of 100,000 steps as they do in the  $\sim 10^6$  step chains used in the final work, it is assumed that the chains have, in fact, converged.

### ORCID iDs

J. Holdship  <https://orcid.org/0000-0003-4025-1552>

N. Jeffrey  <https://orcid.org/0000-0003-2927-1800>

S. Viti  <https://orcid.org/0000-0001-8504-8844>

J. Yates  <https://orcid.org/0000-0003-1954-8749>

### References

- Bevan, A. 2018, *MNRAS*, 480, 4659
- Boogert, A. C. A., Gerakines, P. A., & Whittet, D. C. B. 2015, *ARA&A*, 53, 541
- Buchner, J., Georgakakis, a., Nandra, K., et al. 2014, *A&A*, 564, A125
- Caselli, P., & Ceccarelli, C. 2012, *A&ARv*, 20, 56
- Ceccarelli, C., Caselli, P., Herbst, E., Tielens, A. G. G. M., & Caux, E. 2007, in *Protostars and Planets V*, ed. B. Reipurth, D. Jewitt, & K. Keil (Tucson, AZ: Univ. Arizona Press), 47
- Chang, Q., Cuppen, H. M., & Herbst, E. 2007, *A&A*, 469, 973
- Chuang, K.-J., Fedoseev, G., Ioppolo, S., van Dishoeck, E., & Linnartz, H. 2016, *MNRAS*, 455, 1702
- Collings, M. P., & McCoustra, M. R. S. 2006, in *IAU Symp. 231, Astrochemistry: Recent Successes and Current Challenges*, ed. D. C. Lis, G. A. Blake, & E. Herbst (Cambridge: Cambridge Univ. Press), 405
- Draine, B. 2003, *ARA&A*, 41, 241
- Foreman-Mackey, D., Hogg, D. W., Lang, D., & Goodman, J. 2013, *PASP*, 306
- Fraser, H. J., Collings, M. P., Dever, J. W., & McCoustra, M. R. S. 2004, *MNRAS*, 353, 59
- Fuchs, G. W., Cuppen, H. M., Ioppolo, S., et al. 2009, *A&A*, 505, 629
- Geweke, J. 1992, *Evaluating the Accuracy of Sampling-Based Approaches to the Calculations of Posterior Moments*, Vol. 4 (Oxford: Clarendon)
- Goodman, J., & Weare, J. 2010, *CAMCS*, 5, 65
- Hagen, W., Allamandola, L. J., & Greenberg, J. M. 1979, *Ap&SS*, 65, 215
- Hasegawa, T. I., Herbst, E., & Leung, C. M. 1992, *ApJS*, 82, 167
- Herbst, E., & Van Dishoeck, E. F. 2009, *ARA&A*, 47, 427
- Holdship, J., Viti, S., Jiménez-Serra, I., Makrymallis, A., & Priestley, F. 2017, *AJ*, 154, 38
- Ioppolo, S., Palumbo, M. E., Baratta, G. A., & Mennella, V. 2009, *A&A*, 493, 1017
- Ioppolo, S., van Boheemen, Y., Cuppen, H. M., van Dishoeck, E. F., & Linnartz, H. 2011, *MNRAS*, 413, 2281
- Klein, J. P., & Moeschberger, M. L. 2003, *Survival Analysis: Techniques for Censored and Truncated Data* (Berlin: Springer)
- Lomax, O., Whitworth, A. P., & Hubber, D. A. 2016, *MNRAS*, 458, 1242
- Makrymallis, A., & Viti, S. 2014, *ApJ*, 794, 45
- McElroy, D., Walsh, C., Markwick, A. J., et al. 2013, *A&A*, 550, A36
- Minissale, M., Loison, J.-C., Baouche, S., et al. 2015, *A&A*, 577, A2
- Occhiogrosso, A., Viti, S., Ward, M. D., & Price, S. D. 2012, *MNRAS*, 427, 2450
- Palau, A., Zapata, L. A., Rodríguez, L. F., et al. 2014, *MNRAS*, 444, 833
- Parise, B., Ceccarelli, C., & Maret, S. 2005, *A&A*, 441, 171
- Pirronello, V., Brown, W. L., Lanzerotti, L. J., Marcartonio, K. J., & Simmons, E. H. 1982, *ApJ*, 262, 636
- Pirronello, V., Liu, C., Shen, L., & Vidali, G. 1997, *ApJ*, 475, L69
- Prager, J., Najm, H. N., & Zádor, J. 2013, in *Proc. Combust. Inst.* (New York: Elsevier), 583
- Schmalzl, M., Visser, R., Walsh, C., et al. 2014, *A&A*, 572, A81
- Testi, L., Natta, A., Scholz, A., et al. 2016, *A&A*, 593, A111
- van Dishoeck, E. F. 2004, *ARA&A*, 42, 119
- Watanabe, N., Nagaoka, A., Hidaka, H., & Kouchi, A. 2005, in *Proc. Protostars and Planets V*, ed. B. Reipurth, D. Jewitt, & K. Keil (Houston, TX: LPI), 8244
- Williams, D. A., & Cecchi-Pestellini, C. 2015, *The Chemistry of Cosmic Dust* (Cambridge: Royal Society of Chemistry)