



# Exploiting Network Topology for Accelerated Bayesian Inference of Grain Surface Reaction Networks

Johannes Heyl<sup>1</sup> , Serena Viti<sup>1,2</sup> , Jonathan Holdship<sup>1,2</sup> , and Stephen M. Feeney<sup>1</sup>

<sup>1</sup> Department of Physics and Astronomy, University College London, Gower Street, London WC1E 6BT, UK; [johannes.hey1.19@ucl.ac.uk](mailto:johannes.hey1.19@ucl.ac.uk)

<sup>2</sup> Leiden Observatory, Leiden University, P.O. Box 9513, 2300 RA Leiden, The Netherlands

Received 2020 July 18; revised 2020 September 29; accepted 2020 October 5; published 2020 December 4

## Abstract

In the study of grain-surface chemistry in the interstellar medium, there exists much uncertainty regarding the reaction mechanisms with few constraints on the abundances of grain-surface molecules. Bayesian inference can be performed to determine the likely reaction rates. In this work, we consider methods for reducing the computational expense of performing Bayesian inference on a reaction network by looking at the geometry of the network. Two methods of exploiting the topology of the reaction network are presented. One involves reducing a reaction network to just the reaction chains with constraints on them. After this, new constraints are added to the reaction network and it is shown that one can separate this new reaction network into subnetworks. The fact that networks can be separated into subnetworks is particularly important for the reaction networks of interstellar complex-organic molecules, whose surface reaction networks may have hundreds of reactions. Both methods allow the maximum-posterior reaction rate to be recovered with minimal bias.

*Unified Astronomy Thesaurus concepts:* [Astrostatistics strategies \(1885\)](#); [Astrochemistry \(75\)](#); [Reaction rates \(2081\)](#); [Interstellar abundances \(832\)](#); [Dark interstellar clouds \(352\)](#)

## 1. Introduction

Interstellar dust is a crucial part of the interstellar medium. The dust grains are responsible for much of the rich chemistry observed, as they act as an energy sink, allowing reactions between adsorbed atoms and molecules (van Dishoeck 2014). There is strong evidence to suggest that complex-organic molecules (COMs) form on these interstellar dust grains (Herbst & van Dishoeck 2009). In fact, this rich chemistry is thought to take place long before stars have formed, during the dark cloud phase. In the last couple of decades many experiments have been performed to determine the surface reactions occurring on the icy mantles (e.g., see the review by Linnartz et al. 2015). However, surface reactions in the laboratory are typically only investigated within a narrow range of physical parameters constrained by what is available with the experimental techniques employed and thus often not fully representative of the ISM conditions. Hence the available experimental data for interstellar ices are limited.

Bayesian inference has become a standard tool in astrophysics for determining model parameters from observations. In recent years, it has also become a tool in astrochemistry (Makrymallis & Viti 2014; Holdship et al. 2018; de Mijolla et al. 2019). However, by considering increasingly rich chemistry one must ultimately consider more complicated reaction networks. This results in an increased computational cost. There exists much in the literature regarding chemical network reduction as a means of reducing the computational complexity of the problem being solved (Ayilaran et al. 2019). An understanding of chemical reaction networks and how to simplify them has become increasingly crucial in astrochemistry (Grassi et al. 2012; Xu et al. 2019). However, these methods have primarily focused on simplifying the network for the forward problem. For example, Xu et al. (2019) adopted an iterative approach, where they evaluated the importance of each species at each timestep. The advantage of Bayesian inference is that it provides probability distributions for the parameters of


interest conditioned on the available data, thereby allowing us to quantify the uncertainties on these parameters. However, there is the issue that not all the parameters of interest can be determined. The iterative approach mentioned above only focuses on the reactions to which species are the most sensitive. This allows for the reduction in computational expense, unlike for Bayesian inference. In this work we look to use various features of the network topology to reduce the computational expense of the inference process.

In this work, we build upon the work done by Holdship et al. (2018; hereafter H18) and use the same reaction network to highlight how aspects of the geometry of the network can be exploited to determine the model parameters at reduced computational expense. It should be noted that this reaction network only has 24 reactions with four constraints, but this technique should generalize to larger networks. It is hoped this will prove particularly useful when considering reaction networks of COMs, where the number of reactions is large and the number of constraints is small.

We begin by first presenting the chemical network used in Section 2. Following this, we argue why we can reduce the network with its constraints to a simpler one as well as present explanations for how the positions of particular constraints in the networks are crucial. We present the results of the Bayesian inference for this. We then go on to discuss how specific aspects of the topology of the reaction network are useful to consider and how these influence our choices when we reduce the network. Finally, we look at how we can separate a reaction network into smaller subnetworks.

## 2. The Chemical Network

We use the same chemical network as in H18, which we set out pictorially in Figure 1. We list all reactions in Table 1, assigning them numbers that we use throughout the paper for brevity. For the sake of simplicity, the hydrogen abundance is not a conserved quantity, as it is typically  $\sim 10^4$  times more



**Figure 1.** A diagram of the chemical reaction network considered. For the sake of simplicity, any reactions with hydrogen and oxygen are represented with H and O next to the arrow. For the case where a molecule can be formed in multiple reactions, such as for OCS, the arrow colors pointing to that molecule indicate the reactants. For example, the dashed-dotted orange arrows that point from HS and CO to OCS indicate that these two molecules form OCS. Molecules in blue boxes have constraints on their final abundances. Molecules in white boxes have upper limits on their abundances.

**Table 1**

Table of the Reactions Used in This Work Taken from Holdship et al. (2018)

| Reaction No. | Reaction                           |
|--------------|------------------------------------|
| 1            | $O + H \rightarrow OH$             |
| 2            | $O + H \rightarrow H_2O$           |
| 3            | $CO + OH \rightarrow CO_2$         |
| 4            | $S + H \rightarrow HS$             |
| 5            | $HS + H \rightarrow H_2S$          |
| 6            | $H_2S + S \rightarrow H_2S_2$      |
| 7            | $CS + H \rightarrow HCS$           |
| 8            | $HCS + H \rightarrow H_2CS$        |
| 9            | $CO + S \rightarrow OCS$           |
| 10           | $OCS + H \rightarrow HOCS$         |
| 11           | $H_2S + CO \rightarrow OCS$        |
| 12           | $H_2S + H_2S \rightarrow H_2S_2$   |
| 13           | $H_2S_2 + CO \rightarrow CS_2 + O$ |
| 14           | $H_2S + O \rightarrow SO_2$        |
| 15           | $CS_2 + O \rightarrow OCS + S$     |
| 16           | $CO + HS \rightarrow OCS$          |
| 17           | $S + O \rightarrow SO$             |
| 18           | $SO + O \rightarrow SO_2$          |
| 19           | $SO + H \rightarrow HSO$           |
| 20           | $HSO + H \rightarrow SO$           |
| 21           | $CO + H \rightarrow HCO$           |
| 22           | $HCO + H \rightarrow H_2CO$        |
| 23           | $H_2CO + H \rightarrow H_3CO$      |
| 24           | $H_3CO + H \rightarrow CH_3OH$     |

**Table 2**

A Table Listing All the Various Network Configurations Used and Referred to throughout This Work

| Configuration No. | Reactions Used | Molecules with Constraints                                                                          |
|-------------------|----------------|-----------------------------------------------------------------------------------------------------|
| 1                 | 1–24           | CO, CH <sub>3</sub> OH, CO <sub>2</sub> , H <sub>2</sub> O                                          |
| 2                 | 1–6, 9–24      | CO, CH <sub>3</sub> OH, CO <sub>2</sub> , H <sub>2</sub> O                                          |
| 3                 | 1–3, 21–24     | CO, CH <sub>3</sub> OH, CO <sub>2</sub> , H <sub>2</sub> O                                          |
| 4                 | 1–24           | CH <sub>3</sub> OH, CO <sub>2</sub> , H <sub>2</sub> O                                              |
| 5                 | 1–24           | CO, CO <sub>2</sub> , H <sub>2</sub> O                                                              |
| 6                 | 1–24           | CO, CH <sub>3</sub> OH, CO <sub>2</sub> , H <sub>2</sub> O, SO <sub>2</sub> , OCS, H <sub>2</sub> S |
| 7                 | 4–20           | CO, SO <sub>2</sub> , OCS, H <sub>2</sub> S                                                         |
| 8                 | 4–20           | SO <sub>2</sub> , OCS, H <sub>2</sub> S                                                             |
| 9                 | 1–3, 21–24     | CH <sub>3</sub> OH, CO <sub>2</sub> , H <sub>2</sub> O                                              |

abundant than any other molecule, so its loss in this reaction network is negligible. We also emphasize that this network is a toy model and is not meant to properly reflect a complete grain surface network but simply serves as a proof-of-concept.

A modified version of the open source gas-grain code UCLCHEM was used (Holdship et al. 2017). This version only considered the surface chemistry of a collapsing dark cloud. The cloud is modeled as collapsing from a density of  $10^2 \text{ cm}^{-3}$  to  $2 \times 10^4 \text{ cm}^{-3}$  over 10 million years. The collapse takes place isothermally at 10 K. As chemistry only occurs on the grain surface, freeze-out rates are required. Freeze-out rates are the rates at which atomic and molecular species stick to the dust grains and build up ices (Fraser et al. 2003; Hocuk et al. 2014). The freeze-out rates were found in H18 by running a single-point model of the full UCLCHEM code. The species that were given freeze-out rates were CO, CS, O, H, OH, and S. More details can be found in H18.

As already mentioned above, the small chemical network we use for this work is not meant to represent a comprehensive surface network. Nevertheless, the choice of most of the reactions was based on the results of experimental studies. For example, the successive hydrogenation of CO to form CH<sub>3</sub>OH has been studied in detail and is considered to be the dominant reaction pathway (Chuang et al. 2016). Similarly, CO<sub>2</sub> has been found to be efficiently formed when CO and OH react (Ioppolo et al. 2011). Beside a small network representing the main routes of carbon- and oxygen-bearing species on the ices, we chose to include a small network producing sulfur-bearing species, since there is still much unknown about the form that ultimately sulfur takes on the ices during the cold phase of the star formation process (Woods et al. 2015; Laas & Caselli 2019).

In this work we consider a number of variants of the chemical network shown in Figure 1. These configurations, which differ in terms of the reactions and/or constraints used, are enumerated in Table 2: the configuration numbers will be used throughout the work. The combination of the full reaction network shown in Figure 1 with the abundance measurements as listed in Table 3 is Configuration 1.

### 3. Bayesian Inference

#### 3.1. Introduction to Bayesian Inference

As in H18, our aim is to determine the 24 reaction rates, which we represent as a vector,  $\mathbf{k} = (k_1, k_2 \dots k_{24})$ . We are therefore faced with a 24D inference problem. The grain code used took in these reaction rates and produced the corresponding abundances, which are represented by a vector  $\mathbf{Y} = (Y_1, Y_2 \dots Y_{23})$ . Henceforth,

**Table 3**

The Abundances and Uncertainties for The Molecules with Observed Values Taken from Boogert et al. 2015

| Species            | Abundances Relative to H       |
|--------------------|--------------------------------|
| 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 | $(5.2 \pm 2.4) \times 10^{-6}$ |

we refer to the “forward model” when we input a particular value of  $\mathbf{k}$  to obtain some  $\mathbf{Y}$ .

We know the abundances of four of the molecules in the ices of the network shown in Figure 1, which are in blue boxes. These are taken from Boogert et al. (2015) and listed in Table 3. We are looking to solve the “inverse problem,” i.e., what values of  $\mathbf{k}$  yield values of  $\mathbf{Y}$  that match the observations best? Such a problem naturally lends itself to a Bayesian approach. The inherent degeneracy of this problem should be noted. Observations only exist for 4 of the 24 molecules. This suggests that the rates of the reactions that do not influence the abundances of these four molecules will be poorly constrained (if at all), and there will be many values of these rates that give the same observations. For a discussion of the degeneracies, please refer to H18. Exploiting the low number of constraints in order to speed up the inference process is a crucial point in this work.

We use Bayes’ Law to determine the probability distribution of the values of the reaction rates

$$P(\mathbf{k}|\mathbf{d}) = \frac{P(\mathbf{d}|\mathbf{k})P(\mathbf{k})}{P(\mathbf{d})}, \quad (1)$$

where  $P(\mathbf{k}|\mathbf{d})$  is the posterior probability distribution,  $P(\mathbf{k})$  is the prior,  $P(\mathbf{d}|\mathbf{k})$  is the likelihood, and  $P(\mathbf{d})$  is referred to as the evidence. Here  $\mathbf{d}$  is the data, i.e., the observed values of the abundances. The evidence is a normalizing factor and is typically difficult to evaluate. However, as it is independent of  $\mathbf{k}$ , we can instead just consider

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

which is just the unnormalized posterior.

### 3.2. Implementation

Evaluating the reaction rate posterior requires specification of a prior on the reaction rates and a likelihood. We chose the same prior as in H18: a log-uniform distribution between  $10^{-30}$  and  $10^{-5}$ . From chemical considerations, we know that these reaction rates are ultimately very fast. Typically, we might therefore select a prior that favors higher reaction rates. We chose, however, to ignore this information, instead following the traditional approach of using a log-uniform prior, which equally weights rates over a range of orders of magnitude.<sup>3</sup> It is important to note that, despite our motivation for this prior, we realize that we are in a prior-dominated regime, which we demonstrate in Appendix B with very different prior assumptions. Unlike in the data-dominated regime, these prior-

dominated posteriors differ significantly among themselves. This, however, does not detract from the analysis we conduct in this work.

A Gaussian likelihood function was used, which takes the form

$$P(\mathbf{d}|\mathbf{k}) = \prod_{i=1}^{n_d} \frac{1}{\sqrt{2\pi}\sigma_i} \exp\left(-\frac{(d_i - Y_i)^2}{2\sigma_i^2}\right), \quad (3)$$

where  $n_d$  is the number of observations and  $\sigma_i$  is the uncertainty of the  $i$ th observation. We only multiply over the species that have observed abundances. We refer to these observed abundances as constraints as they constrain the parameter space of our reaction rate posteriors.

In order to determine the posterior, the PyMC3 Python package was utilized (Salvatier et al. 2016). The PyMC3 package includes a range of samplers. Here, we used the Metropolis–Hastings algorithm, a simple Markov Chain Monte Carlo (MCMC) method. A Gaussian proposal distribution was used. Before each run, 500 tuning steps were initially taken to determine the optimal covariance of the proposal distribution. These tuning steps were not included in our analysis and were discarded. For the relatively low number of dimensions (<50), a point-based sampler such as this one is suitable. 50 chains of length  $10^6$  were used to sample the posterior probability space. We created a Python wrapper of the grain code using F2Py that was then fed values of  $\mathbf{k}$  during the sampling process (Peterson 2009).

Though simple sampling methods suffice for the dimensionality of the posteriors considered here, the same will not be true for more complex networks. A key point that needs to be considered is the limitations of many MCMC methods as the number of dimensions increases. Brewer & Foreman-Mackey (2018) discuss how some widely used samplers struggle to give sensible results in higher dimensions. The popular emcee Python package that was used in H18 is discussed to be useful for when the number of dimensions is fewer than 50, with it struggling in higher dimensions (Huijser et al. 2015). Even for the case where the sampler does not struggle as the number of dimensions increases, the time taken to run the inference process will still increase. This increase in computation time will eventually become prohibitive. In Section 6 we will argue that we can split our reaction network up into subnetworks on which we can perform Bayesian inference. By carefully placing the “cut” on the reaction network, we then show that we can reproduce the results of the full reaction network inference with these subnetworks. Each subnetwork has lower dimensionality than the original network, meaning its rates can be inferred more quickly and by simpler samplers. Additionally, the inference process on all the subnetworks can be run in parallel.

### 3.3. Constraints

It is essential to include constraints in order to formulate a likelihood function. In H18, four constraints for molecules in the reaction network were taken. Table 3 shows the abundances of the molecules taken from Boogert et al. (2015). These ice abundances are derived from ice band profiles. The column densities can be calculated using the integrated optical depth as well as the integrated band strength. The latter were determined from laboratory experiments by Boogert et al. (2015). From the table, it is clear that the strongest constraint is on H<sub>2</sub>O, which is

<sup>3</sup> It is worth noting that the use of a single prior to represent complete ignorance (Walley 2000; Norton 2008) has received criticism. Complete ignorance can be represented by repeating the analysis using several priors that significantly differ from one another (see Fischer 2019 for a straightforward application to a simple problem of chemical kinetics).

known to exist at the  $3\sigma$  level, whereas the other molecules’ abundances differ from zero at only  $1.5\text{--}2.2\sigma$ .

H18 also reformulated the likelihood function to include upper bounds on the abundances of OCS,  $\text{H}_2\text{S}$ ,  $\text{SO}_2$ , and  $\text{H}_2\text{CO}$ . This was not found to have a significant effect on the reaction rates determined and so we do not include these upper bounds in the following work. For the rest of this work we use a likelihood function of the form in Equation (3).

## 4. Network Reduction Methods

### 4.1. Overview

Galagali & Marzouk (2019) considered a similar problem to the one discussed in this work in the context of system biology. There, they considered the case of a reaction network with a single observation. The main differences between the networks they considered and the one being considered here is the absence of enzymes as well as the absence of reversible reactions.

They defined the “effective reaction network” as the subset of reactions that must be kept in order to produce the same values of the observable. This is a useful concept to consider, especially in the context of network connectivity. Some subsets of reactions will evolve completely independently of one another, with there being no competition for chemical species. While this may seem unlikely in the context considered here, it is true when one assumes that hydrogen’s abundance is significantly higher than that of any other species. This can be seen in Figure 1, where the successive hydrogenation of CS to form  $\text{H}_2\text{CS}$  is clearly independent of the rest of the reaction network. This reaction chain will be referred to as the “ $\text{H}_2\text{CS}$  chain” from now on.

### 4.2. Network Reduction of Nonconnected Networks

We consider the “ $\text{H}_2\text{CS}$  chain” in greater depth. Under the assumption that there is no competition for hydrogen, we should expect the “ $\text{H}_2\text{CS}$  chain” and the other reaction network with the remaining 22 reactions (Configuration 2) to evolve completely independently, as the species do not interact. We infer the reaction rate posteriors for Configurations 1 and 2. Figure 2 shows the posteriors on the common reactions to be essentially identical. Any differences are due to the specific samples drawn in the MCMC chains. This intuitively makes sense as the ordinary differential equations that govern the two subnetworks will evolve independently of one another. Since none of the reactions in the “ $\text{H}_2\text{CS}$  chain” are constrained in Configuration 1, it stands to reason that these reaction rates are nuisance parameters. As such, removing these two reactions from the inference should make no difference.

While the example given might seem trivial, one needs to consider under what circumstances one might have a reaction network with a disconnected segment. In surface grain chemistry, the molecules must contend with both an activation energy barrier as well as a diffusion energy barrier. Only if both of these can be overcome, is the reaction likely to happen efficiently. If either one of these barriers is too high, then one can in fact approximate that reaction as not happening and “cut” off that reaction. This might lead one to separate a reaction chain from the rest of the network. In this example, one could conceivably imagine a hypothetical reaction being possible between  $\text{H}_2\text{CS}$  and any other molecule in

Configuration 2, but the activation energy barrier is too high to overcome at 10 K. It is therefore simpler to exclude it.

## 5. Further Network Reduction

In the previous section, we observed that the network connectivity of a chemical reaction network can allow us to discard reactions that do not influence the values of observed abundances. In this section, we develop this idea further by arguing that the locations of the constraints in the reaction network allow us to discard more reactions.

We wish to emphasize once again that we are not seeking to make quantitative predictions about the reaction rates. Instead, we are looking to develop a qualitative understanding of the kinetics as well as develop an intuition for how the methods that will be discussed in this work can be applied to other astrochemical modeling scenarios.


### 5.1. Reducing the Network

We begin by briefly returning to the posteriors in Figure 2. The uncertainties on the rates of reactions 1 and 2 are significantly smaller than for reactions 3 and 21–24. This relates to the uncertainties on the abundances of the molecules, as discussed in Section 3.3. In the limit that the abundances of all molecules involved in a reaction are perfectly known, one would expect the posterior distribution of the reaction rate to approach a Dirac delta function. The greater confidence level in water’s presence is therefore responsible for the tighter constraints on the rates of Reactions 1 and 2. Improving the precision of the abundance measurements of CO,  $\text{CO}_2$ , and  $\text{CH}_3\text{OH}$  would in turn tighten the constraints on their reaction rates.

In Figure 1, we observe that all the reactions whose reaction rates are constrained have a constraint at the end of their respective chain. Consider the successive hydrogenation of carbon monoxide to form methanol, henceforth referred to as the “methanol chain.” The fact that there is a constraint present at the end is significant. By constraining the amount of methanol, one effectively constrains the reaction rates of its precursors,  $\text{CH}_3\text{O}$  and  $\text{H}_2\text{CO}$  from below, as the existence of methanol requires its precursors to have been produced. If the reaction rate of these reactions is too high, then too much methanol will be produced. However, there is an inherent degeneracy in the rates of the intermediate reactions, so it is unclear how the rates are partitioned between the two reactions. What one finds is that these intermediate reaction rates are coupled, by observing their joint probability distributions. One reaction will serve as the rate-limiting reaction, with the other compensating for this by being significantly faster to produce a sufficient amount of the final product, in this case methanol. This is discussed in more depth for the specific reactions in H18.

Additionally, the constraint on carbon monoxide constrains the abundances of these molecules from above. One can then qualitatively say that as there are now constraints on their abundances, there are therefore constraints on their reaction rates, as the reaction rates reflect how much of these molecules forms over time.

In the network shown in Figure 1, none of the sulfur-bearing molecules have constraints on them. The reaction rates of these reactions can be treated as unconstrained parameters, despite the fact that carbon monoxide, a central molecule in the sulfur-centric network, is constrained. This suggests that the subnetwork consisting of the 10 nonsulfur bearing species can be




**Figure 2.** Plots of the posterior probability distribution for the original reaction network considered in Holdship et al. (2018) as well as the 22D effective network by removing the “H<sub>2</sub>CS chain.” We observe good agreement in the shapes of the posterior distributions, with any differences due to specific samples drawn from the MCMC chains. The configuration 1 posteriors match those from H18.

treated independently. We exclude these sulfur-bearing reactions for now and only consider reactions 1–3 and 21–24. This is Configuration 3, which is a “subnetwork” of Configuration 1 as it contains a subset of the reactions.

To investigate the impact of excluding the sulfur-bearing reactions, we rerun our Bayesian inference on this configuration, producing the posterior probability distribution functions in Figure 3. With the exception of  $\text{H} + \text{CO} \rightarrow \text{HCO}$ , we recover the maximum-posterior reaction rates obtained previously. However, we see that the variances have decreased for reactions 3 and 21–24, resulting in more peaked distributions, and reactions 1 and 2 have posteriors that are less peaked than for Configuration 1.

We would like to emphasize that this is a purely artificial effect. By eliminating the sulfur-bearing reaction rates, which were essentially nuisance parameters, the variance of reactions 3 and 21–24 have decreased. It should be noted that these reactions compete for CO with the removed sulfur subnetwork. Removing the sulfur-based reactions means the nonsulfur reactions have to use up more CO. We observe that of all the reactions, the reaction  $\text{H} + \text{CO} \rightarrow \text{HCO}$  sees the greatest change between Configuration 1 and Configuration 3. In fact, we observe that a significant portion of the posterior mass is shifted from the reaction rates below the peak to the reaction rates above the peak. This, coupled with the increase in the maximum-posterior rate, suggests that the excess CO, that



**Figure 3.** Plots of the posteriors of Configuration 1, Configuration 3, and Configuration 3 with the dummy reaction  $X + \text{CO} \rightarrow \text{XCO}$ . We observe that the inclusion of this additional dummy reaction provides a better approximation to the Configuration 1 posterior than the Configuration 3 posterior does.

would normally be consumed by the sulfur reactions, is stored in the methanol chain. In particular, the fact that only the hydrogenation of CO experiences a significant change in the posterior suggest that the excess CO is stored as HCO.


At this point, it is unclear why Reactions 1 and 2 see increases in the variance of their posteriors. For these reactions, the decrease in the posterior mass under the peaks is compensated for by an increase in the posterior masses for reaction rates slower than the maximum-posterior rate. This suggests that the reactions can proceed at slower rates and still produce sufficient OH that goes on to produce  $\text{H}_2\text{O}$  and  $\text{CO}_2$ .

The MCMC runs for Configuration 3 are 2.3 times shorter than for Configuration 1, with the time taken for the runs decreasing

from 30 to 13 hr. By excluding the unconstrained reactions, we are able to reduce walltime drastically, at the cost of moderate changes to the posterior. In the next subsection, we discuss a method for recovering the full posterior. Finally, it should be emphasized that this dimension reduction must only be considered when solving the inverse problem. For a full picture of the chemistry one must include all reactions in the forward model.

### 5.2. Recovering the Full Variance

We noted previously that the decrease in the variance of the posteriors was an artificial effect. Recovering as much of this original variance as possible is critical, as without it the



**Figure 4.** The posterior probability distributions for reactions 21–24 when CO and methanol are separately removed. The original distributions are also included for comparison. We observe that for reactions 21–24, removing CO neither affects the position of the peak of the distribution nor the shape of the distribution. Removing methanol does not change reaction 21’s maximum-posterior rate, but removes all information about the reaction rates of reactions 22–24. We do not include reactions 1–3, as their posteriors are unchanged when the constraints are removed.

precision of the inferred rates will be overstated. A consideration of the reactions that were removed is a good starting point. Looking at Figure 1, one can see that from the perspective of Configuration 3, the removal of the reactions affected the CO depletion. In other words, Configuration 3 only sees that we removed CO-depleting reactions. The products of these CO-depleting reactions are not constrained, and these reactions are therefore responsible for the larger uncertainty in the posteriors of Configuration 1 compared to Configuration 3.

In order to recover this variance, one must account for the artificially altered CO depletion. As a first, simple approximation we add a fake reaction  $X + \text{CO} \rightarrow \text{XCO}$ , where  $X$  is meant to encompass all the removed reactions that consumed CO, and XCO is simply the network of products. The abundances of both  $X$  and XCO remain unconstrained. Even though  $X$  is not meant to represent a particular molecule, it still requires a freeze-out rate, as this grain-code only considers reactions that take place on the grain surface. The freeze-out rate for sulfur was used. The posteriors are shown in Figure 3. Adding a single unconstrained reaction clearly yields a good (though still imperfect) approximation to the full set of sulfur-consuming reactions in this setting, matching the variance of the full network’s rate constants more closely, and removing the bias on the inferred rate of  $\text{H} + \text{CO} \rightarrow \text{HCO}$ . Further work should be done to investigate whether increasing the number (and architecture) of “dummy reactions” aids in recovering the full posterior.

### 5.3. Network Topology Considerations

We now discuss how the placement of the constraints in the network can be significant. From the above analysis, it is not entirely clear what constraints are the most “essential.” To shed light on this problem, we focus on the methanol chain, which has a constraint on molecules at both ends.

We perform Bayesian inference on the full reaction network twice:

1. once without the CO constraint (Configuration 4), and
2. once without the  $\text{CH}_3\text{OH}$  constraint (Configuration 5).

The resulting posterior probability distribution functions for Reactions 21–24 are shown in Figure 4 alongside the distributions for the original network. The posteriors for Reactions 1–3 are not included, because these did not change significantly.

We observe that removing the constraint on CO has no effect on the reaction rates. The reaction rate posteriors for reactions 21–24 are broadly identical. This can be easily explained. In this case, one knows that a fixed amount of methanol is produced. As such, over the period of 10 Myr, a certain amount of CO has to be consumed. This results in the successive hydrogenations being constrained, which is why the reaction rate posteriors do not change.

However, removing the constraint on methanol results in the loss of constraints on three of the four reaction rates of the methanol chain, with only the hydrogenation of CO being recovered. This appears to suggest that there is a notion of “distance” between a constraint and the reaction rate of interest. Information about the subsequent reaction rates in the methanol chain is lost due to methanol’s abundance being unconstrained. In this model, we know that CO (as one of the adsorbed species) is present on the grains, but there is no information about how much of it goes into making methanol. However, it is interesting that reaction 21 remains constrained. One possible interpretation is that we know how much CO is used in reaction 3 and this possibly helps constrain how much CO is used in reaction 21. However, using this reasoning, one cannot explain why reactions 9 and 13 are not constrained, as these are also CO depletion reactions.

This suggests that in any reaction chain, some knowledge of the abundance of the end-products is required, which might be problematic when the species are undetected. However, one

can still provide theoretical predictions for these abundances that could be used.

## 6. Application to the Network with Artificial Sulfur Constraints

As a proof of concept, we wish to apply the insight from the previous section to a new grain surface network. However, this is difficult due to the limited number of observations that exist for grain-surface molecules. Boogert et al. (2015) provided upper limits for several molecules in the ice. We chose to artificially transform the upper limits on OCS, H<sub>2</sub>S, and SO<sub>2</sub>, into weak measurements by taking their abundances to be half the respective upper limit with an uncertainty of one-quarter of the upper limit. We would like to emphasize again here that we are simply trying to demonstrate how the location of these three additional constraints provides us with more knowledge of  $k$ . We do not claim this to be an accurate representation of sulfur chemistry on the ices in a dark cloud. Many theoretical and modeling studies have been recently performed investigating the sulfur depletion and the reactions on surfaces involving sulfur-bearing species and we refer the reader to such studies for a comprehensive review on the subject (e.g., Jiménez-Escobar et al. 2014; Woods et al. 2015; Vidal & Wakelam 2018; Laas & Caselli 2019).

### 6.1. The Full Network

Bayesian inference was performed for the full network with the new artificial constraints. This is Configuration 6. Despite adding these constraints elsewhere in the network, the maximum-posterior reaction rates of reactions 1–3 and 21–24 (none of which involve sulfur compounds) were found to be unchanged and the posteriors were largely similar. This fact strongly implies that the sulfur-based and non-sulfur-based reactions can be separated into subnetworks, whose reaction rates can be inferred independently, even when constraints on the sulfur-based products become available.

We also identify eight new reactions for which the marginalized posterior probability distributions deviate from uniformity. These are all reactions that involve several of the molecules whose abundances have now been constrained. The posterior probability distributions are shown in Figure 5.

### 6.2. Including the CO Constraint 1

In the following, we investigate the optimal way of splitting the full network into sulfur- and non-sulfur-based subnetworks. The two subnetworks compete over CO, one of the molecules with a constraint. We know from Section 5 that we can include the full CO constraint in the nonsulfur subnetwork without significantly biasing the inferred reaction rates. To test if this is the case for the sulfur subnetwork, we consider two cases, performing Bayesian inference on the sulfur subnetwork with the full CO constraint (Configuration 7) and leaving the CO abundance as a free parameter (Configuration 8). The time taken for Configuration 6 is about 30 hr, while the runs for Configurations 7 and 8 took about 23 hr.

In Figure 5 we compare the rate posteriors for the artificially constrained sulfur-based reactions obtained with Configurations 6–8. For the eight new rate posteriors obtained in Figure 5, we observe better agreement when the CO constraint is not included. The main explanation for this is that as several

CO depletion reactions in Configurations 7 and 8 have been discarded, the other reactions in the network are required to produce more molecules that will react to deplete the CO abundance. This can be seen by the fact that the reaction rates for the successive hydrogenation of sulfur to produce H<sub>2</sub>S are greater when the CO constraint is included. This is because H<sub>2</sub>S reacts with CO to produce OCS. To deplete the excess CO, more H<sub>2</sub>S must be produced.


It is interesting to note that one can apply the full CO constraint in the nonsulfur subnetwork, but not in the sulfur subnetwork. This is due to the relative sizes of the abundances of the constrained species. As shown in Table 4, the abundances of the sulfur-based molecules that are added are between 2 and 3 orders of magnitude less abundant than the constrained species in the other network. It should be noted that it is assumed that CO is already present to begin with and can only be consumed. No CO-formation reactions are present. As such, the contribution of the sulfur-network in depleting the CO is small compared to the nonsulfur network.

To get a better idea of the amount of CO that is used up by the subnetworks, we ran the forward models of the grain-code. By setting the reaction rates to zero, the only rates left were the freeze-out rates. These gave an idea of the total amount of CO available. This was found to be  $4.0 \times 10^{-5}$ . From Table 4, we know that the amount of CO that should be left is  $(1.2 \pm 0.8) \times 10^{-5}$ . Running the forward model of the sulfur-only network with CO as a free parameter, shows that the final CO abundance is about  $(3.3 \pm 1.1) \times 10^{-5}$ . This suggests that the nonsulfur network consumes four times as much CO as the sulfur-centric network does. By considering how the two subnetworks rely on the common constrained molecule, CO, we find that we can easily separate them.

Leaving the CO abundance unconstrained reduces the bias in the maximum-posterior reaction rates, but this does not perfectly reproduce the full network’s posterior. This is an issue, because it means that the variance of the full network is not preserved. An additional reaction of the form  $X + CO \rightarrow XCO$  could potentially be used, just as was done in Section 5.2. However, unlike before this reaction will be replacing a reaction subnetwork with constraints. This makes the problem more complicated than before, as one might need to consider how to combine the constraints to create a “constraint” for XCO. One could simply give XCO an abundance equal to the sum of all the constraints that have been replaced. One might also need to provide several “dummy” reaction chains of varying length to best recover the original posterior. Considerations of the architecture will be discussed in future work.

### 6.3. Including the CO Constraint 2

In the previous subsection, we argued that by virtue of the fact that the constraints in one subnetwork were orders of magnitude greater than in the other subnetwork, we could simply take the full CO constraint and use it in the former. Specifically, we want the CO constraint to be comparable to that of OCS, a molecule that depletes CO. However, we now consider the case where the constraints in each subnetwork are comparable in nature. To do this, we artificially reduce the abundance of CO in the system. This is done by reducing the freeze-out rate by a factor of  $10^4$ . As only grain surface




**Figure 5.** Plots of the posterior probability distribution that deviate from uniformity for the expanded reaction network. We compare the posterior distributions of Configuration 6 with those of Configurations 7 and 8. We observe better agreement of the sulfur subnetwork when we leave CO’s abundance as a free parameter, which corresponds to Configuration 8.

reactions are considered, this reduces the amount of CO available for grain-reactions by the same factor.

Obviously changing the amount of CO present in the model has an effect on reactions obtained. With so little CO, it is impossible to match the abundances of methanol and carbon dioxide. As a result of this, the posteriors of these reaction chains are close to uniform. However, the point of these simulations is not to model the chemistry accurately. We want to know what we should do with the CO constraint in each subnetwork. To aid understanding, we will consider the two subnetworks separately. We would like to note that even though the CO constraint has been reduced, the configuration for the full network still corresponds to Configuration 6.

### 6.3.1. The Nonsulphur Subnetwork

Recall that the nonsulfur subnetwork consists of reactions 1–3 and 21–24. The question is whether or not one wishes to include the reduced CO constraint to recover the posterior of the full-network (Configuration 6). Including the CO constraint gives Configuration 3 and excluding it gives Configuration 9. The posterior distributions are shown in Figure 6. We observe that the only posteriors that deviate from nonuniformity are those for the reaction rates of Reactions 1 and 2. We observe that Configuration 9, which corresponds to excluding the reduced CO constraint, matches the posterior distribution of Configuration 6 the best for Reaction 1, not only in terms of the location of the maximum-posterior but also in terms of the



**Figure 6.** The posterior distributions for Reactions 1 and 2 when the initial CO abundance is reduced by a factor of  $10^4$ . We compare the posterior distributions of Configuration 6 with those of Configurations 3 and 9. We observe the best agreement between Configurations 6 and 9, suggesting that for this subnetwork, it is better to exclude the reduced CO constraint.

**Table 4**

The Abundances and Uncertainties Taken for the Network with Artificial Sulfur Constraints

| Species            | Abundances                     |
|--------------------|--------------------------------|
| 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 | $(5.2 \pm 2.4) \times 10^{-6}$ |
| OCS                | $(6.0 \pm 3.0) \times 10^{-8}$ |
| SO <sub>2</sub>    | $(2.0 \pm 1.0) \times 10^{-6}$ |
| H <sub>2</sub> S   | $(8.0 \pm 4.0) \times 10^{-7}$ |

**Note.** For the first four species, the abundances were taken in their present form from Boogert et al. (2015). Boogert et al. (2015) provided upper bounds for the listed sulfur-based species. For the analysis in this section, the abundances of the sulfur-based species were taken to be half the upper bound value. Their uncertainties were taken to be 50%.

posterior shape. Configuration 3 is found to recover the posterior better for Reaction 2.

We also observe that the maximum-posterior reaction rates for Reactions 1 and 2 match those obtained previously in the work, even though the CO abundance was greatly reduced. This adds support to the idea that Reactions 1 and 2 form their own subnetwork and are ultimately independent.

#### 6.3.2. The Sulfur Subnetwork

Figure 7 shows the nonuniform posteriors of the reaction rates for the sulfur subnetwork. We find that the maximum-posterior rates for these reactions do not change when we discard Reactions 1–3 and 21–24, regardless of whether we include CO’s new abundance constraint. As before, however, the precise forms of some of the posteriors are very different. This did not appear to be as much of an issue in Figure 6, which might be related to the relative levels of uncertainty on the relevant species, as discussed in Section 5.1.

#### 6.4. Comments on the Topology of the Sulfur Subnetwork

We notice that adding the artificial constraints on sulfur-based species results in the rate of the reaction between H and OCS being constrained. This is interesting, as the abundance of the product, HOCS, is not constrained. Instead it is the penultimate molecule in the reaction chain, OCS, that is constrained. This suggests that it is not necessary for the end of


a reaction chain to be constrained to constrain the reaction rates, as was observed with methanol in Section 5.2. It seems that having a constraint on the penultimate molecule is sufficient. Constraining an earlier molecule would not do the trick, as was demonstrated when methanol’s constraint was removed but CO’s was kept. There is a notion of distance that needs to be considered. However, this would need to all be reconsidered for the case where there is more than one depletion mechanism for OCS. It is likely that having two depletion mechanisms, each of whose end product is unconstrained, would have a different effect, as there will be uncertainty about the branching ratio of each depletion route.

## 7. Conclusions

In this work, we have proposed new methods for performing Bayesian inference on chemical reaction networks that have very few constraints. We find that reducing the reaction network to just the reaction chains whose ends are constrained allows us to greatly reduce the computational expense. Despite the simplification, our most likely reaction rate values are mostly unchanged. We also find that we can separate chemical reaction networks into subnetworks, which can be analyzed in parallel. We believe that such network reduction techniques will be of great use when looking at grain surface reaction networks, where there are few constraints on the molecules. However, it should be noted that the results of such a simplified chemical model can only provide a qualitative understanding of the chemistry.

We briefly summarize some general observations we have made that might prove useful in reducing reaction networks for Bayesian inference:

1. Reducing the network reduces the computational expense of the inference process. The time taken for our inference runs scales roughly linearly with the number of reactions. However, the network connectivity is also likely to be a significant factor. This warrants further investigation.
2. Reducing the network comes at the cost of artificially changing the variance of the posteriors. However, this variance can be partially recovered by adding a dummy reaction where the product is unconstrained, though there is the risk that the joint posterior distribution becomes more unrealistic. This needs to be investigated in further depth.
3. When considering a reaction chain, it is important to include constraints for the final or penultimate molecule



**Figure 7.** Plots of the obtained posteriors for Configurations 6–8 when the CO abundance is reduced by a factor of  $10^4$ . We observe that for this case, where the CO constraint is not several orders of magnitude greater than the abundances of the molecules in the sulfur subnetwork, that it does not matter whether we include the CO abundance or not. Both cases allow us to recover the reaction rate with a very small bias.

produced, as this ensures (for the case of a linear reaction chain) that the reaction rates of intermediate reactions are constrained. This might provide a general idea for future observations in terms of which molecules in the ices to look for. However, for more complex reaction networks, the intermediate reactions may play a more significant role.

4. A network can be “separated” into subnetworks. This is a potentially very useful tool, especially when looking at the grain surface chemistry of COMs, where the networks themselves are very large. For example, Garrod (2013) provides a potential surface reaction network with around 200 reactions. In principle, this network could be split up into smaller subnetworks and Bayesian inference could then be performed on each subnetwork in parallel. A potential general strategy would be to perform the network splitting at the point in the network with the highest network connectivity. For the network considered in this paper, this was the CO molecule. By making appropriate arguments about the relative magnitudes of the subnetworks and placing appropriate cuts in the networks, one could repeat the procedure as above. In order to decide what to do with a constrained molecule that is shared by the two subnetworks, one can make arguments about the relative abundances of the molecules in each subnetwork. There are two cases to consider:
5. For the case where the shared constrained molecule has a significantly larger abundance than the molecules in one of the subnetworks, one can include its constraint in the higher-abundance subnetwork.
6. For the case where the shared constrained molecule is roughly the same as the abundances in either of the two subnetworks, one can choose to include it in either or both subnetworks.

For the case of a more interconnected networks with linear components (see Figure 11 in Linnartz et al. 2015 for an example of such a network), it makes sense to separate out the linear reaction chains first, as we found in Section 5.3 that their topology is easy to understand intuitively. By removing these linear reaction chains, one could treat the interconnected subnetworks separately. Depending on their topology, one can employ the separation strategies discussed in this work.

Further work will need to focus on recovering the posterior distributions better. A quantitative approach is needed to better compare the posteriors inferred with full networks and subnetworks as well as explore how to use “dummy reactions” of the form  $X + \text{CO} \rightarrow \text{XCO}$  to recover the variance for the case where the reaction replaces a subnetwork with constraints.

Future work will also need to look at reaction networks with more complex geometries. The example considered here is fairly simple, with a relatively low degree of connectivity. As the complexity of the reaction networks considered increases, there will need to be more well-defined notions of how the position of a constraint influences the inference of related reaction rates. The guidelines we have presented are applicable to simple networks. Investigation of more interconnected networks is the focus of ongoing work. We aim to come up with a set of criteria to determine how to best separate more complex networks.

The authors thank the referees for constructive comments that greatly improved this work. J.H. was supported by the STFC UCL Centre for Doctoral Training in Data Intensive Science (grant number ST/P006736/1). This work was also supported by European Research Council (ERC) Advanced Grant MOPPEX 833460. S.V. acknowledges support from the European Unions Horizon 2020 research and innovation program under the Marie Skłodowska-Curie grant agreement

No. 811312 for the project ‘‘Astro-Chemical Origins (ACO). S. M.F. is supported by the Royal Society. This work used computing equipment funded by the Research Capital Investment Fund (RCIF) provided by UKRI, and partially funded by the UCL Cosmoparticle Initiative.

*Software:* UCLCHEM (Holdship et al. 2017), PyMC3 (Salvatier et al. 2016), F2PY (Peterson 2009).

## Appendix A Convergence

Any MCMC chain needs to be checked for convergence. In the limit of an infinitely long chain, the sampled posterior distribution can be said to approximate the true posterior. However, when dealing with a finitely long chain, one needs to check that the posterior is not changing by much. We made use of two diagnostics to check. However, it should be emphasized that these diagnostics do not guarantee that the chains are converged. Rather, if these checks fail, then we know the chain has not converged. Satisfying the conditions of the diagnostics simply lends credence to the hypothesis that the chains have converged.

### A.1. Geweke Diagnostic

The Geweke diagnostic calculates a  $z$ -score between two sections of a chain, typically the first 10% and the final 50% (Geweke 1991; Roy 2020). The  $z$ -score is calculated by

$$z = \frac{\bar{\theta}_a - \bar{\theta}_b}{\sqrt{\sigma_a^2 + \sigma_b^2}}, \quad (\text{A1})$$

where the quantities with subscript  $a$  refer to the first 10% of the chain and the quantities with subscript  $b$  refer to the final 50%. In the limit of the chain length going to infinity, the Geweke diagnostic is expected to follow a normal distribution (Cowles & Carlin 1996). In Figure 8, we plot the distribution of the Geweke diagnostic for the chains of configuration 1 along with a normal distribution overlaid. We observe that the vast majority of points in the diagnostic stay within one standard deviation of the mean.


### A.2. Gelman–Rubin

The Gelman–Rubin diagnostic provides a means of comparing the variance across all chains with the variance of the individual chains (Gelman & Rubin 1992; Hogg & Foreman-Mackey 2018). There are several quantities of interest at play here. They are calculated for each scalar parameter of interest separately (Gelman & Rubin 1992; Gelman & Shirley 2011).

For  $m$  chains, assumed to be of length  $n$ , the between-chain variance is defined as

$$B = \frac{n}{m-1} \sum_{i=1}^m (\hat{\theta}_i - \hat{\theta})^2, \quad (\text{A2})$$

where  $\hat{\theta}_i$  is the estimator of the mean for chain  $i$  and  $\hat{\theta}$  is the estimator of the mean of the sample. The latter is simply the average of all chain mean estimators.



**Figure 8.** A plot of the distribution of the Geweke diagnostic for the samples obtained in Configuration 6. We find that most of the points are within two  $z$ -scores of the mean. A normal distribution with zero mean is overlaid to show that the Geweke diagnostic is approaching a normal distribution.

The within-chain variance,  $W$ , is defined as the average of the variances of all chains

$$W = \frac{1}{m} \sum_{i=1}^m \hat{\sigma}_i^2, \quad (\text{A3})$$

where  $\hat{\sigma}_i^2$  is the estimator of the variance for chain  $i$ .

The pooled variance estimate is defined as

$$\hat{V} = \frac{n-1}{n} W + \frac{B}{n}. \quad (\text{A4})$$


The quantity of interest is

$$\hat{R} = \frac{\hat{V}}{W}, \quad (\text{A5})$$

which is referred to as the potential scale reduction factor (PSRF). In the limit of infinitely long chains,  $\hat{R}$  tends to 1 from above. The closer  $\hat{R}$  is to 1, the better. In practice, a cutoff is used to determine convergence, typically 1.1 though there is some debate surrounding this (Vats & Knudson 2018; Roy 2020). The PSRF is related to the autocorrelation time that was used in H18, in that a value of  $\hat{R}$  that satisfies the criterion corresponds to a chain length greater than the autocorrelation time. There exists some debate as to whether to discard the first half of the chain when evaluating  $\hat{R}$  (see Roy 2020 and Gelman & Shirley 2011 for two opposing views on the matter). Regardless of whether we do this, we find  $\hat{R} \leq 1.09$  for the reaction rates of the nonconstrained reactions and  $\hat{R} \leq 1.03$  for the constrained reactions.

## Appendix B Bayesian Sensitivity Analysis


Bayesian parameter inference depends on two fundamental quantities: the likelihood and the prior. It is important to understand the relative information content of these two quantities to determine whether one’s conclusions are driven by one’s initial beliefs or the data at hand. In this work, we have used a log-uniform prior in the reaction rates,  $k$ , encoding our ignorance of their order of magnitude (Jeffreys 1946). One



**Figure 9.** Posterior probability distributions for Configuration 1 using three different priors. Alongside the uniform prior in  $y = \log(k)$ , we also consider uniform priors in the variables  $t = \frac{1}{\log(k)}$  and  $u = k$ .

of the key aspects when performing Bayesian inference is to determine if the posterior distributions are driven by the data or by the prior. This was considered in detail by Fischer (2019) in the context of chemical kinetics, but has also been discussed with regards to climate models in Tomassini et al. (2007). In this work, we motivated our choice of prior using chemical considerations and therefore used a log-uniform prior in the reaction rates  $k$  (Ioppolo et al. 2011; Chuang et al. 2016). Fischer (2019) argues that if we are completely ignorant about  $k$ , then we are also ignorant in  $f(k)$ , where  $f$  is an arbitrary function in  $k$ . We could equivalently use a prior that is uniform in  $f(k)$ . If the posterior distributions differ depending on the choice of prior, this means that the posteriors are “prior-driven,” as opposed to “data-driven.”

In this appendix, we repeat a portion of our analysis using two alternative priors in order to demonstrate the impact of different prior assumptions. Figure 9 shows the posterior distributions for Configuration 1 using three different priors. Alongside the uniform prior in  $y = \log(k)$  that we have used throughout this paper, we also use uniform priors in  $t = \frac{1}{\log(k)}$  and  $u = k$ . We observe that the posteriors can differ significantly depending on the prior. However, it is interesting to note that the posteriors for reactions 1 and 2 are relatively similar when uniform priors in  $y$  and  $t$  are used. As was discussed in Section 5.1, this is due to the fact that these reactions are related to the formation of  $\text{H}_2\text{O}$ , the abundance of which is known to differ from zero at the  $3\sigma$  level. This appears



**Figure 10.** Posterior probability distributions for Configuration 2 using three different priors. Alongside the uniform prior in  $y = \log(k)$ , we also consider uniform priors in the variables  $t = \frac{1}{\log(k)}$  and  $u = k$ .

to suggest that the marginal probability distributions for Reactions 1 and 2 are more “data-driven” than the posteriors for the other reactions. The other reactions are associated with the weaker abundance constraints. Their associated posterior distributions do not rule out much of the parameter space, so it is unsurprising that the posteriors are affected by the priors. We observe similar trends for the posterior distribution for Configuration 2, shown in Figure 10. We have cropped the ordinate at 1 in order to make sure the posteriors that came from using uniform priors in  $y$  and  $t$  are more visible.

#### ORCID iDs

Johannes Heyl <https://orcid.org/0000-0003-0567-8796>  
 Serena Viti <https://orcid.org/0000-0001-8504-8844>

Jonathan Holdship <https://orcid.org/0000-0003-4025-1552>  
 Stephen M. Feeney <https://orcid.org/0000-0003-2268-2519>

#### References

- Ayilaran, A., Hanicinec, M., Mohr, S., & Tennyson, J. 2019, *PIST*, **21**, 064006
- Boogert, A. A., Gerakines, P. A., & Whittet, D. C. 2015, *ARA&A*, **53**, 541
- Brewer, B. J., & Foreman-Mackey, D. 2018, *J. Stat. Software*, **86**, 1
- Chuang, K. J., Fedoseev, G., Ioppolo, S., van Dishoeck, E. F., & Linnartz, H. 2016, *MNRAS*, **455**, 1702
- Cowles, M. K., & Carlin, B. P. 1996, *J. Am. Stat. Assoc.*, **91**, 883
- de Mijolla, D., Viti, S., Holdship, J., Manolopoulou, I., & Yates, J. 2019, *A&A*, **630**, A117
- Fischer, M. 2019, in *ISIPTA 2019—Int. Symp. on Imprecise Probabilities Theories and Applications*, Vol. 103, 203, <https://hal.archives-ouvertes.fr/hal-02292325>

- Fraser, H., Williams, D., Sims, I., Richards, A., & Yates, J. 2003, *A&G*, **44**, 4.29
- Galagali, N., & Marzouk, Y. M. 2019, *J. R. Soc. Interface*, **16**, 20180766
- Garrod, R. T. 2013, *ApJ*, **765**, 60
- Gelman, A., & Shirley, K. E. 2011, *Handbook of Markov Chain Monte Carlo* (London: Chapman and Hall/CRC), 189
- Gelman, A., & Rubin, D. B. 1992, *StaSc*, **7**, 457
- Geweke, J. F. 1991, *Evaluating the Accuracy of Sampling-based Approaches to the Calculation of Posterior Moments*, Tech. Rep. 148
- Grassi, T., Bovino, S., Gianturco, F. A., Baiocchi, P., & Merlin, E. 2012, *MNRAS*, **425**, 1332
- Herbst, E., & van Dishoeck, E. F. 2009, *ARA&A*, **47**, 427
- Hocuk, S., Cazaux, S., & Spaans, M. 2014, *MNRAS*, **438**, L56
- Hogg, D. W., & Foreman-Mackey, D. 2018, *ApJS*, **236**, 11
- Holdship, J., Jeffrey, N., Makrymallis, A., Viti, S., & Yates, J. 2018, *ApJ*, **866**, 116
- Holdship, J., Viti, S., Jiménez-Serra, I., Makrymallis, A., & Priestley, F. 2017, *AJ*, **154**, 38
- Huijser, D., Goodman, J., & Brewer, B. J. 2015, arXiv:1509.02230
- Ioppolo, S., van Boheemen, Y., Cuppen, H. M., van Dishoeck, E. F., & Linnartz, H. 2011, *MNRAS*, **413**, 2281
- Jeffreys, H. 1946, *RSPSA*, **186**, 453
- Jiménez-Escobar, A., Muñóz Caro, G. M., & Chen, Y.-J. 2014, *MNRAS*, **443**, 343
- Laas, J. C., & Caselli, P. 2019, *A&A*, **624**, A108
- Linnartz, H., Ioppolo, S., & Fedoseev, G. 2015, *Int. Rev. Phys. Chem.*, **34**, 205
- Makrymallis, A., & Viti, S. 2014, *ApJ*, **794**, 45
- Norton, J. D. 2008, *Philos. Sci.*, **75**, 45
- Peterson, P. 2009, *Int. J. Comput. Sci. Eng.*, **4**, 296
- Roy, V. 2020, *AnRSA*, **7**, 387
- Salvatier, J., Wiecki, T. V., & Fonnesbeck, C. 2016, *PeerJ Comput. Sci.*, **2**, e55
- Tomassini, L., Reichert, P., Knutti, R., Stocker, T. F., & Borsuk, M. E. 2007, *JCli*, **20**, 1239
- van Dishoeck, E. F. 2014, *FaDi*, **168**, 9
- Vats, D., & Knudson, C. 2018, arXiv:1812.09384
- Vidal, T. H. G., & Wakelam, V. 2018, *MNRAS*, **474**, 5575
- Walley, P. 2000, *Int. J. Approximate Reasoning*, **24**, 125
- Woods, P. M., Occhiogrosso, A., Viti, S., et al. 2015, *MNRAS*, **450**, 1256
- Xu, R., Bai, X.-N., Öberg, K., & Zhang, H. 2019, *ApJ*, **872**, 107