UNIVERSITY COLLEGE LONDON

Faculty of Mathematics and Physical Sciences

Department of Physics & Astronomy

CHEMICAL AND STATISTICAL MODELS OF THE INTERSTELLAR MEDIUM AND STAR–FORMING REGIONS

Thesis submitted for the Degree of Doctor of Philosophy
of the University of London

by

Richard Peter Rollins

Supervisors:
Prof. Jonathan M. C. Rawlings
Prof. Serena Viti

Examiners:
Prof. Michael J. Barlow
Dr Claudia J. Cyganowski

July 15, 2015

To my father.

I, Richard Peter Rollins, confirm that the work presented in this thesis is my own. Where information has been derived from other sources, I confirm that this has been indicated in the thesis.

Abstract

Astrochemistry is a field currently in a data–rich era, but observations alone cannot provide a complete description of the objects that we see; sophisticated modelling techniques are needed to extract the maximum astrophysical understanding from the data. Chemical models of interstellar gas and ices are continually being developed to provide more accurate descriptions of the chemical composition of our universe. Furthermore, Bayesian inference and statistical learning methods are starting to be incorporated into the analysis of both observational data and astrochemical models. These are the motivations for this thesis.

We develop a Markov chain Monte Carlo Bayesian statistics code based on nested sampling for parameter estimation and model comparison with optimal parallel performance. We assess its performance for a set of test problems in terms of accuracy, reliability, time–to–solution and sampling efficiency on two of the country's top supercomputers. The software is applied to a chemical model for diffuse clouds in the interstellar medium. Properties of the gas such as its density and temperature are estimated from chemical observations in a statistically robust manner which enables us to comment on the importance of various chemical mechanisms such as photochemistry and grain–surface chemistry in controlling the chemical composition.

We also present a study of the effects of mutual shielding by atomic carbon in promoting nitrogen chemistry. By solving the radiative transfer we are able to show that the carbon continuum can have a chemically significant effect on the photoionization and photodissociation rates of a number of species including molecular nitrogen and CN. Finally, we develop a chemical model for carbon monoxide and HCO^+ in protostellar outflows where hot outflowing material mixes turbulently and reacts with the cold molecular envelope, allowing us to address an apparent bias towards outflows with large opening angles.

Acknowledgements

There are many individuals to whom I owe thanks for their contributions both to this work and to my time at University College London. To my supervisor Jonathan for bringing me from an astrochemistry amateur, through conference talks, published papers and now this thesis, to where I am today. To Serena for her guidance and wisdom through my time as a student. To Filipe for helping to broaden my academic horizons, not to mention some sensational barbecue. To Sree for his mentorship, teaching me what feels like everything I now know about computers. To Jeremy and the DiRAC consortium for making a significant portion of my research possible and offering me the chance to join them. To my family for their endless love and support, especially through hard times. To all the friends I made along the way who offered welcome distractions and made the journey that much more enjoyable. And finally to the STFC for my studentship, a significant portion of which was spent over several nights at the Euston Tap supporting the Lagunitas Brewing Company.

Before, I wandered as a diversion. Now I wander seriously...

W. Percy

Contents

List of Figures

List of Tables

Chapter 1

Introduction

Space is far from the empty void it is typically held to be in popular understanding. Between the stars, planets and other celestial bodies, our galaxy is occupied by gas, dust grains, photons and cosmic rays, all of fundamental astrophysical importance. While pressures can be orders of magnitudes less than any artificial vacuum on Earth, on galactic time–scales a range of processes can drive chemical reactions leading to a diverse gas of atoms, ions and molecules as well as ices covering grain surfaces. Astrochemistry, the multidisciplinary study of those extra–terrestrial chemical species and processes, is a young field of research currently thriving in a data–rich era with the latest generations of telescopes such as the *Herschel Space Observatory* (Pilbratt et al. 2010) and the *Atacama Large Millimeter Array* (ALMA, Wootten & Thompson 2009) providing a wealth of sensitive, high–resolution images and spectra. Understanding the chemical complexity of interstellar gas is of importance to understanding a range of phenomena in the Milky Way and other galaxies. For example, it plays a key role in physical processes such as heating and cooling (Goldsmith & Langer 1978, Richings et al. 2014*a*), radiative transfer (Boisse 1990, Efstathiou et al. 2000, Offner et al. 2009) and the optical properties of dust (Whittet et al. 2001). In cold, dense clouds the chemical state of the gas is coupled to the process of gravitational collapse and structure formation (Hartquist et al. 1993), while highly energetic processes such as shocks and turbulence can lead to chemical enrichments (Lesaffre et al. 2013, Godard et al. 2014). Astrochemistry is also of importance in its own right, from understanding the delivery of water to the Earth's oceans by asteroids or comets (Morbidelli

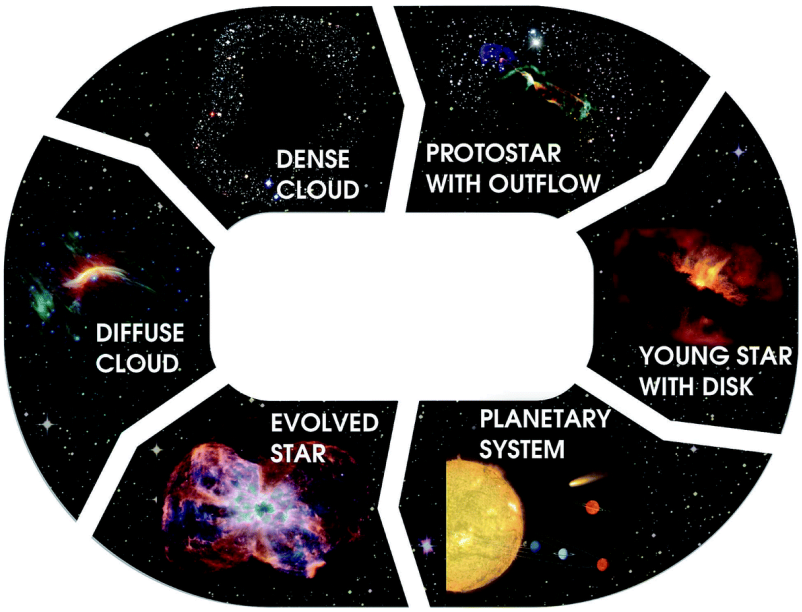

Figure 1.1. Schematic depiction of the cycling of gas and dust between various phases and structures in the interstellar medium. Image credit: M. Persson, NASA, ESA, ESO, ALMA.

et al. 2000, Hartogh et al. 2011, Altwegg et al. 2015) to the formation of complex prebiotic molecules that are the precursors of life on Earth (Jørgensen et al. 2012, Majumdar et al. 2012).

Observations of molecular rotational transitions in the radio, submillimetre and far–infrared wavelengths as well as vibrational transitions in the mid–infrared allow us to make inferences about the physical state of the gas such as its density and temperature. However, these observations face other fundamental limitations, such as our ability to only observe spatial distributions in the plane of the sky and radial velocities, when in fact the objects we observe live in a six dimensional phase space of position and velocity. Making inferences from this plentiful yet incomplete data relies on sophisticated modelling techniques such as data reduction, chemical modelling, dynamical simulations, radiative transfer and statistical methods to explore what are still largely ill–understood parameter spaces. This section is intended to serve as an introduction to these and other observational and theoretical concepts important to the study of the astrochemistry of our galaxy before considering the scope of this thesis.

Table 1.1. Typical densities and temperatures for the thermal equilibrium (top) and non–equilibrium (bottom) phases of gas in the interstellar medium. Adapted from Ferrière (2001) and Williams & Viti (2014).

Component	$n_{\rm H}\,/\,{\rm cm}^{-3}$	$T\,/\,{\rm K}$
Hot Ionized Gas	10^{-3}	10^{6}
Warm Ionized / Neutral Gas	10^{-1}	10^{4}
Cold Neutral Gas	10^{1}	10^{2}
HII Regions	$10^{2} - 10^{6}$	10^{4}
Molecular Clouds	10^{4}	10

1.1 Interstellar Gas

Interstellar gas is continually cycled between many different states throughout the galaxy by the processes of gravitational collapse and stellar birth and death as depicted in Figure 1.1. The large range of densities and temperatures of some of those states is demonstrated in Table 1.1. Field et al. (1969) were the first to demonstrate that gas heated by cosmic rays (energetic particles traveling through interstellar space at close to the speed of light) and cooled by atomic fine structure lines split into two phases in pressure equilibrium. McKee & Ostriker (1977) expanded on the description to include hot, rarefied gas from the ejecta of supernovae to give a three phase model for the interstellar medium including the cold neutral medium (CNM), warm neutral medium (WNM) and hot ionized medium (HIM). While the HIM occupies the majority of the volume of interstellar space, the bulk of the gas is in the CNM, the so called diffuse gas clouds of the interstellar medium. Some of the gas in the WNM can also be photoionized by the presence of massive stars. In addition, there are structures in the interstellar medium that are not in thermal pressure equilibrium such as HII regions, molecular clouds and protostars, all of which harbor interesting physics.

The transition from diffuse gas to molecular clouds is mechanical, relying on the accumulation of gas swept up by stellar winds (Hollenbach et al. 1976) or radiation pressure from massive stars (Koenig et al. 2008) into more compact clouds. These new clouds are gravitationally bound and the raised density encourages the formation of and cooling by molecular hydrogen and other species. These new molecular clouds have very inhomogeneous structures including dense cores that may remain stable and "starless" due to the support of turbulent or magnetic pressure (Tan et al. 2013) while others may become gravitationally unstable as they exceed the Jeans mass leading to isothermal collapse on the

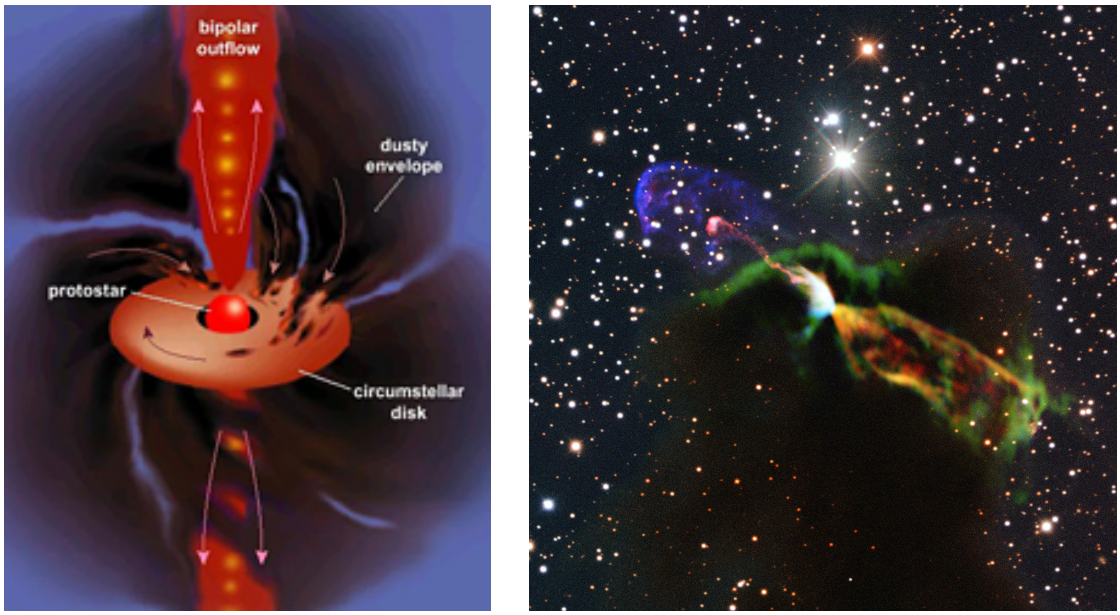

Figure 1.2. Left: Morphology diagram for a Class I protostellar object embedded in a dusty envelope, demonstrating a central source fed by accretion from a circumstellar disc and ejecting material as a collimated bipolar outflow. Image credit: Greene (2001). Right: Colour composite image of HH46/47 Class I protostellar object observed in the CO(1–0) rotational transition (ALMA; red, orange, green) and in the I, SII, V and B infrared and optical bands (*New Technology Telescope*; pink and blue). Image Credit: ESO, ALMA, Arce et al. (2013), Bo Reipurth.

gravitational free–fall time–scale to form a protosteller core. This first stage of collapse was initially modelled as being spherically symmetric (Shu 1977) but new evidence of filamentary structures within molecular clouds (André et al. 2010, Molinari et al. 2010) leads to a picture of anisotropic inflow (Seifried et al. 2015).

The case of low–mass ($\lesssim 8\,M_{\odot}$) star formation from protostellar cores proceeds through a series of well understood stages. Initially, the density of the central object rises to the point that energy from the gravitational collapse cannot be released through optically thick cooling lines causing the core to heat adiabatically while it collapses. The halt of contraction of the central source due to thermal pressure and the formation of a hydrostatic core defines Class 0 objects for which the core mass is less than the envelope mass and energetic bipolar outflows may be present. In Class I objects (see Figure 1.2) much of the gas envelope has been excavated by an outflow and the remainder forms an extended disc. By Class II accretion and outflow processes have mostly stopped and the protoplanetary disc is fully formed as in the classic T Tauri type star while in Class III radiation pressure,

among other processes, has cleared much of the gas from the disc while remaining dust forms a debris disc that is the precursor to the formation of planets. More rigorously, the classes are distinguished by their spectral indices in the near– and mid–infrared (Lada 1987). High–mass star formation is distinguished from the low–mass case in that the Kelvin–Helmholtz time–scale to radiate away the energy of gravitational collapse is shorter than the accretion time–scale. As such, high–mass stars reach thermal equilibrium on the zero–age main sequence and begin nuclear fusion while accretion is still ongoing leading to complex dynamics between the accretion flow and radiation pressure which is understood less clearly.

When stars die, they can return matter to their local interstellar environment by one of a few energetic processes dependent on their mass. Low–mass stars may evolve into red giant stars or asymptotic–giant–branch stars with strong winds while massive stars may undergo core–collapse followed by a violent supernova explosion. Both of these processes return stellar material to the interstellar medium as well as potentially sweeping up diffuse gas to begin the process of forming molecular clouds again. Diffuse clouds in the cold neutral medium therefore represent a stockpile of gas recycled from the death of old stars available to be converted into new stars and are of significant importance in understanding the evolution of galaxies right down to the formation of our own planet.

1.2 Interstellar Chemistry

1.2.1 Gas–Phase Reactions

The simplest chemical reactions that can take place in interstellar gas are two body reactions between neutral species. In many cases, they tend to be rather inefficient due to small cross–sections, the need for high densities to promote collisions and in some cases thermal activation barriers demanding high temperatures. However, certain radical species can be reactive even in molecular clouds and dense cores (Herbst et al. 1994, Smith et al. 2004). In very rare cases three body reactions are also possible, such as in the gas around evaporating ice mantles (Cecchi-Pestellini et al. 2010, Rawlings et al. 2013*b*). In general though, two–body reactions between ionic and neutral species are much faster and play a larger role in the chemistry of interstellar gas. The cross–sections for these ion–neutral reactions are enhanced due to the induction of a temporary electric dipole in the neutral species by the ion leading to attractive forces between the two. H_3^+ is perhaps the most

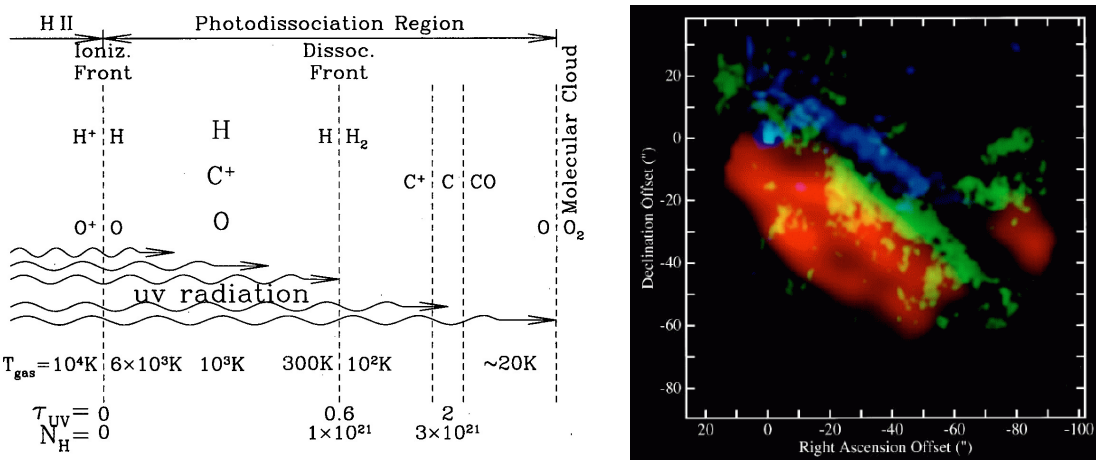

Figure 1.3. Left: Schematic depth–dependent chemical transitions from HI → H_2 and $C^+ \rightarrow C \rightarrow CO$ with increasing extinction A_V in one dimensional photon dominated region. Image credit: Draine (2011). Right: Stratified PAH 3.3 μm (blue), H_2 rovibrational (green) and CO(1-0) rotational (red) emission in the Orion Bar PDR due to the bright star θ^1 Ori C to the north–west. Image credit: Tielens et al. (1993).

important ion in dark interstellar clouds since it can donate protons and hence its charge to neutral species and thereby initiate further ion–neutral reactions to more complex species.

In order for ion–neutral reactions to be viable in interstellar gas a source of ionization is needed with the most obvious candidate being the ultraviolet (UV) photons emitted by massive stars. Regions of interstellar space where the chemistry is driven primarily by photoionization and photodissociation reactions are known as photon dominated regions or photodissociation regions (PDRs). While gas in the WNM is too diffuse to initiate chemical reactions and the many magnitudes of dust extinction mean UV photons cannot penetrate dark molecular clouds, diffuse gas, the surfaces of dense clouds and the interior edges of ionized nebulae are commonly described and modelled as PDRs. Photons with wavelengths shortwards of 912 Å tend not to be present in the interstellar radiation field as they maintain the ionization balance of atomic hydrogen in the HIM and WNM. As such, atomic hydrogen, oxygen and nitrogen are not photoionized in PDRs since their ionization potentials lie at or below 912 Å. However, many other species can be photoionized or photodissociated at the edge of PDRs including atomic carbon, H_2 and CO. As the attenuation of the UV radiation field by dust increases with depth into the gas, photorates fall and those species become more abundant, leading to a stratified chemical structure (see Figure 1.3). More complex chemistry in PDRs is a fine balance between sufficient

ionization to initiate ion–neutral reactions and sufficient attenuation by dust to prevent photodissociation of relevant molecules. In high–extinction regions where there is no incident flux of UV photons, ion–neutral reactions can still take place due to cosmic rays that readily ionize molecular hydrogen, initiating a range of chemistries as described in Section 1.2.2 below. In addition, a small but significant flux of UV photons is still present via the mechanism detailed by Prasad & Tarafdar (1983) whereby the electrons released in cosmic ray ionization events excite hydrogen molecules that then emit as they return to the ground state.

1.2.2 Molecular Hydrogen

Hydrogen is the most abundant element in the universe and similarly molecular hydrogen is the most abundant molecule. It is well established that star formation on galactic scales is driven by the availability of gas (Schmidt 1959, Kennicutt 1998). H_2 therefore represents the most significant reservoir of raw materials for star formation making it probably the most important molecule in interstellar gas. The formation of molecular hydrogen by two body reactions in the gas phase is prohibitively slow since most collisions between hydrogen atoms are elastic. However, hydrogen atoms can bind to dust grains where they are incredibly mobile, allowing them to come together and react to form H_2 before being desorbed back into the gas phase (see Section 1.2.4).

Other species depend critically on the availability of molecular hydrogen to initiate their gas–phase chemistry. In photon dominated regions, high fluxes of UV photons ionize atomic carbon which then undergoes successive fast ion–neutral reactions with molecular hydrogen (see Figure 1.4), each time gaining one hydrogen atom until the CH_5^+ ion is formed. One final reaction with molecular hydrogen sees the ion loose a proton to form the saturated species methane from which more complex molecules can be formed. The chemistry of atomic nitrogen is similarly initiated when it reacts with UV–pumped molecular hydrogen to form NH which can be photoionized and undergo further fast ion–neutral hydrogenation reactions to ammonia similar to the case of carbon. While oxygen tends not to be photoionized since its ionization potential is greater than that of the far more abundant atomic hydrogen, in even the darkest clouds molecular hydrogen is still ionized by penetrating cosmic rays and then reacts with further hydrogen to yield H_3^+. The trihydrogen cation is responsible for initiating the fast ion–neutral chemistry of atomic oxygen (as well as carbon and nitrogen in dark clouds) by donating a proton to form OH^+ from

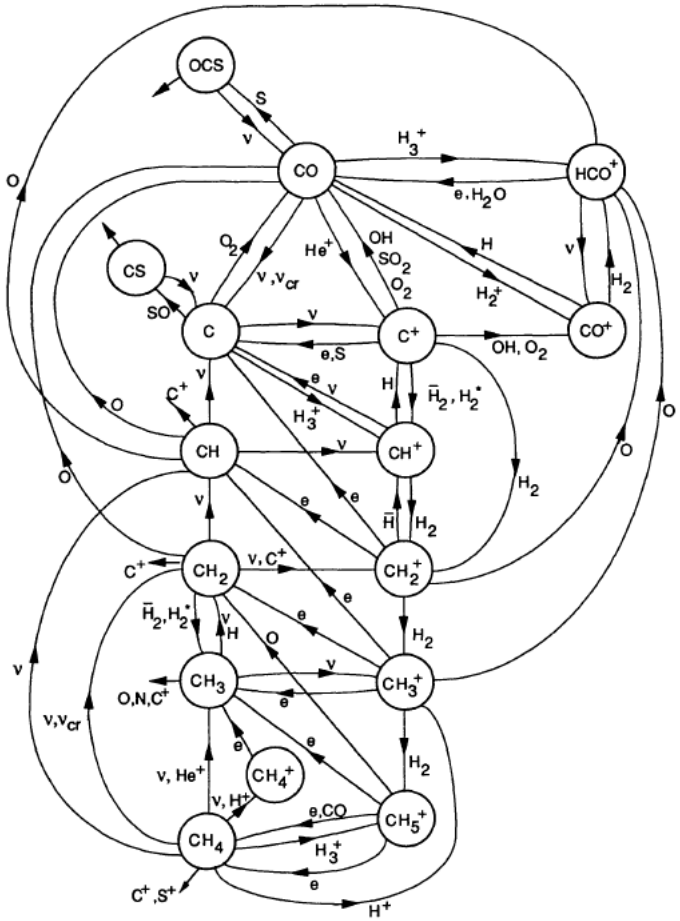

Figure 1.4. A partial chemical network tracing the gas–phase carbon chemistry in photon dominated regions through reactions with hydrogen, oxygen, electrons, photons and cosmic rays. Image credit: Sternberg & Dalgarno (1995).

which successive hydrogenation reactions again follow as in the case of carbon and nitrogen up to the formation of water.

Unfortunately, the lack of a permanent dipole for radiative rotational transitions makes molecular hydrogen incredibly difficult to observe. Although in hotter gas it can be an efficient coolant through rovibrational transitions, we require other tracers in order to be able to observe and study the cold, dense gas that is the precursor to star formation.

1.2.3 Carbon Monoxide

Despite being unstable in terrestrial conditions, carbon monoxide is the second most abundant molecule in interstellar space, constituting the majority of gaseous interstellar carbon in certain environments. Once carbon chemistry has been initiated, the CH radical reacts with oxygen to form the molecular ion HCO^+ which then recombines dissociatively with

electrons to form CO. Unlike H_2, CO has a permanent dipole moment which combined with its high abundance makes it readily observable. The assumption that all carbon is locked up in carbon monoxide means that the CO abundance should be fixed by the elemental abundance of carbon relative to hydrogen. This suggests that observed column densities of carbon monoxide trace unobservable molecular hydrogen (and hence the total gas content) whose columns can be inferred using a conversion factor from CO, $X_{\rm CO}$, although many studies show that $X_{\rm CO}$ is far from constant (see Bolatto et al. 2013 for a review). In many cases, CO is still our best tracer of interstellar gas and research into the $X_{\rm CO}$ factor is ongoing.

Carbon monoxide plays an important role in the star formation process as a strong coolant through its low–J rotational line transitions that allow the gas temperature to be bought as low as 10 Kelvin. While other molecular rotations (water) and atomic fine structure transitions ([CII]) can give similar levels of cooling, the abundance of CO tends to make it the dominant coolant in prestellar cores. As the temperature in the core falls the gas contracts without thermal support until it exceeds the Jean's mass and begins the runaway gravitational collapse that initiates star formation.

1.2.4 Interstellar Dust

Interstellar gas is also composed of a significant number of solid particles nucleated in the outflows from supernovae explosions and the winds of red giant and asymptotic–giant–branch stars. These "dust grains" range in size from large conglomerations of carbonaceous and silicate species not dissimilar to terrestrial soot and sand down to atomic carbon macromolecules known as polycyclic aromatic hydrocarbons (PAHs). The distribution of the sizes a of grains was found by Mathis et al. (1977) to be a power law of the form $n(a) \propto a^{-3.5}$ from a fit to the interstellar extinction curve. While gas to dust ratios of 100 are typically assumed in the interstellar medium (Savage & Mathis 1979) making dust another important tracer of interstellar hydrogen, this conversion factor is again somewhat variable. Dust is important in the interstellar medium for its optical properties, absorbing photons at optical and UV wavelengths and re–emitting the energy in the infrared, as seen in Figure 1.5. It is therefore also an important source of opacity, regulating photochemical reactions from destroying molecules in dense cores beyond the PDR.

As gas temperatures fall below approximately 100 K, molecules from the gas phase start to stick to the surfaces of dust grains by either van der Waals interactions (physisorption) or

Figure 1.5. The optically dark molecular universe revealed by submillimetre observations. Left: Optical image of the Barnard 211 and Barnard 213 regions of the Taurus molecular cloud complex obscuring light from background stars. Right: Superimposed submillimetre continuum observations (orange) showing thermal emission from cold dust grains tracing the dense molecular gas component in the filamentary structure. Image credit: ESO, APEX (MPIfR, ESO, OSO), A. Hacar et al., Digitized Sky Survey 2. Acknowledgment: Davide De Martin.

the formation of chemical bonds (chemisorption). Not only are these processes important because they alter the chemical composition of the gas phase, but they also bring together atoms and molecules in the ice mantle where they can react. Dust grains therefore act as interstellar catalysts for chemical reactions. There are three main mechanisms by which this can occur: Langmuir–Hinshelwood where two species on the grain surface diffuse towards each other to react, Eley–Rideal where a gas–phase species collides directly with a species in the ice mantle and Harris–Kasemo where a gas–phase species is adsorbed but undergoes fast diffusion to react with a more distant adsorbate before thermalizing to the dust temperature. Atomic hydrogen is also the most mobile of mantle species due to its low mass which along with its high abundance means that the majority of grain surface reactions are hydrogenations. This implicitly covers the formation of molecular hydrogen from two hydrogen atoms on grain surfaces which is far more rapid than any gas–phase route to the molecule. In the case of H_2, the exothermicity of the chemical bond creation is enough to break the bond between the molecule and the dust grain, releasing it straight into the gas phase where it can initiate more complex chemistry. Other hydrogenated species

common on grain mantles include methane, ammonia, water and methanol. However, a number of non-hydrogenated species are also abundant in interstellar ices including carbon monoxide and carbon dioxide. Depleted species are released back into the gas phase by either thermal heating of the ice mantle allowing species to break their bond with the grain surface or by a range of non–thermal desorption processes including destruction of the bond by UV photons, localized heating by the formation of molecular hydrogen and heating of the whole grain upon impact by a cosmic ray.

1.3 Chemical Modelling

1.3.1 Solving the Chemical Rate Equations

The process of chemical modelling involves the calculation of the densities n_{i} for all chemical species i belonging to the set of species $\mathcal{S}$ to be modelled in the system. The densities are often calculated relative to n_{H}, the total density of hydrogen nucleons, giving the relative chemical abundances $Y_{\mathrm{i}} = n_{\mathrm{i}}/n_{\mathrm{H}}$. On a microscopic level, the densities of species change when individual atoms and molecules take part in chemical reactions, either as reactants or products. We model those changes macroscopically by considering the time rate of change of the density $\dot{n}_{\mathrm{i}}$ for a given species i as a function of the rate of all of the chemical reactions in which it is involved:

$$\dot{n}_{\mathrm{i}} = \sum_{\mathrm{j}\in\mathcal{N}} \left[\Delta_{\mathrm{ij}}\, k_{\mathrm{j}} \prod_{\mathrm{k}\in\mathcal{R}_{\mathrm{j}}} n_{\mathrm{k}} \right], \tag{1.1}$$

where $\mathcal{N}$ represents a network of chemical reactions, k_{j} is the rate coefficient of the j'th reaction in the network, $\mathcal{R}_{\mathrm{j}}$ is the set of reactants for that reaction and Δ_{ij} is the change in the number of molecules of species i in a single reaction of j. Equation 1.1 is subject to the constraint that the total density of atoms of any given element l in the set of all elements $\mathcal{E}$ is fixed by the relative atomic abundance for that element, X_{l}:

$$\sum_{\mathrm{i}\in\mathcal{E}} N_{\mathrm{li}} n_{\mathrm{i}} = n_{\mathrm{H}} X_{\mathrm{l}} \tag{1.2}$$

where N_{li} is the number of atoms of the element l in a single molecule of the species i.

Solving for the time–dependent chemistry is an initial value problem requiring the solution to the set of non–linear ordinary differential equations (ODEs) defined by Equation 1.1. In general, the equations are numerically stiff due to the range of reaction rates spanning many orders of magnitude, leading to a range of eigenvalues in the Jacobian matrix making it difficult to invert. The solution to Equation 1.1 therefore requires stable integration schemes and matrix methods specific to stiff problems. In this thesis all such equations are solved using the LSODE routine in the ODEPACK package of ODE solvers (Hindmarsh & Petzold 1995). It implements a range of backward differentiation integration schemes of which the simplest is the backward Euler method that casts the solution vector $\boldsymbol{n}^{\mathrm{t}}$ at time–step t as an implicit function $\boldsymbol{G}$ of the solution at previous steps:

$$\boldsymbol{G}(\boldsymbol{n}^{\mathrm{t}}) = \boldsymbol{n}^{\mathrm{t}} - \boldsymbol{n}^{\mathrm{t}-1} - h\dot{\boldsymbol{n}}^{\mathrm{t}} = 0, \tag{1.3}$$

where h is the step size. Equation 1.3 is solved iteratively using Newton's method:

$$(\mathbf{I} - h\mathbf{J})(\boldsymbol{n}^{\mathrm{t}} - \boldsymbol{n}^{\mathrm{t}-1}) = \mathbf{0}, \tag{1.4}$$

where $\mathbf{I}$ is the identity matrix and $\mathbf{J} = \frac{\partial \dot{\boldsymbol{n}}^{\mathrm{t}}}{\partial \boldsymbol{n}^{\mathrm{t}}}$ is the Jacobian of Equation 1.1 with respect to $\boldsymbol{n}^{\mathrm{t}}$ providing a linearized approximation. Equation 1.4 is solved by direct LU decomposition of the matrix $\mathbf{I} - h\mathbf{J}$. However, the stability of the solution is not always guaranteed; in many cases it has been identified that Equation 1.1 for the system of chemical interaction rates can demonstrate bistable equilibrium solutions (Le Bourlot et al. 1993, Pineau Des Forêts & Roueff 2000). It has been demonstrated that the domain of bistability is sensitive to the presence of certain chemical processes, particularly gas–grain interactions such as depletion (Lee et al. 1998) and surface hydrogenation (Le Bourlot et al. 1995) suggesting the effect is intrinsic to the chemistry and not numerical. Le Bourlot et al. (1995) also identified the two solutions as representing high– and low–ionization phases of the gas, although to date observational verification of the phenomenon remains elusive.

1.3.2 Parametrized Rate Coefficients

There are a number of databases available tabulating laboratory and theoretical rate coefficients for many chemical reactions under astronomical conditions from which a model chemical network $\mathcal{N}$ can be constructed. Popular examples include the University of Manchester Institute of Science and Technology Database for Astrochemistry (UMIST RATE06, Woodall et al. 2007, RATE12, McElroy et al. 2013) and the KInetic Database for Astrochemistry (KIDA, Wakelam et al. 2012, Wakelam et al. 2015). While some studies advocate a minimal chemical network for well studied species where the key chemical pathways are understood (e.g. Keto & Caselli 2008, Caselli et al. 2010) due in part to the significant uncertainties of many reaction rates, models of more complex molecules inherently require more complex networks and the choice of individual reactions to include can be important (e.g. Woods et al. 2012).

Databases typically provide parametrized analytic functions for the rate coefficients k where the constants have been fit to either experimental or theoretical data. For example, in RATE12 the rate coefficients for two–body gas–phase reactions take the form of modified Arrhenius equations which are functions of the kinetic gas temperature $T_{\rm gas}$:

$$k = k_0 \left(\frac{T_{\rm gas}}{300\,{\rm K}}\right)^{\alpha} \exp\left(-\frac{T_{\rm a}}{T_{\rm gas}}\right)\ {\rm cm^3\,s^{-1}}, \tag{1.5}$$

where k_0, α and $T_{\rm a}$ are the constants fit to the data, representing the rate coefficient at 300 K, temperature dependence index and thermal activation barrier respectively. Fits to an equation of this form are typically only valid over the range of gas temperatures for which the original laboratory or theoretical study was carried out since extrapolations outside this range can be unreliable. Those ranges are also quoted in the database.

For photochemical reactions, rate coefficients $k_{\rm pd}$ are determined as a function of the specific flux of UV photons $I(\lambda)$ at wavelength λ from the cross–sections σ of transitions directly into the dissociative continuum:

$$k_{\rm pd}^{\rm cont} = \int_{912\,\text{Å}}^{\infty} \sigma(\lambda) I(\lambda)\, d\lambda\, s^{-1}, \tag{1.6}$$

or from the oscillator strengths $f_{\rm ul}$ of unstable line transitions from lower level l to upper

level u:

$$k_{\rm pd}^{\rm line} = \frac{\pi e^2}{mc^2}\lambda_{\rm ul}^2 f_{\rm ul}\eta_{\rm u}x_{\rm l}I(\lambda_{\rm ul})\,s^{-1}, \tag{1.7}$$

where $\lambda_{\rm ul}$ is the wavelength of the transition, $\eta_{\rm u}$ is the dissociation efficiency of the upper level and $x_{\rm l}$ is the occupation fraction of the lower level. The solutions to Equations 1.6 and 1.7 depend crucially on the local specific photon flux of the UV radiation field $I(\lambda)$ which comes from the solution of the equations of radiative transfer throughout the entire model geometry. In order to simplify the calculation, k is again represented by a parametric model fit to theoretical solutions of the coupled radiative transfer – photochemical rate equations for a given incident radiation field (e.g. Habing 1968, Draine 1978, or a black body radiation field) and simplified geometry (typically a semi–infinite plane parallel slab or a spherical cloud). It is usually assumed that the attenuation due to dust extinction and chemical photoprocesses are separable since dust attenuation is approximately independent of wavelength over the linewidths of individual photochemical transitions. In RATE12 the parametric form of $k_{\rm pd}$ for photoprocesses is given as:

$$k_{\rm pd} = G_0\,\Theta(\boldsymbol{N},\phi)\,k_0\exp\left(-\alpha A_{\rm V}\right), \tag{1.8}$$

where again k_0 and α are the constants fit to the theoretical solution representing the unattenuated rate for the Draine interstellar radiation field (Draine 1978) and the dust attenuation coefficient for the assumed dust properties respectively. The rate is also scaled by G_0, the photon flux of the incident radiation field relative to the Draine radiation field and the exponential term represents a separable extinction $A_{\rm V}$ of the incident UV radiation by dust scattering and absorption. $\Theta(\boldsymbol{N},\phi)$ denotes a shielding function of the incident radiation field as photons are "consumed" effecting photoreactions along their incident path, which depends on the vector of column densities $\boldsymbol{N}$ of all species along the ray of the incident radiation field and ϕ denoting other physical effects such as Doppler line widths. Approximations to the shielding functions tend to only include self–shielding (attenuation of the incident radiation field at the wavelengths of a certain photoreaction by the same photoreaction along the incident path of the photons) for a few of the most important species, again to eliminate the need for computationally expensive radiative transfer

calculations. Examples include the self–shielding of the photodissociation of molecular hydrogen and carbon monoxide (van Dishoeck & Black 1988, Lee et al. 1996, Visser et al. 2009) and the photoionization of atomic carbon (Kamp & Bertoldi 2000). However, it has been identified that mutual shielding (attenuation of the radiation field at the wavelengths of one photoreaction by another photoreaction at the same wavelengths along the incident photon path) can also be significant, such as the shielding of nitrogen bearing species by the carbon ionization continuum (Rollins & Rawlings 2012, see Chapter 4).

By comparison, ionization reactions by cosmic rays take a much simpler form since they penetrate into even the darkest clouds without interacting. The rate coefficients are therefore assumed to be a factor α (tabulated in RATE12) relative to the rate coefficient for molecular hydrogen, ζ, which in the Milky Way takes a typical value between $10^{-17}\,\mathrm{s}^{-1}$ and $10^{-15}\,\mathrm{s}^{-1}$ (e.g. Indriolo et al. 2007, Indriolo & McCall 2012) and orders of magnitude higher in the galactic centre (Yusef-Zadeh et al. 2007). However, as discussed in Section 1.2, cosmic ray ionization events can also lead to the emission of UV photons by excited molecular hydrogen which also lead to photochemistry. The rate coefficients for photoprocesses due to these secondary photons require calculating the branching ratio for the processes relative to absorption by dust as an integrated function over the hydrogen emission spectrum per cosmic ray ionization event $P(\nu)$ as in Sternberg et al. (1987) and Gredel et al. (1989):

$$\mathrm{k_i} = \frac{\zeta}{(1-\omega)} \int \frac{\sigma_\mathrm{i}(\nu)\mathrm{P}(\nu)}{2\sigma_\mathrm{g}} \mathrm{d}\nu\,\mathrm{cm}^{-3}\mathrm{s}^{-1}, \tag{1.9}$$

where $\sigma_i(\nu)$ is the cross–section for the photoprocess and it has been assumed that UV photons are primarily absorbed by dust grains with an albedo ω and absorption cross–section of σ_g per hydrogen nucleon.

Chemical interactions between molecules and dust grains are modelled with varying degrees of sophistication. The rate per unit volume $\dot{n}_\mathrm{i}^\mathrm{ice}$ at which a species i depletes from the gas phase to form an ice mantle is governed by kinetic theory (Spitzer 1978, Umebayashi & Nakano 1980):

$$\dot{n}_\mathrm{i}^\mathrm{ice} \approx \sqrt{\frac{3k_\mathrm{B}T}{m_\mathrm{i}}} \sigma_\mathrm{g} d_\mathrm{g} n_\mathrm{H} n_\mathrm{i}, \tag{1.10}$$

where k_B is the Boltzmann constant, m_i is the mass of the species, σ_g is the mean surface

area of each grain and d_g is the dust–to–gas number ratio. Additional sticking efficiencies are normally used to approximate physisorption, chemisorption and negative electric charges of grains. Once the ice is formed, its chemical behavior is very uncertain. One approach used in the UCL_CHEM chemical model (Viti & Williams 1999, Viti et al. 2004, Viti et al. 2011) is to assume that since the most abundant chemical species in regions of the interstellar medium harboring ices is either atomic or molecular hydrogen, ice species undergo hydrogenation reactions on time–scales much faster than ice desorption processes. They therefore treat species freezing out as instantly forming an ice of their maximally hydrogenated molecular counterpart (e.g. oxygen freezes out to water ice, nitrogen to ammonia ice etc) with the necessary hydrogen taken from the gas–phase reservoir with no consideration of any detailed mechanisms by which this might occur. This implicitly covers the formation of molecular hydrogen on grain surfaces, although an additional temperature dependence for the sticking probability such as in Buch & Zhang (1991) is also typically used. Such approximations allow for a rapid calculation of the interstellar ice abundances that can readily be coupled with the gas–phase chemical evolution and other physics. Some models consider mechanisms such as Langmuir–Hinshelwood in more detail to investigate the formation of more complex molecules (e.g. ethylene oxide, Occhiogrosso et al. 2012). The most sophisticated models such as Taquet et al. (2012) and Garrod (2013) use a Monte Carlo treatment of the sticking potentials, diffusion pathways and pairwise chemical reactions for species on a simulated grain to track the formation and evolution of the ice structure in detail.

The release of ices from grains back into the gas phase is also a complex process. A range of laboratory experiments attempt to characterize the thermal desorption temperatures (i.e. the temperatures at which the bonds between ice species and the grain or remaining mantle are overcome by thermal motions) of interstellar ice analogues (e.g. Fayolle et al. 2011, Noble et al. 2012) although such temperatures depend critically on the heating rate and ice morphology and are not trivial to adapt for astrochemical modelling (Collings et al. 2004). By contrast, non–thermal desorption processes are well modelled by studies such as Roberts et al. (2007) and Garrod et al. (2007) for simple layered structure approximations. However, Monte Carlo grain models demonstrate significant inhomogeneous ice compositions highlighting that the desorption rate for individual species is a function of the chemical history of the ice, with reactive species such as carbon monoxide tending to be trapped close to the grain surface under many other layers of ice (Taquet

et al. 2012).

1.3.3 Coupled Physical Processes

In general, it is straightforward to solve Equation 1.1 for fixed physical input parameters such as the density and temperature of the gas, the incident radiation field and the cosmic ray ionization rate. However, it is also possible to couple these equations to other physics that governs those parameters self–consistently. The most common example is the solution of the thermal balance due to heating and cooling processes (either time–dependent or in thermal equilibrium) to yield the kinetic gas temperature.

In cold molecular gas, diffuse interstellar clouds and photon dominated regions, it has been identified that the major coolants are collisionally excited rotational transitions of molecular species such as molecular hydrogen, carbon monoxide, water and OH and fine structure transitions of atomic species such as ionized carbon and atomic oxygen (Richings et al. 2014*a*, Richings et al. 2014*b*). These form multi–level quantum systems whose level populations are the solution of the rates of radiative and collisional processes:

$$\sum_{\mathrm{j}>\mathrm{i}} (B_{\mathrm{ij}} I_{\mathrm{ij}} + C_{\mathrm{ij}}) n_{\mathrm{i}} + \sum_{\mathrm{j}<\mathrm{i}} (A_{\mathrm{ij}} + B_{\mathrm{ij}} I_{\mathrm{ij}} + C_{\mathrm{ij}}) n_{\mathrm{i}} = \sum_{\mathrm{j}>\mathrm{i}} (A_{\mathrm{ji}} + B_{\mathrm{ji}} I_{\mathrm{ji}} + C_{\mathrm{ji}}) n_{\mathrm{j}} + \sum_{\mathrm{j}<\mathrm{i}} (B_{\mathrm{ji}} I_{\mathrm{ji}} + C_{\mathrm{ji}}) n_{\mathrm{j}}, \quad (1.11)$$

for all levels i where for a transition from levels $\mathrm{i} \rightarrow \mathrm{j}$, the densities of each level are n_{i} and n_{j}, A_{ij} and B_{ij} are the Einstein A and B coefficients, C_{ij} are the collisional (de–)excitation rates and I_{ij} is the mean specific intensity of radiation at the wavelength of the transition. Equation 1.11 is subject to the constraint

$$\sum_{\text{Levels i}} n_{\mathrm{i}} = n_{\mathrm{H}} Y, \quad (1.12)$$

where n_{H} is the total density of hydrogen nucleons and Y is the relative abundance of the species. For atomic species these equations typically simplify to a two or three level system. However for molecular species a full treatment includes a ladder of many rovibrational energy levels. These equations define a cooling function Λ as a sum over all downward

radiative transitions j → i and the energy $E_{\rm ji}$ carried by the emitted photons:

$$\Lambda = \sum_{\rm i} \sum_{\rm j>i} (A_{\rm ji} + B_{\rm ji} I_{\rm ji}) E_{\rm ji} n_{\rm j} \, {\rm erg \, cm^3 \, s^{-1}}. \tag{1.13}$$

A more complete treatment would also consider the radiative transfer of the emitted photons to determine if they truly escape the gas or are reabsorbed along with the energy that they carry. This process is typically modelled using the escape probability formalism with the large velocity gradient approximation (Sobolev 1960) to the full radiative transfer.

Gas heating is typically due to either photons or cosmic rays. Energetic electrons ejected by the photoionization of chemical species (e.g. atomic carbon, Black 1987, more generally as in Richings et al. 2014*a*), cosmic ray ionization of molecular hydrogen (Tielens & Hollenbach 1985) and photoelectric ejection from dust grains and PAHs (Bakes & Tielens 1994) rapidly thermalise, transferring their excess energy to the gas through collisions. The binding energy of molecular hydrogen when it is formed on and then released from grain surfaces also adds to the rovibrational and kinetic energy of the molecule which again is readily thermalised. Energy in turbulent motions of the gas can also be converted to heat by dissipative processes such as viscosity (Black 1987) and ion–neutral friction in the presence of magnetic fields (Godard et al. 2009). The difference between the total heating and cooling rates of these (and other) processes defines the net heating rate $\Gamma_{\rm net}$ from which the gas temperature can then either be time–integrated in parallel with the chemistry or solved iteratively to find the equilibrium temperature T satisfying $\Gamma_{\rm net}(T) = 0$.

Until this point all discussion of photoprocesses has relied on approximations to the radiative transfer within the cloud. However it is also possible to couple chemical models to accurate calculations of the radiation transport by solving the equation of radiative transfer for the specific intensity of radiation I_ν as a function of optical depth τ_ν:

$$\frac{{\rm d}I_\nu}{{\rm d}\tau_\nu} = S_\nu - I_\nu, \tag{1.14}$$

where S_ν is the source term including emission cross–sections, while the optical depth accounts for the cross–sections of attenuating processes including absorption, photoprocesses and scattering. A range of algorithms exist to solve Equation 1.14 with Accelerated

Lambda Iteration being one of the most popular for multilevel systems (Rybicki & Hummer 1991). A successful application to chemical modelling is the MEUDON PDR code (Le Petit et al. 2006) which calculates the wavelength–dependent attenuation of UV photons by dust and line transitions of H, H_2 and CO. Such a treatment requires accurate cross–section data but allows for more accurate calculations of photochemical reaction rates and photoheating rates.

Finally, to lift the approximation of simple, symmetric geometries, chemical codes can be coupled to three dimensional hydrodynamic codes, either for integration in parallel or post processing. The 3D-PDR code (Bisbas et al. 2012) is an example of a recent chemical code that can couple to the outputs of smoothed–particle or grid–based hydrodynamic codes to calculate the photon–dominated chemistry for an arbitrary three dimensional density distribution. It has been successfully coupled to the ORION adaptive–mesh magneto–hydrodynamic code (Truelove et al. 1998) to produce some of the most realistic chemical simulations of molecular clouds to date (Offner et al. 2014).

1.4 Bayesian Statistics

Bayesian statistics is a robust framework for the statistical analysis of theoretical models with free parameters and experimental datasets with their errors. It allows researchers to tackle common problems such as parameter estimation and model comparison in a consistent manner, as well as more sophisticated ones such as the training of statistical models and learning algorithms. Underlying the field is Bayes' theorem for the conditional probability between two independent random variables A and B:

$$P(A|B) = \frac{P(B|A)P(A)}{P(B)}, \tag{1.15}$$

where $P(A)$ represents the probability distribution on the random variable A and $P(A|B)$ is the conditional probability on A given that B is true. Translating those variables to the vector of free parameters $\boldsymbol{\theta}$ for a theoretical model $\mathcal{M}$ and an observational dataset $\mathcal{D}$

yields the fundamental definition of the posterior distribution $P(\boldsymbol{\theta}|\mathcal{D},\mathcal{M})$:

$$P(\boldsymbol{\theta}|\mathcal{D},\mathcal{M}) = \frac{P(\mathcal{D}|\boldsymbol{\theta},\mathcal{M})P(\boldsymbol{\theta}|\mathcal{M})}{P(\mathcal{D}|\mathcal{M})} = \frac{\mathcal{L}(\boldsymbol{\theta})\pi(\boldsymbol{\theta})}{\mathcal{Z}}, \tag{1.16}$$

where $\mathcal{L}(\boldsymbol{\theta}) = P(\mathcal{D}|\boldsymbol{\theta},\mathcal{M})$ is the likelihood distribution describing the uncertainties of the dataset relative to outputs of the model, $\pi(\boldsymbol{\theta}) = P(\boldsymbol{\theta}|\mathcal{M})$ is the prior distribution representing our initial beliefs about the value of $\boldsymbol{\theta}$ from previous experimental or theoretical studies and $\mathcal{Z}$ is the evidence, a normalization factor for the posterior distribution with some important statistical properties. It is the posterior distribution that we wish to determine and not the likelihood distribution as is often used for so–called maximum likelihood analyses. This is because once we have conducted an experiment or observation, any inferences we make about our model must be conditional on the data we have taken.

With the posterior distribution, the two most common inference problems can be tackled. Parameter estimation, or the identification of the most likely $\boldsymbol{\theta}$ for the given data, requires calculating the expectation value of $\boldsymbol{\theta}$ weighted over the posterior distribution:

$$E[\boldsymbol{\theta}] = \int \boldsymbol{\theta}\, \mathrm{d}P(\boldsymbol{\theta}|\mathcal{D},\mathcal{M}), \tag{1.17}$$

while model comparison, the selection of one of two competing models $\mathcal{M}^{\mathrm{A}}$ and $\mathcal{M}^{\mathrm{B}}$ of the same dataset, depends on the ratio of the evidence values $\mathcal{Z}^{\mathrm{A}}$ and $\mathcal{Z}^{\mathrm{B}}$ for the two models:

$$\frac{P(\mathcal{M}^{\mathrm{A}}|\mathcal{D})}{P(\mathcal{M}^{\mathrm{B}}|\mathcal{D})} = \frac{P(\mathcal{D}|\mathcal{M}^{\mathrm{A}})P(\mathcal{M}^{\mathrm{A}})}{P(\mathcal{D}|\mathcal{M}^{\mathrm{B}})P(\mathcal{M}^{\mathrm{B}})} = \frac{\mathcal{Z}^{\mathrm{A}}}{\mathcal{Z}^{\mathrm{B}}}\frac{P(\mathcal{M}^{\mathrm{A}})}{P(\mathcal{M}^{\mathrm{B}})}, \tag{1.18}$$

where the evidence is given by the integral:

$$\mathcal{Z} = \int \mathcal{L}(\boldsymbol{\theta})\pi(\boldsymbol{\theta})\, \mathrm{d}\boldsymbol{\theta}. \tag{1.19}$$

When studying a new problem the prior function is typically uniform over a broad domain given by theoretical considerations while $\frac{P(\mathcal{M}^{\mathrm{A}})}{P(\mathcal{M}^{\mathrm{B}})} = 1$ since there is no reason to support one model over the other. However, when analysing new datasets in light of old ones, the

posterior distribution from the previous study can be used as the prior distribution for the new study. Similarly, the old ratio of evidence values can be used for the ratio $\frac{P(\mathcal{M}^{\mathrm{A}})}{P(\mathcal{M}^{\mathrm{B}})}$. This represents the process of incorporating previous results into prior knowledge.

The challenge in the above analyses lies in the evaluation of the integrals in Equations 1.17 and 1.19. As models become more complex and the dimensionality of the parameter space spanned by $\boldsymbol{\theta}$ grows, more evaluations of the posterior distribution are required to adequately sample the entire domain while many sampling algorithms become less efficient. This is commonly referred to as the "curse of dimensionality". The most effective class of algorithms are Markov chain Monte Carlo algorithms where newly sampled points are generated randomly from a probability distribution dependent on the previous sample(s). Such algorithms are naturally self–adapting, tending towards regions of maximum posterior with their sampling schemes. Two examples of widely used algorithms are the Metropolis–Hastings (Metropolis et al. 1953, Hastings 1970) and simulated annealing (Kirkpatrick et al. 1983, Černý 1985) algorithms. Each have found noteworthy applications in the field of astrochemistry. Makrymallis & Viti (2014) used the Metropolis–Hastings to constrain a number of parameters such as the gas density and the incident radiation field in the UCL_CHEM chemical model given observations of gas and ice abundances in dark molecular clouds. By comparison, Keto et al. (2004) constrained similar parameters but for their radiative transfer model of the N_2H^+ emission spectrum of a starless core. These separate studies ten years apart hint at the range of problems within the field of astrochemical modelling that may be tackled using these very general methods (see Chapter 3), yet also at how slow the uptake of these methods has been. Part of this may relate to the demand for even more efficient sampling schemes to limit the computational expense of running what are already very complex models many times as is required by the algorithms (see Chapter 2).

1.5 This Thesis

This thesis presents work on new chemical models for various regions of interstellar space, including the development and application of a Bayesian statistical code for making robust inferences. The content is outlined below:

Chapter 2: PLINY: A Scalable Parallel Implementation of the Nested Sampling Algorithm for High–Performance Computing Systems

We present PLINY; an implementation of the nested sampling algorithm with a highly optimized MPI parallelization scheme designed for performance on large computer clusters. While a range of popular nested sampling software packages already exist and find use throughout many fields of scientific research, our benchmark testing demonstrates superior results for a number of performance metrics against the industry standard software. We show that PLINY evaluates the Bayesian evidence integral for a range of unimodal posterior distributions to significantly better accuracy from fewer likelihood evaluations while also demonstrating linear parallel scaling to a greater number of CPUs before becoming communication bound on both distributed– and shared–memory architectures. This represents a great improvement in the accuracy vs time–to–solution trade off as nested sampling is used to tackle larger and more complex problems in the era of big–data and petascale computing.

Chapter 3: Probing the Physical State of Diffuse Gas Towards G10.6-0.4 with Nested Sampling

Recent observations made by the *Herschel Space Observatory* have helped to reveal the chemical complexity of diffuse interstellar gas clouds. Many attempts have been made to model the chemical formation pathways to individual species via processes such as photochemistry, reactions on dust grain surfaces, ion–neutral decoupling and cosmic rays. However, no study has attempted to comment on the relative importance of these processes in a self–consistent way by fitting to a range of observations, presumably due to the computational complexity. In this work, we make parametric approximations to the physical state of the diffuse gas clouds towards G10.6-0.4 (W31C) as well as a range of potential chemical mechanisms in order to model and fit to the observed column densities of NH, NH_2, NH_3, CH^+, SH^+, OH^+, H_2O^+ and H_3O^+ simultaneously. We use nested sampling, a Bayesian Markov chain Monte Carlo algorithm, to explore the free parameters of the model in an unbiased and statistically rigorous way and make inferences about their expected values and uncertainties. The results show that the gas towards G10.6-0.4 is relatively dense and cool with grain surface processes playing an important role in the formation of nitrogen hydrides and significant ion–neutral decoupling during a magnetically active phase. However, photochemical processes that have been considered in previous studies appear to be ruled out by the data while no definitive distinction can be made

between turbulently dissipating vortices and a single strong shock as the origin of the molecular ions.

Chapter 4: The Chemical Effects of Mutual Shielding in Photon Dominated Regions

We investigate the importance of the shielding of chemical photorates by molecular hydrogen photodissociation lines and the carbon photoionization continuum deep within models of photon dominated regions. In particular, the photodissociation of N_2 and CN are significantly shielded by the H_2 photodissociation line spectrum. We model this by switching off the photodissociation channels for these species behind the HI $\rightarrow$ H_2 transition. We also model the shielding effect of the carbon photoionization continuum as an attenuation of the incident radiation field shortwards of 1102 Å. Using recent line and continuum cross–section data, we present calculations of the direct and cosmic–ray–induced photorates for a range of species, as well as optically thick shielding factors for the carbon continuum. Applying these to a time–dependent PDR model we see enrichments in the abundances of N_2, N_2H^+, NH_3 and CN by factors of $\sim 3-100$ in the extinction band $A_v = 2.0-4.0$ for a range of environments. While the precise quantitative results of this study are limited by the simplicity of our model, they highlight the importance of these mutual shielding effects, neither of which has been discussed in recent models.

Chapter 5: Champagne Flutes and Brandy Snifters: Modelling Protostellar Outflow–Cloud Chemical Interfaces

A rich variety of molecular species has now been observed towards hot cores in star–forming regions and in the interstellar medium. An increasing body of evidence from millimetre interferometers suggests that many of these form at the interfaces between protostellar outflows and their natal molecular clouds. However, current models have remained unable to explain the origin of the observational bias towards wide–angled "brandy snifter" shaped outflows over narrower "champagne flute" shapes in carbon monoxide imaging. Furthermore, these wide–angled systems exhibit unusually high abundances of the molecular ion HCO^+. We present results from a chemo–dynamic model of such regions where a rich chemistry arises naturally as a result of turbulent mixing between cold, dense molecular gas and the hot, ionized outflow material. The injecta drives a rich and rapid ion–neutral chemistry in qualitative and quantitative

agreement with the observations. The observational bias towards wide–angled outflows is explained naturally by the geometry–dependent ion injection rate causing rapid dissociation of CO in the younger systems.

Chapter 2

Pliny: A Scalable Parallel Implementation of the Nested Sampling Algorithm for High–Performance Computing Systems

The work presented in this chapter is based on the paper by Rollins et al., submitted to MNRAS, in collaboration with S. T. Balan, F. B. Abdalla and S. A. Thomas.

Bayesian statistics provides a powerful and rigorous framework for the analysis of observational data and theoretical models in astrophysics and cosmology. Markov chain Monte Carlo (MCMC) methods such as the Metropolis–Hastings algorithm (Metropolis et al. 1953, Hastings 1970) allow for the exploration of large parameter spaces by generating samples from the posterior distribution, allowing users to make inferences about both the input and derived parameters of a parametric model. However, as problems become more complex, including larger parameter spaces of higher dimensionality as well as multimodality and curved degeneracies in the posterior distribution, they can quickly become

inefficient with high rejection rates and requiring many iterations to fully explore the parameter space. Furthermore, calculation of the Bayesian evidence statistic by this method is typically very computationally expensive for the necessary level of precision required to make reliable model comparisons. While other sampling schemes such as simulated annealing (Kirkpatrick et al. 1983, Černý 1985), Gibbs sampling (German & German 1984) and Hamiltonian sampling (Duane et al. 1985) as well as the technique of thermodynamic integration (Ó Ruanaidh & Fitzgerald 1996) for evidence calculation have found some success in addressing these issues they are often insufficient and there is clearly demand for more sophisticated and powerful algorithms to make these analyses simpler as well as more accurate and computationally efficient.

Nested sampling is an example of such a modern algorithm that has gained a lot of popularity over the last decade. Although designed to evaluate Bayesian evidence integrals for model comparison, it also generates weighted samples of the posterior distribution that can be used for parameter estimation, making it exceptionally versatile. It uses a set of "live points" and a rejection scheme based on the minimum likelihood of those live points to define an adaptive scheme that rapidly restricts sampling to regions of high likelihood. In one of the earliest applications to the field of cosmology, Mukherjee et al. (2006) demonstrated that for a typical model selection problem nested sampling required orders of magnitude fewer likelihood evaluations than thermodynamic integration.

MultiNest (Feroz & Hobson 2008, Feroz et al. 2009, Feroz et al. 2013) is the current most widely used nested sampling software package that is freely available under license. It boasts a clustering algorithm that partitions live points based on the minimum volume set of bounding ellipsoids that allows for an efficient treatment of multimodal posteriors or those with strong curving degeneracies, as well as a Message Passing Interface (MPI) parallelization scheme. A recent update implemented the importance nested sampling scheme (Feroz et al. 2013) for improved accuracy while the popular BAMBI software package (Graff et al. 2012) has combined MultiNest with an artificial neutral network for fast approximation of the likelihood function. To date, MultiNest has found wide application in the fields of astrophysics and cosmology (e.g. radio interferometry, Lochner et al. 2015, exoplanet characterization, Placek et al. 2014, pulsar timing, Vigeland & Vallisneri 2014, inflationary cosmology, Martin et al. 2014, active galactic nuclei, Han & Han 2012) as well as particle physics (SuperBayeS, Ruiz de Austri et al. 2011).

The repetitive nature of evaluating likelihood functions for many parameters lends itself naturally to the data parallelism paradigm. However, while large numbers of processors are readily available to researchers, from local clusters through to TOP500 supercomputers and cloud services, naive parallelization schemes can quickly become communication bound, preventing the effective utilization of further resources. In this paper we present PLINY, an open source implementation of the nested sampling algorithm in C++ designed with optimal parallel scaling on both distributed– and shared–memory clusters as a primary goal. This is achieved through an MPI communication scheme highly optimized for nested sampling. Points are efficiently sampled from an accurate minimum volume covering ellipsoid fit to the live points, although the templated design pattern makes implementation of a user–defined sampler possible. The software is available under the Mozilla Public License 2.0 from `https://github.com/Astrophysics-UCL/pliny` and has recently found its first application in the field of astrochemistry (Rollins & Rawlings, submitted).

In this paper we compare the performance of PLINY and MULTINEST using a number of metrics including accuracy of the evidence integral, time–to–solution for parallel computation and optimal sampling acceptance rates. We present our work in five Sections. In Section 2.1 we describe the nested sampling algorithm and outline our implementation in PLINY, including the ellipsoid sampling algorithm in Section 2.1.1 and the MPI parallelization scheme in Section 2.1.2. We outline our performance benchmarks in Section 2.2; tests for the accuracy of the evidence integral calculation in Section 2.2.1 and for the parallel scalability on distributed– and shared–memory systems in Section 2.2.2. The results of those tests are presented in Section 2.3 and discussed in Section 2.4 before drawing our conclusions in Section 2.5.

2.1 The PLINY Nested Sampling Algorithm

Nested sampling is a recent MCMC algorithm (Skilling 2004) designed to calculate the Bayesian evidence integral $\mathcal{Z}$ for a parametric model $\mathcal{M}$ of a dataset $\mathcal{D}$:

$$\mathcal{Z} = \int \mathcal{L}(\boldsymbol{\theta})\pi(\boldsymbol{\theta})\,\mathrm{d}\boldsymbol{\theta}, \tag{2.1}$$

where $\mathcal{L}(\boldsymbol{\theta}) = P(\mathcal{D}|\boldsymbol{\theta}, \mathcal{M})$ is the likelihood distribution and $\pi(\boldsymbol{\theta}) = P(\boldsymbol{\theta}|\mathcal{M})$ is the prior distribution of the vector of free parameters $\boldsymbol{\theta}$ belonging to $\mathbb{R}^n$ for an n–dimensional prob-

lem. It also represents the normalization factor of the posterior probability distribution of $\boldsymbol{\theta}$ which is given by Bayes' theorem:

$$P(\boldsymbol{\theta}|\mathcal{D},\mathcal{M}) = \frac{P(\mathcal{D}|\boldsymbol{\theta},\mathcal{M})P(\boldsymbol{\theta}|\mathcal{M})}{P(\mathcal{D}|\mathcal{M})} = \frac{\mathcal{L}(\boldsymbol{\theta})\pi(\boldsymbol{\theta})}{\mathcal{Z}}. \tag{2.2}$$

$\mathcal{Z}$ is of statistical significance since Bayes' theorem also shows that the ratio of posterior probabilities for two models $\mathcal{M}^A$ and $\mathcal{M}^B$ of the same data is given as:

$$\frac{P(\mathcal{M}^{\mathrm{A}}|\mathcal{D})}{P(\mathcal{M}^{\mathrm{B}}|\mathcal{D})} = \frac{P(\mathcal{D}|\mathcal{M}^{\mathrm{A}})P(\mathcal{M}^{\mathrm{A}})}{P(\mathcal{D}|\mathcal{M}^{\mathrm{B}})P(\mathcal{M}^{\mathrm{B}})} = \frac{\mathcal{Z}^{\mathrm{A}}}{\mathcal{Z}^{\mathrm{B}}}\frac{P(\mathcal{M}^{\mathrm{A}})}{P(\mathcal{M}^{\mathrm{B}})}, \tag{2.3}$$

demonstrating that the relative ability of two models to reproduce an observed dataset is proportional to the ratio of their evidence values. Those values are typically compared using the scales of either Jeffreys (1961) or Kass & Raftery (1995). Accurate calculation of evidence values is therefore fundamental to problems of model comparison. Equation 2.1 is an n–dimensional integral, meaning that MCMC algorithms can become inefficient at sampling the entire parameter space to sufficient accuracy in higher dimensions, commonly referred to as the "curse of dimensionality". Although also MCMC based, the nested sampling algorithm naturally adapts its sampling scheme to regions of high likelihood leading to efficient sampling from the bulk of posterior mass and hence accurate determination of $\mathcal{Z}$.

The basic algorithm begins by sampling a set $\mathcal{S}$ of N "live points" from the prior distribution π and their likelihood values are calculated. The minimum likelihood of points in $\mathcal{S}$ defines the critical likelihood value $\mathcal{L}^*$. Further points are then sampled from π until one is found with likelihood satisfying $\mathcal{L} > \mathcal{L}^*$. That point is then added to the set of live points, the live point with the minimum likelihood is removed and added to a chain $\mathcal{C}$ and the critical likelihood $\mathcal{L}^*$ is recalculated. The point moved to $\mathcal{C}$ is also assigned a weight equal to the change in prior volume of the region satisfying $\mathcal{L} > \mathcal{L}^*$ between the old and new values of $\mathcal{L}^*$, which can be determined statistically (Skilling 2004). The points in $\mathcal{C}$ and their weights have the property that they form a weighted sample from π and as such Equation 2.1 can be cast as sum over the likelihoods $\mathcal{L}_{\mathrm{i}}$ of points i in $\mathcal{C}$ and their weights

w_i:

$$\mathcal{Z} = \sum_{\mathrm{i}\in\mathcal{C}} w_\mathrm{i}\mathcal{L}_\mathrm{i}. \tag{2.4}$$

The procedure of sampling and replacing points is repeated until some implementation–dependent convergence criterion is met.

Different implementations of nested sampling employ their own algorithms for sampling points from π satisfying $\mathcal{L} > \mathcal{L}^*$ efficiently. With PLINY, we initially sample points from the entire space using rejection sampling against the distribution π until we can fit a minimum volume covering ellipsoid to the live points with volume less than the total volume of the prior space at which point we switch to using an ellipsoid sampling algorithm (see Section 2.1.1). Currently PLINY is only able to sample from a single ellipsoid fit to all of the live points, although extending the code to cluster live points and fit many sampling ellipsoids (e.g. Feroz & Hobson 2008, Feroz et al. 2009) would be possible. We also use an optimized MPI algorithm that allows for multiple computer processors to calculate the likelihood values of many candidate live points in parallel in a time–efficient manner, outlined in Section 2.1.2. In calculating the evidence integral we use the statistical expectation for the remaining prior volume after i steps $X_\mathrm{i} = e^{-\mathrm{i}/N}$. We employ a trapezoidal integration scheme so that the prior weighting of the i'th point in the chain is $w_\mathrm{i} = e^{-\mathrm{i}/N}\sinh(1/N)$. Convergence of the algorithm occurs after I steps when the ratio between the contribution to the evidence of the live points $\mathcal{S}$ and the chain points $\mathcal{C}$ falls below a user–input tolerance t:

$$\frac{e^{-\frac{I}{N}}}{N}\sum_{\mathrm{j}\in\mathcal{S}}\mathcal{L}_\mathrm{j} < t\sum_{\mathrm{i}\in\mathcal{C}} w_\mathrm{i}\mathcal{L}_\mathrm{i}. \tag{2.5}$$

The PLINY algorithm features two important unique features; an accurate algorithm for the bounding ellipsoid of the set of live points and an efficient MPI communication scheme. Together they represent a significant improvement in efficiency compared with existing implementations. We describe these two features in detail below.

2.1.1 Ellipsoid Sampling

Restricting the volume of prior space from which candidate live points are sampled is important in order to ensure a reasonable sampling acceptance rate (the ratio of live points accepted to the total number of candidate points sampled). MULTINEST takes the approach of finding one or more ellipsoids that together bound all of the live points and treating their surface as an approximation to the critical iso–likelihood contour $\mathcal{L}^*$. New candidate live points are then sampled uniformly from within their volume and rejected according to the usual criteria of $\mathcal{L} > \mathcal{L}^*$. The ellipsoids should have as small a volume as possible while still bounding the $\mathcal{L} > \mathcal{L}^*$ region so as to minimize the rate of rejected points. MULTINEST approximates the minimum–volume covering ellipsoid (MVCE) of the live points by the inverse of their covariance matrix. The volume of the ellipsoid is then fixed to the statistical expectation for the prior space volume X_i and subsequently scaled by a further user defined target acceptance factor to be somewhat larger than the true MVCE due to fluctuations in X_i. Shaw et al. (2007) use a slightly different scheme of first enlarging the ellipsoid so that it bounds all live points explicitly and then further expanding by the target acceptance factor.

In general, the true MVCE in n dimensions is a fit to up to $n(n+3)/2$ points that lie on the ellipsoid surface (John 1948). However in MULTINEST their covariance method does not fit the surface of the ellipsoid to any points while Shaw et al. (2007) only fit the surface to a single point. As such, in both cases the ellipsoid surface does not necessarily closely trace the true iso–likelihood surface $\mathcal{L} = \mathcal{L}^*$. This can lead to an ellipsoid with volume much larger than the volume of the true MVCE and hence poor acceptance, or where regions of the prior space satisfying $\mathcal{L} > \mathcal{L}^*$ are not enclosed which could bias the sampling to higher likelihoods.

PLINY also implements an ellipsoidal sampler but with a more accurate algorithm for determining the MVCE of the live points. The problem of determining MVCEs is an important topic in constrained optimization that has attracted a lot of research interest and more optimal algorithms exist (Khachiyan 1996). We use the dual reduced Newton (DRN) algorithm of Sun & Freund (2004) which solves the dual of the optimization problem with a logarithmic barrier function by determining the gradient towards the optimal solution at each iteration using Newton's method. We use convergence tolerances for feasibility of the solution of $\epsilon_1 = 10^{-2}$ and for the optimality gap of $\epsilon_2 = 10^{-2}$ (as defined in Sun &

Freund 2004) since we find that higher accuracy is not needed for our application. In Sun & Freund (2004) they discuss possible "active set" schemes that initially fit to only a subset of the points that hopefully lie closest to the MVCE surface, then adjust the active set and refit iteratively until all points are enclosed. Since the main computational cost of fitting to m points is the inversion of an $m \times m$ matrix, such schemes can greatly speed up the algorithm. We note that in nested sampling, the points lying close to the MVCE surface are typically those that have the lowest likelihoods. Denoting the set of N live points by $\mathcal{S} = \{\boldsymbol{u_1}, \boldsymbol{u_2}, \boldsymbol{u_3} \dots \boldsymbol{u_N}\}$ we form the initial active set $\mathscr{S}_0$ from the $n^{1.5}+1$ live points with the lowest likelihoods and find the initial solution $(\boldsymbol{c_0}, \mathbf{Q_0})$ for the bounding ellipsoid centre and shape matrix respectively by running the DRN algorithm on those points in $\mathscr{S}_0$ only. The final solution is then found by iteratively updating the active set and recalculating the ellipsoid using DRN until it bounds all of the live points. At each iteration k, we test if all live points are enclosed by the current solution $(\boldsymbol{c_k}, \mathbf{Q_k})$ such that:

$$(\boldsymbol{u_i} - \boldsymbol{c_k})^T \mathbf{Q_k} (\boldsymbol{u_i} - \boldsymbol{c_k}) < 1 + \epsilon_1. \tag{2.6}$$

If Equation 2.6 holds for all live points $\boldsymbol{u_i} \in \mathcal{S}$ then the algorithm is converged. If not, we form a new active set $\mathscr{S}_{k+1}$ from the active points $a_i \in \mathscr{S}_k$ that are identified as being close to the surface of the previous solution by satisfying:

$$(\boldsymbol{a_i} - \boldsymbol{c_k})^T \mathbf{Q_k} (\boldsymbol{a_i} - \boldsymbol{c_k}) > 1 - \sqrt{\epsilon_1}, \tag{2.7}$$

where the limit $1 - \sqrt{\epsilon_1}$ is seen from our testing to be a suitable value. In addition, up to $n(n+3)/2$ of the live points $\boldsymbol{\theta_i}$ that failed to satisfy Equation 2.6 are also added to $\mathscr{S}_{k+1}$. If there are more than $n(n+3)/2$ such points then those points with the lowest likelihoods are used since they are likely to lie closest to the true MVCE boundary. Once the iterative procedure is converged we stretch each axis of the ellipsoid by a factor of $f^{-1/n}$ where f is the user–input target acceptance rate of the sampler (between 0 and 1), giving us a vector $\boldsymbol{c}$ for the centre and a matrix $\mathbf{Q}$ for the shape of the ellipsoid. The complete algorithm is summarised in Algorithm 1. While this algorithm is not free from the same risks of not bounding the entire region $\mathcal{L} > \mathcal{L}^*$, it certainly produces a more accurate bounding ellipsoid where the surface is fit to as many as $n(n+3)/2$ points and with a smaller volume

leading to improved sampling efficiencies over covariance based methods set to the same target acceptance rate.

Data: Set of live points $\mathcal{S}$. Target acceptance rate f.
Result: A sampling ellipsoid with centre $\boldsymbol{c}$ and shape $\mathbf{Q}$ bounding the live points $\mathcal{S}$.
$\mathscr{S}_0 \leftarrow \min(\mathrm{N}, n^{1.5}+1)$ lowest $\mathcal{L}$ live points;
$(\boldsymbol{c_0}, \mathbf{Q_0}) \leftarrow DRN(\mathscr{S}_0)$;
$k \leftarrow 0$;
Converged $\leftarrow$ Equation 2.6 $\forall\, \boldsymbol{u_i} \in \mathcal{S}$;
while *not Converged* **do**
  $\mathscr{S}^{\mathrm{new}}_{\mathrm{k+1}} \leftarrow$ Up to $\frac{n(n+3)}{2}$ lowest likelihood points not in $\mathscr{S}_{\mathrm{k}}$ and not satisfying Equation 2.6;
  $\mathscr{S}^{\mathrm{old}}_{\mathrm{k+1}} \leftarrow$ All points in $\mathscr{S}_{\mathrm{k}}$ satisfying Equation 2.7;
  $\mathscr{S}_{\mathrm{k+1}} \leftarrow \mathscr{S}^{\mathrm{new}}_{\mathrm{k+1}} \cup \mathscr{S}^{\mathrm{old}}_{\mathrm{k+1}}$;
  $(\boldsymbol{c_{k+1}}, \mathbf{Q_{k+1}}) \leftarrow DRN(\mathscr{S}_{\mathrm{k+1}})$;
  $k \leftarrow k+1$;
  Converged $\leftarrow$ Equation 2.6 $\forall\, \boldsymbol{u_i} \in \mathcal{S}$;
end
$\boldsymbol{c} \leftarrow \boldsymbol{c_k}$;
$\mathbf{Q} \leftarrow f^{2/n}\mathbf{Q_{k+1}}$;

Algorithm 1: PLINY minimum volume bounding ellipsoid employing a custom minimum likelihood active–set strategy for the DRN algorithm (Sun & Freund 2004)

We sample random points uniformly from within the volume of the ellipsoid through three stages as in Shaw et al. (2007). First a point is sampled uniformly from the surface of a unit n–sphere by sampling an n–dimensional standard Gaussian variate (which is rotationally invariant, Knuth 1969). Next, a random radius is drawn from the distribution $\frac{\mathrm{d}P}{\mathrm{d}r} = r^{1-n}$ and the previous point is projected inside the n–sphere to radius r giving a point $\mathbf{p}$. Finally, $\mathbf{p}$ is transformed to a point $\boldsymbol{\theta}$ in the prior space bounded by the ellipsoid:

$$\boldsymbol{\theta} = \boldsymbol{c} + \mathsf{X}\boldsymbol{\Lambda}\mathbf{p}, \tag{2.8}$$

where X is the matrix of (column) eigenvectors of $\mathbf{Q}$ and $\boldsymbol{\Lambda}$ is the diagonal matrix with each element given by the eigenvalues of $\mathbf{Q}$ raised to the power of -0.5, equivalent to the ellipsoid axes lengths. Each $\boldsymbol{\theta}$ is then rejection sampled against the prior probability $\pi(\boldsymbol{\theta})$ to ensure that the ellipsoid sampled points are prior distributed before their likelihoods are calculated and tested against $\mathcal{L}^*$. With every step taken by the nested sampler $\mathbf{Q}$ is scaled by a factor of $e^{2/(nN)}$ representing the expected statistical shrinkage of the prior volume. Once $N \ln f^{-1}$ nested sampler steps have been taken the statistical shrinkage will

have matched the original stretch by the target acceptance rate meaning further symmetric statistical shrinkage would be unreliable and so we refit the ellipsoid.

2.1.2 Parallel Scheme

In general, Bayesian MCMC algorithms spend the majority of their time repeatedly evaluating the likelihood function for new points. Being able to spread this computational cost over multiple processors can lead to dramatic speedup. Nested sampling lends itself to a very natural parallelization scheme whereby every iteration each processor samples its own candidate live point, calculates its likelihood and sends the results to a master process. The master process can then test all candidate points sequentially to see if they lie above $\mathcal{L}^*$. The same list of candidates can still be used even after taking intermediate nested sampling steps since the only change to the target distribution is an increase in the likelihood threshold. This is broadly the scheme followed by MULTINEST. However, this scheme has two shortcomings that ultimately limit its performance, namely that the master process must spend time sending data about the sampling ellipsoid to all other processors and those processors have to send back not only the likelihood of the point but also the point itself.

In PLINY there is no master process; all processors keep and maintain their own identical local copies of all data structures, including but not limited to the live points and the sampling ellipsoid. This way, there is never any need to send information about the ellipsoid sampler between processors; each processor independently runs Algorithm 1 on its own local copy of the live points to generate its own local copy of the ellipsoid sampler. The obvious trade–off is the memory requirement of duplicating data across processors, although PLINY's serial memory footprint is in general negligibly small. At each iteration when running the code on p processors, each processor samples p new candidate points by the combination of ellipsoid sampling and prior rejection sampling detailed in Section 2.1.1 and stores them locally. Since each processor has its own random number generator and they are all seeded to the same value, each processor generates the same p points and so there is never any need to send new points between processors. Instead, each processor calculates the likelihood for one of the p points and then distributes only that value (plus any other derived parameters the user wishes to save from their likelihood calculation) to all other processes via the MPI_BCAST method. Each processor then accepts or rejects each live point locally. This algorithm ensures that the same points are sampled and

accepted by each processor so that they remain synchronized. By only parallelizing the computationally dominant bottleneck of calculating likelihoods and duplicating all other computations across the system to maintain local copies of all data structures we avoid unnecessary extra communications and the overheads that they incur.

2.2 Benchmarking

In all of our tests we use the most up to date versions of each software, Pliny 1.0 and MultiNest 3.9 respectively. MultiNest has a number of features not implemented in Pliny, such as clustering the live points for multimodal posteriors to form more than one sampling ellipsoid and improve the acceptance rate for new live points with highly non-Gaussian posteriors. It also has an importance sampling mode that can improve the accuracy of the evidence integral without sampling extra points and a constant efficiency mode where the volume of the sampling ellipsoids is varied dynamically to achieve the users target acceptance rate. In order to make a direct comparison of the fundamental algorithms of nested sampling and ellipsoidal sampling in terms of accuracy and parallel performance of the two implementations, we consider only broadly unimodal benchmarks and switch off the multimodal sampling, importance sampling and constant efficiency features of MultiNest unless stated otherwise.

2.2.1 Evidence Benchmarks

Of fundamental importance to any implementation of the nested sampling algorithm is its ability to accurately and reliably calculate the evidence integral $\mathcal{Z}$. We test our code Pliny against the well established MultiNest implementation to compare their relative capabilities in calculating the evidence of nine analytic test likelihood functions specified in Table 2.1. In each case, we assume a uniform prior over the domain of the function and run the nested sampler using 1000 live points, a convergence tolerance on $\ln \mathcal{Z}$ of $t = 10^{-2}$ and a target sampling acceptance rate of 1.1^{-n} motivated by Shaw et al. (2007) where n is the dimensionality of the parameter space. We repeat the calculations for each function with each code 1000 times using different seeds for the random number generators in order to investigate the distribution on output evidence values.

Seven of the functions (Ackley, Beale, Booth, Goldsten–Price, Lévi N.13, Matyas, Mc-

Table 2.1. Names, dimensionalities n, domains and functional forms for the likelihoods of our nine evidence benchmarks. In each case the prior function is uniform over the domain and the posterior mass is dominated by a single maxima. All functions are taken from the Virtual Library of Simulation Experiments: Test Functions and Datasets (`http://www.sfu.ca/~ssurjano`, Surjanovic & Bingham 2015) except for the Gaussian model which was taken from Feroz et al. (2010).

Name	n	Domain	$\ln \mathcal{L}(\boldsymbol{\theta})$
Ackley	2	$\boldsymbol{\theta}_1, \boldsymbol{\theta}_2 \in [-33, 33]$	$20 \exp\left(-0.2\sqrt{0.5(\boldsymbol{\theta}_1^2 + \boldsymbol{\theta}_2^2)}\right)$ $+ \exp\left(\frac{1}{2}\cos(2\pi\boldsymbol{\theta}_1) + \frac{1}{2}\cos(2\pi\boldsymbol{\theta}_2)\right) - 20 - \mathrm{e}$
Beale	2	$\boldsymbol{\theta}_1, \boldsymbol{\theta}_2 \in [-4.5, 4.5]$	$-(1.5 - \boldsymbol{\theta}_1 + \boldsymbol{\theta}_1\boldsymbol{\theta}_2)^2 - (2.25 - \boldsymbol{\theta}_1 + \boldsymbol{\theta}_1\boldsymbol{\theta}_2^2)^2$ $-(2.625 - \boldsymbol{\theta}_1 + \boldsymbol{\theta}_1\boldsymbol{\theta}_2^3)^2$
Booth	2	$\boldsymbol{\theta}_1, \boldsymbol{\theta}_2 \in [-10, 10]$	$-(\boldsymbol{\theta}_1 + 2\boldsymbol{\theta}_2 - 7)^2 - (2\boldsymbol{\theta}_1 + \boldsymbol{\theta}_2 - 5)^2$
Goldstein –Price	2	$\boldsymbol{\theta}_1, \boldsymbol{\theta}_2 \in [-2, 2]$	$-(1 + (\boldsymbol{\theta}_1 + \boldsymbol{\theta}_2 + 1)^2$ $\times(19 - 14\boldsymbol{\theta}_1 + 3\boldsymbol{\theta}_1^2 - 14\boldsymbol{\theta}_2 + 6\boldsymbol{\theta}_1\boldsymbol{\theta}_2 + 3\boldsymbol{\theta}_2^2))$ $\times(30 + (2\boldsymbol{\theta}_1 - 3\boldsymbol{\theta}_2)^2$ $\times(18 - 32\boldsymbol{\theta}_1 + 12\boldsymbol{\theta}_1^2 + 48\boldsymbol{\theta}_2 - 36\boldsymbol{\theta}_1\boldsymbol{\theta}_2 + 27\boldsymbol{\theta}_2^2))$
Lévi N.13	2	$\boldsymbol{\theta}_1, \boldsymbol{\theta}_2 \in [-10, 10]$	$-\sin^2(3\pi\boldsymbol{\theta}_1) - (\boldsymbol{\theta}_1 - 1)^2(1 - \sin^2(3\pi\boldsymbol{\theta}_2))$ $-(\boldsymbol{\theta}_2 - 1)^2(1 - \sin^2(2\pi\boldsymbol{\theta}_2))$
Matyas	2	$\boldsymbol{\theta}_1, \boldsymbol{\theta}_2 \in [-10, 10]$	$-0.26(\boldsymbol{\theta}_1^2 + \boldsymbol{\theta}_2^2) + 0.48(\boldsymbol{\theta}_1\boldsymbol{\theta}_2)$
McCormick	2	$\boldsymbol{\theta}_1 \in [-1.5, 4.0]$ $\boldsymbol{\theta}_2 \in [-3.0, 4.0]$	$-\sin(\boldsymbol{\theta}_1 + \boldsymbol{\theta}_2) - (\boldsymbol{\theta}_1 - \boldsymbol{\theta}_2)^2 + 1.5\boldsymbol{\theta}_1 - 2.5\boldsymbol{\theta}_2 - 1$
Gaussian	10	$\boldsymbol{\theta}_\mathrm{i} \in [0, 1]$	$-\sum_{\mathrm{i}=1}^{10} \ln(0.002\pi\mathrm{i}) - \sum_{\mathrm{i}=1}^{10} 0.5\left(\frac{\boldsymbol{\theta}_\mathrm{i} - 0.5}{0.001\mathrm{i}}\right)^2$
Sphere	10	$\boldsymbol{\theta}_\mathrm{i} \in [-5, 5]$	$-\sum_{\mathrm{i}=1}^{10} \boldsymbol{\theta}_\mathrm{i}^2$

Cormick) are well established two–dimensional test functions. The functional forms given in Surjanovic & Bingham (2015) have a single global minimum and are typically used for benchmarking optimization problems and so we use the negative of those functions in order to provide a single global maximum likelihood that contains the bulk of the posterior mass. The iso–likelihood contours of these functions are non–ellipsoidal and as such will test the ellipsoidal fitting algorithms of the two codes in their ability to accurately construct a bounding ellipsoid to the critical likelihood contour $\mathcal{L}^*$ given the constraint of the target acceptance. Being two–dimensional, the evidence integrals for these functions can be trivially evaluated using standard numerical integration schemes for comparison with the outputs of the nested sampling codes.

To complement these functions, we also consider two ten dimensional functions; the Sphere function (Surjanovic & Bingham 2015) and an anisotropic Gaussian function (Feroz et al. 2010) with a maximal aspect ratio of 10. Although in these cases the iso–likelihood contours are ellipsoids, the functions test the ability of the codes to approximate those

contours when they are more sparsely sampled in higher dimensions. These functions have analytical evidence integrals again allowing for direct comparison with the code outputs.

Although the underlying nested sampling algorithm is identical in the implementations of Pliny and MultiNest and both utilize ellipsoidal sampling, we expect potential differences in the outputs due to the differing algorithms used to construct the sampling ellipsoid. In particular, when the ellipsoid fails to bound the entire region satisfying $\mathcal{L} > \mathcal{L}^*$ due to either algorithm accuracy, sparsity of live points or using too low of a target acceptance, we would expect the evidence value calculated to be higher than the analytic values due to biasing of the sampled live points to higher likelihoods and faster shrinkage of the prior volume relative to the statistical expectation leading to overweighting of the points in the chain.

2.2.2 Parallel Performance

For more complex problems including time consuming likelihood calculations, higher dimensionality parameter spaces and greater numbers of live points, it becomes of practical advantage to utilize parallel algorithms that can distribute the sampling of new live points over multiple central processing units (CPUs), from multi–core processing chips all the way up to world class supercomputing facilities. In this regime the time–to–solution becomes an important performance diagnostic as users look to make optimal use of the available computing resources. Whilst near linear parallel scaling with CPU count can be expected when the computation time is dominated by a very slow likelihood function, for faster likelihood calculations it is important to understand the parallelization scheme and identify the regime in which the algorithm becomes communication bound.

To test the relative parallel time–to–solution performance of Pliny and MultiNest we run a number of tests on a multi–dimensional Gaussian likelihood function:

$$\ln \mathcal{L}(\boldsymbol{\theta}, n) = -\frac{n}{2} \ln(2\pi\sigma^2) - \sum_{\mathrm{i}=1}^{n} \frac{1}{2} \left(\frac{\boldsymbol{\theta}_{\mathrm{i}} - 0.5}{\sigma} \right)^2, \tag{2.9}$$

Table 2.2. Combinations of dimensionality, n, number of live points, N and artificial delay on the likelihood function, τ, considered for the Gaussian parallel performance benchmark function defined by Equations 2.9 and 2.10.

Test	n	N	τ/s
A	1	10^3	10^{-3}
B	1	10^3	10^{-2}
C	1	10^3	10^{-1}
D	10	10^3	10^{-3}
E	10	10^3	10^{-2}
F	10	10^3	10^{-1}
G	10	10^2	10^{-3}
H	10	10^4	10^{-3}
I	20	10^3	10^{-3}
J	30	10^3	10^{-3}
K	50	10^3	10^{-3}

$$\pi(\boldsymbol{\theta}) = \begin{cases} 1.0 & \mathbf{0} \leq \boldsymbol{\theta} \leq \mathbf{1} \\ 0.0 & \text{otherwise} \end{cases} \tag{2.10}$$

where the number of dimensions n can be changed, $\sigma = 0.001$ is the scalar standard deviation and $\mathbf{0}$ and $\mathbf{1}$ are the vectors of all zeros and ones respectively. We also implement an artificial time delay τ in the code for the likelihood function such that it will not return until a specified time inside the function has passed based on the MPI_WTIME function. As in the benchmarks in Section 2.2.1, we use an evidence convergence tolerance of $t = 10^{-2}$ and a target sampling acceptance rate of 1.1^{-n}. Since the analytic evidence value for the evidence is $\ln \mathcal{Z} = 0$ independent of the dimensionality of the Gaussian, we use the time (again taken from the MPI_WTIME function) for the evidence integral to reach $\ln \mathcal{Z} = -1$ as our time–to–solution performance diagnostic; a compromise accounting for variability in the final calculated evidence values but still sampling sufficiently to reach the bulk of posterior mass. This allows us to standardise the likelihood function across both codes so that only the internal algorithm implementations of the code are tested, including the MPI communication scheme, ellipsoidal decomposition speed and likelihood evaluations needed to reach the solution. Table 2.2 details the combinations of dimensionality, number of live points and likelihood calculation times investigated.

The performance of a particular parallel scheme is also critically dependent on the system architecture it is run on and in particular the interconnect network through which

Table 2.3. Analytic evidence values $\mathcal{Z}_{\mathrm{True}}$ and the mean and standard deviations of the evidence values of 1000 runs each by both PLINY ($\mathcal{Z}_{\mathrm{PLINY}}$) and MULTINEST ($\mathcal{Z}_{\mathrm{MULTINEST}}$) for the nine evidence benchmarks defined in Table 2.1.

Function	$\mathcal{Z}_{\mathrm{True}}$	$\mathcal{Z}_{\mathrm{PLINY}}$	$\mathcal{Z}_{\mathrm{MULTINEST}}$
Ackley	-9.74	-9.75 ± 0.08	-9.64 ± 0.08
Beale	-4.23	-4.24 ± 0.06	-4.20 ± 0.05
Booth	-5.94	-5.94 ± 0.07	-5.94 ± 0.07
Goldstein–Price	-10.42	-10.43 ± 0.08	-10.42 ± 0.08
Lévi N.13	-5.64	-5.64 ± 0.07	-5.64 ± 0.07
Matyas	-2.55	-2.55 ± 0.04	-2.55 ± 0.04
McCormick	-0.83	-0.83 ± 0.04	-0.82 ± 0.04
Gaussian	0.0	0.0 ± 0.2	0.4 ± 0.2
Sphere	-17.30	-17.3 ± 0.1	-17.32 ± 0.05

MPI messages must be sent to share data in local memory. We run our test functions on two supercomputers available by peer reviewed time allocation through the DiRAC consortium to particle, nuclear and astrophysics researchers based at universities in the United Kingdom; the Complexity machine based in Leicester, UK and the COSMOS machine based in Cambridge, UK. Complexity is a typical "distributed memory" architecture supercomputing cluster comprising 272 compute nodes each with two 8–core 2.6 GHz Xeon E5–2670 processors and 128GB RAM connected by a $56\,\mathrm{Gb\,s^{-1}}$ FDR Infiniband network. By contrast, the COSMOS "shared memory" machine comprises 192 compute nodes each with an 8–core 2.6 GHz Xeon E5-4650L and 64GB RAM connected by a NUMAlink 6 cache–coherent interconnect system that makes the entire 12.2 TB of memory globally addressable and is capable of offloading MPI communication overheads from the CPUs. We run each calculation using between 1 and either 512 (COSMOS) or 1024 (Complexity) processors based on available resources and repeat each run 10 times to get the mean time–to–solution profile as a function of processor count and architecture. For all of these tests we switch off all I/O functionality including the writing of chain– and resume–files since the system I/O nodes represent a shared resource with other users which can lead to significant unpredictable bottlenecking while queuing behind other processes.

2.3 Results

Table 2.3 shows the mean calculated evidence values for each of the nine benchmark functions specified in Table 2.1 calculated from 1000 runs using unique seeds by both PLINY

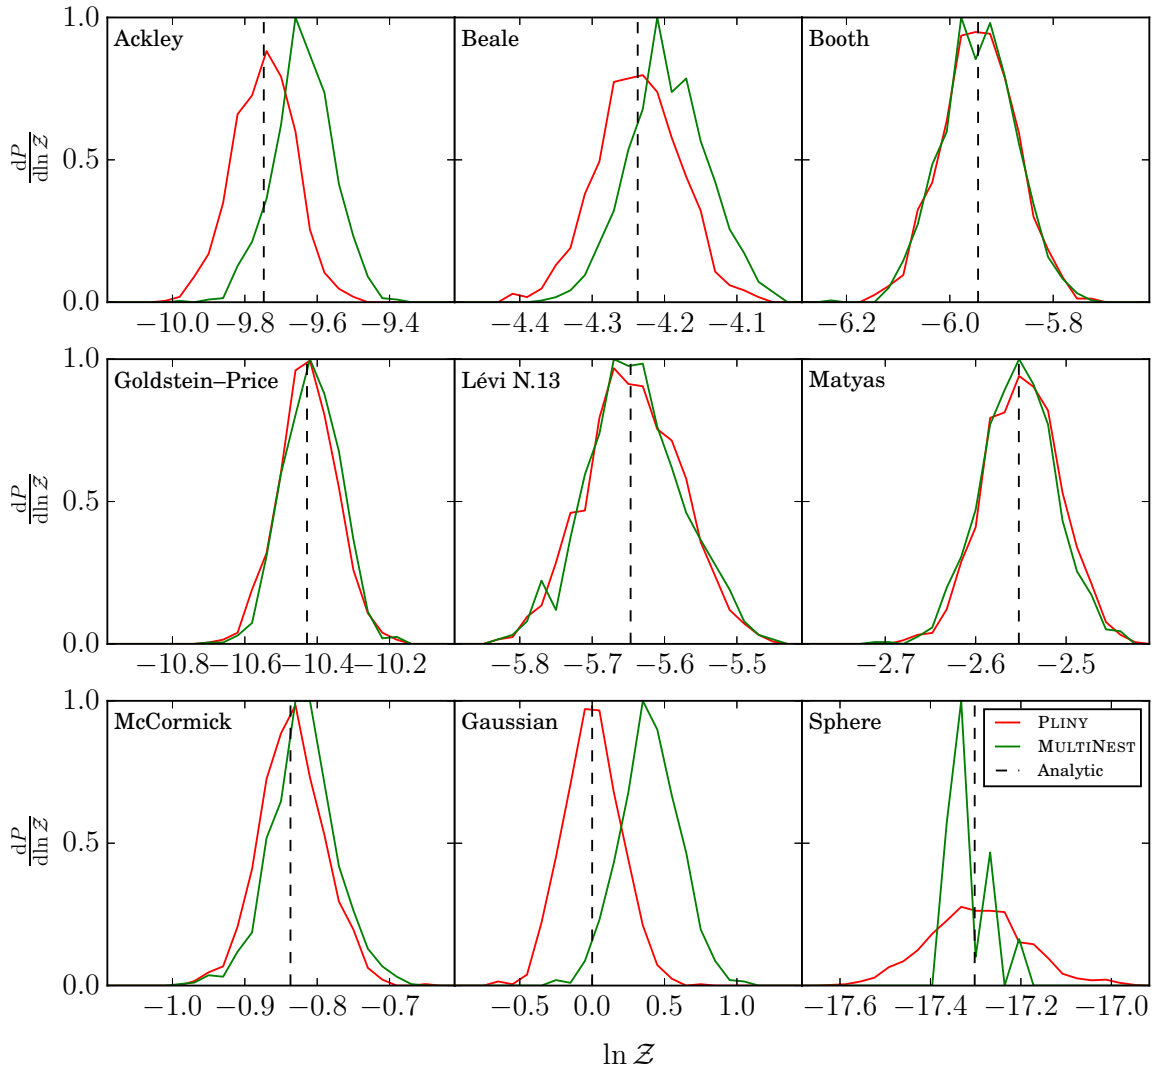

Figure 2.1. Probability distributions $dP/d\ln \mathcal{Z}$ for the natural logarithm of the evidence values $\mathcal{Z}$ calculated by Pliny (red) and MultiNest (green) for 1000 runs of each of the nine evidence benchmarks defined in Table 2.1. Distributions are normalized to a peak value of 1.0 and the analytic evidence values are shown by dashed black lines.

and MultiNest along with the analytic values. Figure 2.1 similarly shows the distributions of the evidence values calculated by both Pliny and MultiNest over the nine test functions. We see that in all cases Pliny is able to accurately calculate the evidence statistic with variations as expected from the Poisson uncertainty in the weighting of each nested sampling step (Skilling 2004). MultiNest achieves equally accurate and precise results for four of the test functions (Booth, Goldstein–Price, Lévi N.13 and Matyas) where the distributions of evidence values for both codes appear to be identical normal distributions centred on the analytic evidence value and with the statistically expected errors. This reinforces that the underlying algorithms are the same and work correctly for these functions. However, in the cases of the Ackley, Beale, McCormick and Gaussian functions, the evidence values calculated by MultiNest are seen to bias significantly higher than the analytic results. We would expect this to happen when the sampling ellipsoid fails to bound the entire region satisfying $\mathcal{L} > \mathcal{L}^*$ excluding regions near the $\mathcal{L}^*$ contour from being sampled which has two effects on the evidence integral. First, new samples are biased towards higher likelihood values near the centre of the region bounded by the critical likelihood contour. In addition the prior volume occupied by the live points shrinks faster than expected statistically leading the algorithm to weight each point in the chain higher than its true weighting. Therefore, if the sampling efficiency becomes too high the combination of inflated weights and likelihoods increases each point's contribution to the evidence integral and biases the final evidence value to higher values. Clearly ellipsoidal sampling is not the cause of this issue since Pliny is able to utilize it and correctly calculate evidence values for non–Gaussian distributions without bias. We investigate the potential cause of this effect in Section 2.4.

In addition, for the ten dimensional Sphere problem the distribution of evidence values calculated by MultiNest has three peaks. The strongest peak is biased slightly below the analytic evidence value while the other two are biased slightly high; none of the peaks coincide with the analytic value and the standard deviation across all three peaks is less than expected statistically from the Poisson uncertainty in the weighting of each nested sampling step. This behavior appears quite different from the systematic bias to higher evidence values seen in four of the other functions and it is particularly curious to see that for a symmetric unimodal posterior with ellipsoidal iso–likelihood contours the distribution should come out distinctly trimodal with widths significantly smaller than expected statistically. We make no attempt to understand or explain this behavior further but note

that again MultiNest has failed to accurately calculate the evidence value for another benchmark in which Pliny is successful.

In Figure 2.2 we show the time–to–solution performance of both Pliny and MultiNest as a function of the number of processors for the parallel performance benchmarks A, B and C i.e. a 1 dimensional Gaussian likelihood with each likelihood calculation taking an artificially slow 10^{-3}, 10^{-2} and 10^{-1} seconds respectively. For all three cases on both the distributed– and shared–memory architectures the serial time–to–solution is identical for both codes and we observe an initial linear scaling as the total computation time falls with increasing number of processors. As the wall–time spent on parallel likelihood calculations falls and the inter–processor communication costs rise with the number of processors, each code eventually reaches its peak performance beyond which the growing communication overhead dominates and the time–to–solution rises. For test A on the Complexity distributed–memory machine this peak in performance was reached at approximately 16 processors for both codes and for test C on the COSMOS shared–memory machine the performance peak was not reached by 512 processors. In the remaining four cases Pliny is able to make use of twice as many processors before reaching peak performance and in all six cases the minimum time–to–solution was significantly less than for MultiNest by up to a factor of three. Although this gap narrows as the likelihood delay time increases and differences in the parallelization schemes are masked by increased time spent calculating likelihoods, these results demonstrate the superior parallel scalability of Pliny for these test problems. As we consider the ten–dimensional equivalent problem for tests D, E and F also in Figure 2.2 we see qualitatively similar results, although the number of processors at which peak performance is achieved is typically a factor of two to four higher for both codes.

The codes also display strong performance scaling with the number of live points used. From Figure 2.3 we see that for test G (10 dimensional Gaussian with 10^{-3} s delay and only 100 live points) the serial computation time of MultiNest on the distributed memory machine using 32 processors (peak performance) is three times longer than Pliny as its acceptance rate of new points fails to reach the target rate of 39%, instead typically falling below 20%. We suggest that this is due to the covariance based bounding ellipsoid method failing to produce an accurate representation of the critical likelihood contour when the prior space is sparsely sampled by fewer live points. By comparison, Pliny's more accurate DRN ellipsoid algorithm actually leads to sampling efficiencies above the target acceptance

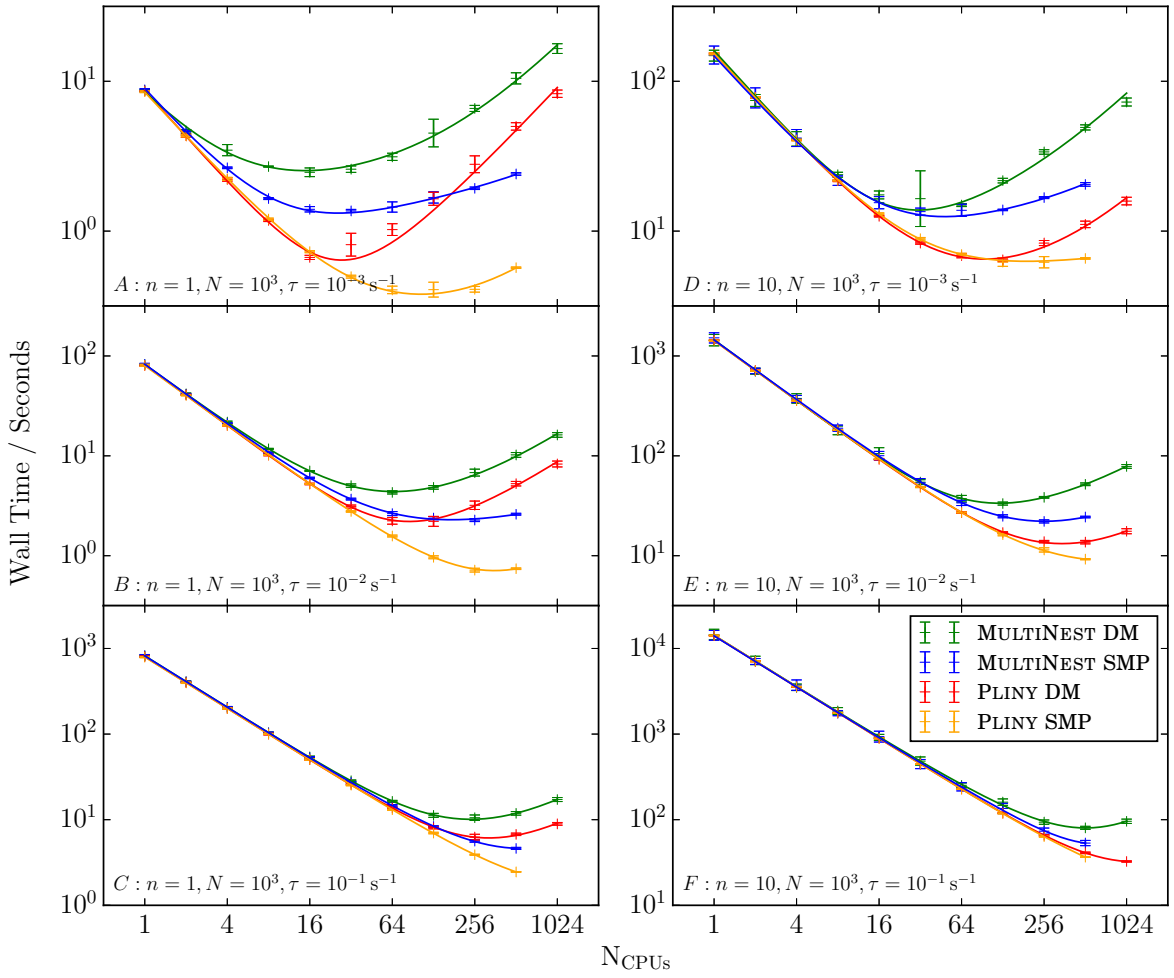

Figure 2.2. Results from the parallel performance benchmark tests A (top–left: dimensionality $n = 1$, number of live points $N = 10^3$, likelihood function delay $\tau = 10^{-3}\,\mathrm{s}^{-1}$), B (middle–left: $n = 1$, $N = 10^3$, $\tau = 10^{-2}\,\mathrm{s}^{-1}$), C (bottom–left: $n = 1$, $N = 10^3$, $\tau = 10^{-1}\,\mathrm{s}^{-1}$), D (top–right: $n = 10$, $N = 10^3$, $\tau = 10^{-3}\,\mathrm{s}^{-1}$), E (middle–right: $n = 10$, $N = 10^3$, $\tau = 10^{-2}\,\mathrm{s}^{-1}$) and F (bottom–right: $n = 10$, $N = 10^3$, $\tau = 10^{-1}\,\mathrm{s}^{-1}$). Plotted are the mean wall times for both Pliny and MultiNest to reach $\mathcal{Z} = -1$ with error bars giving the standard deviations from ten repeated runs. Each test was carried out using between 1 and 1024 processors (N_{CPUs}) on both the Complexity distributed memory system (DM, Pliny in red, MultiNest in green) and the COSMOS shared memory system (SMP, Pliny in orange, MultiNest in blue). Solid lines are maximum likelihood fits to the function $t = t_{\mathrm{s}} + t_{\mathrm{p}}/N_{\mathrm{CPUs}} + t_{\mathrm{c1}} \ln N_{\mathrm{CPUs}} + t_{\mathrm{c2}} N_{\mathrm{CPUs}}$ for the wall time t with serial component t_{s}, parallel component t_{p} and communication costs t_{c1} and t_{c2}.

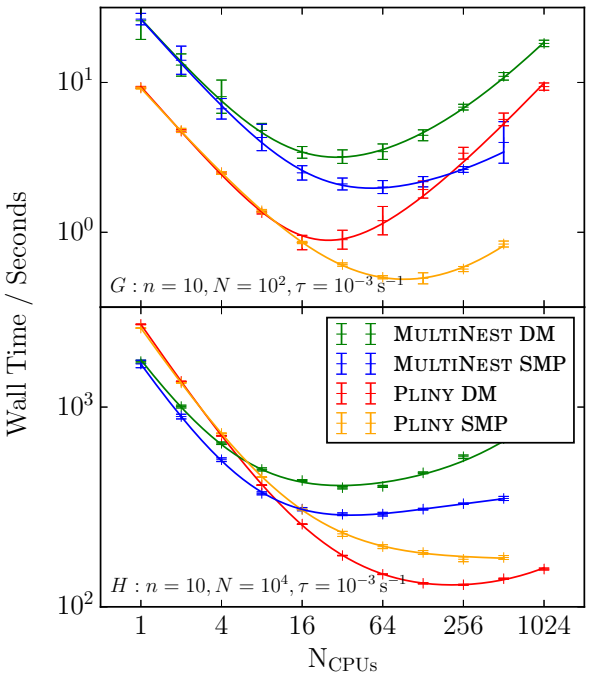

Figure 2.3. As for Figure 2.2 except for tests G (top: dimensionality $n = 10$, number of live points $N = 10^2$, likelihood function delay $\tau = 10^{-3}\,\mathrm{s}^{-1}$) and H (bottom: $n = 10$, $N = 10^4$, $\tau = 10^{-3}\,\mathrm{s}^{-1}$).

rate at close to 60% since the volume of the bounding ellipsoid is significantly smaller than that of the $\mathcal{L} > \mathcal{L}^*$ region in the sparse limit. Both codes reach peak performance at 64 processors on the shared memory machine, although again Pliny's time–to–solution is more than three times faster. For test H running with 10,000 live points, we see a very different scaling. The serial time for MultiNest is somewhat faster than Pliny due to the computation time for the covariance based bounding ellipsoid algorithm scaling more favourably than DRN for larger sets of live points. However, as the number of processors is increased Pliny reaches the solution significantly faster, reaching peak performance on the distributed memory system at 128 processors with a computation time four times faster than that of MultiNest. MultiNest performs somewhat better on the shared memory architecture but is still unable to match the scalability or peak performance of Pliny.

In Figure 2.4 we see behaviour in the 20 and 30 dimensional cases reminiscent of the 10 dimensional case except with both codes able to utilize more processors at peak performance as the dimensionality is increased. However for the 50 dimensional problem (Test K) there is some evidence for Pliny performing better for all processor counts. The time–to–solution for MultiNest is much more variable than statistically expected as in

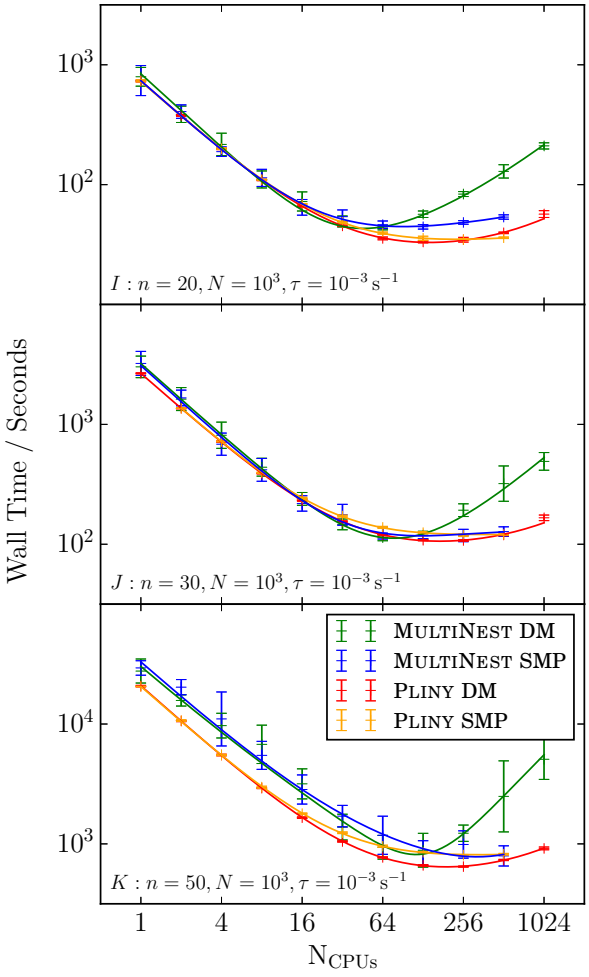

Figure 2.4. As for Figure 2.2 except for tests I (top: dimensionality $n = 20$, number of live points $N = 10^3$, likelihood function delay $\tau = 10^{-3}\,\mathrm{s}^{-1}$), J (middle: $n = 30$, $N = 10^3$, $\tau = 10^{-3}\,\mathrm{s}^{-1}$) and K (bottom: $n = 50$, $N = 10^3$, $\tau = 10^{-3}\,\mathrm{s}^{-1}$).

some cases it makes up to 30% more calls to the likelihood function than Pliny before converging. We presume that this is again an issue relating to the sparsity of live points when fitting the bounding ellipsoid in higher dimensional spaces, suggesting a fundamental limitation on the performance of covariance based ellipsoid samplers for many–dimensional problems.

2.4 Discussion

In Figure 2.1 we saw that the evidence values calculated by MultiNest for some of the benchmark functions were biased systematically higher than the correct values determined by Pliny. We suppose that the reason for this is because the volume of their sampling

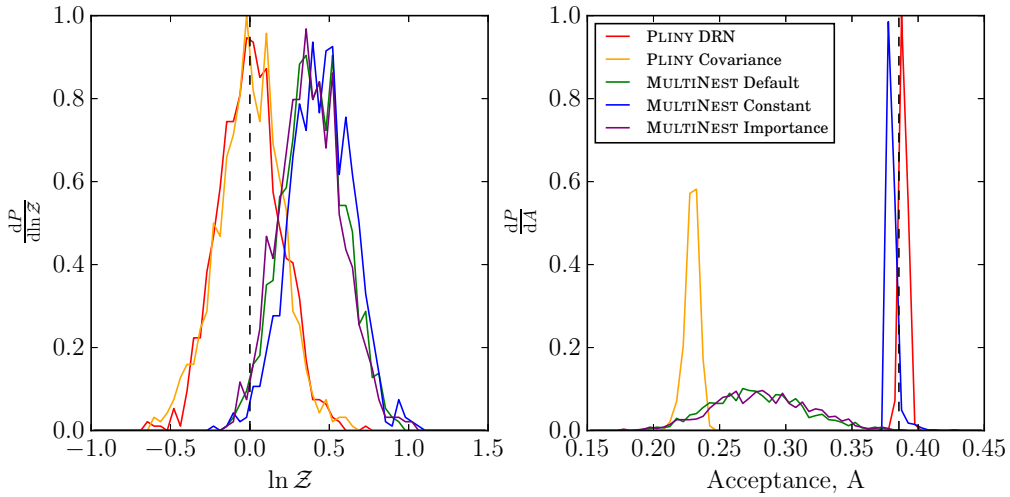

Figure 2.5. Probability distributions $\mathrm{d}P/\mathrm{d}\ln\mathcal{Z}$ and $\mathrm{d}P/\mathrm{d}A$ for the natural logarithm of the evidence values $\mathcal{Z}$ (left) and operational acceptance rates A (right) for 1000 runs of the Gaussian benchmark function in Table 2.1 calculated by PLINY (DRN ellipsoid algorithm in red, covariance ellipsoid algorithm in orange) and MULTINEST (default in green, constant efficiency in blue, importance nested sampling in purple). Distributions are normalized to a peak value of 1.0 and the analytic evidence value and input target acceptance rate are shown by dashed black lines.

ellipsoid follows the statistical expectation for the prior volume X occupied by the live points while in reality that volume fluctuates significantly. It is possible that at some steps in the algorithm, the sampling ellipsoid does not bound the entire likelihood, biasing new live points to higher likelihoods and weightings in the evidence integral. MULTINEST's constant efficiency mode instead fixes the ellipsoid volume at each iteration to ensure an acceptance rate as close as possible to the target value. Similarly, the importance sampling mode claims to be able to increase the accuracy of the evidence integral via a pseudo–importance–sample including points rejected under the likelihood constrained sampling scheme (Feroz et al. 2013). To investigate this we repeat the calculation of the Gaussian function in Table 2.1 (for which we see the largest discrepancy) running MULTINEST in its default, constant efficiency and importance sampling modes separately 1000 times each to see how the distributions of evidences and acceptance rates vary and compare with PLINY. We also implement the covariance based ellipsoid fitting algorithm from Shaw et al. (2007) inside PLINY for an additional comparison. The distributions for the evidence values and operational acceptance rates are shown in Figure 2.5. It is instantly clear that in none of the three modes is MULTINEST able to accurately calculate

the correct evidence value. Without constant efficiency mode, the operational acceptance rate of MultiNest has significant scatter. However it is always lower than in the constant efficiency mode, meaning that the default sampling ellipsoids are larger and therefore more likely to bound the entire critical likelihood contour and be free from bias yet this is not reflected in the distributions of evidences. Covariance based ellipsoid fitting is also not the root cause since Pliny is able to use such an algorithm to produce an accurate evidence distribution in spite of working at a low acceptance. All we can conclude is that there is some unidentified feature intrinsic to the MultiNest implementation that can in some cases lead to a significant bias towards higher likelihoods. In addition we note that the importance sampling mode fails to make any noticeable difference to either the peak or width of the output evidence distribution and its performance appears identical to the default MultiNest algorithm.

In assessing the parallel performance in Section 2.3 we had used a fixed target acceptance rate across the two codes to create a like–for–like comparison. However, in light of the apparent evidence bias it is likely that for many of those tests MultiNest will have failed to calculate the correct evidence values. It appears important in evaluating peak performances to consider the maximum target acceptance rate at which each code can calculate accurate evidence values for a fixed problem. In Figure 2.6 we show the output evidence values for 10^4 Pliny and MultiNest runs of the Gaussian benchmark function in Table 2.1 plotted as a function of a randomly input target acceptance rate between 0.1 and 1.0 and also as a function of the number of likelihood evaluations required to converge. The target acceptance represents the user input for the code while the number of likelihood evaluations is a proxy for the time–to–solution in the limit of computationally slow likelihood functions. We see that Pliny is able to operate at optimal accuracy (less than one percent bias in $\ln \mathcal{Z}$) for this problem at target acceptance values of less than 0.6 while MultiNest does not reach the same level of accuracy until the target acceptance is reduced all of the way to 0.1. For reference the target acceptance used for the ten dimensional Gaussian in Figures 2.1 and 2.2 was 0.3855. In terms of number of likelihood evaluations, Pliny is able to achieve the same level of accuracy as MultiNest with a factor of four fewer evaluations. In the regime where slow likelihood functions dominate ellipsoid computation times and communication overheads this equates to a significant speedup. Pliny is always at least as accurate as MultiNest except for target acceptance rates approaching one, although that regime represents a poor choice of user input and

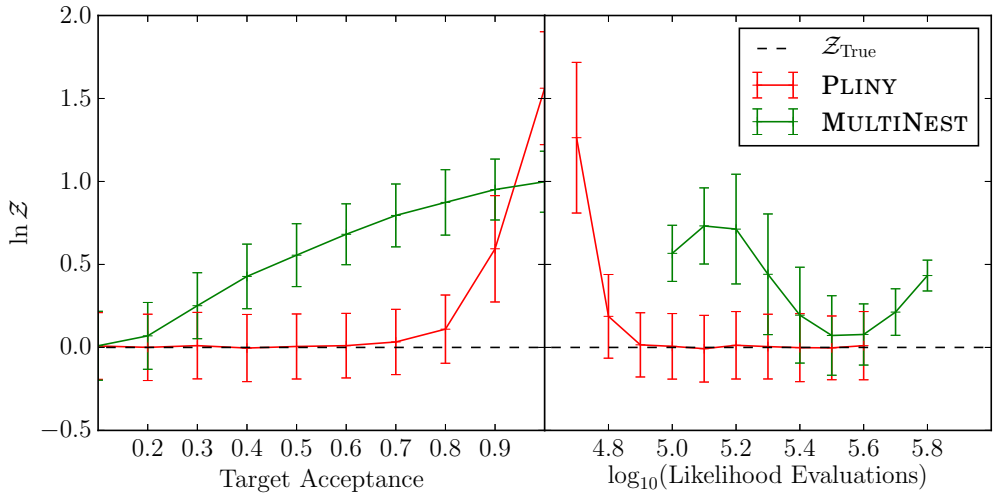

Figure 2.6. Binned evidence values $\mathcal{Z}$ as a function of the input target acceptance and the number of likelihood evaluations required to reach convergence for 10^4 runs of the Gaussian benchmark function for both PLINY and MULTINEST. The analytic evidence value $\mathcal{Z}_{\rm True}$ is plotted as a dashed black line. PLINY is seen to reach its peak accuracy in $\ln \mathcal{Z}$ (less than one percent bias) for target acceptances of less than 0.6 while MULTINEST requires less than 0.1 for the same accuracy, equating to a factor of four more likelihood evaluations.

MULTINEST also fails to calculate the correct evidence for those target acceptance values.

In light of the two codes having different target acceptance thresholds for calculating the evidence accurately we conduct one further test of their relative performance when working at their optimal acceptance rates. Although in general the optimal target acceptance rate is not known before starting a new calculation, it could always be estimated in a manner similar to Figure 2.6 assuming a Gaussian likelihood with dimensionality and number of live points matching the problem. Figure 2.7 shows the parallel time–to–solution performance for PLINY and MULTINEST working at target acceptance rates of 0.6 and 0.1 respectively on Test F from Table 2.2. We recall that in Figure 2.2 the two codes performed identically when working with the same acceptance rate of 0.3855 although the evidence values calculated by MULTINEST were biased high. However when working at its superior optimal acceptance rate, PLINY outperforms MULTINEST by almost an order of magnitude for any given number of processors that we tested.

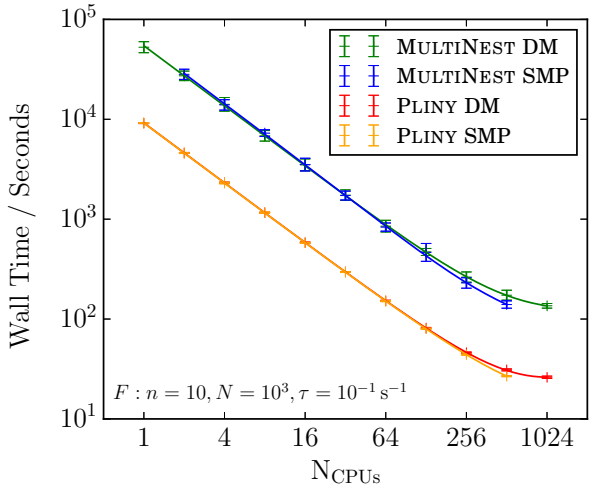

Figure 2.7. As in Figure 2.2 for test F (dimensionality $n = 10$, number of live points $N = 10^3$, likelihood function delay $\tau = 10^{-1}\,\mathrm{s}^{-1}$) but with the optimal acceptance values of 0.6 for Pliny and 0.1 for MultiNest as determined from Figure 2.6. Pliny's time–to–solution is seen to be up to an order of magnitude faster than MultiNest in all cases.

2.5 Conclusions

We have developed Pliny, an implementation of the nested sampling algorithm utilizing an accurate bounding ellipsoid algorithm for ellipsoidal sampling and an optimized parallelization scheme. In this paper we compared Pliny to the current standard software MultiNest for a range of performance metrics. We have shown that Pliny is capable of calculating the Bayesian evidence integral for a range of unimodal posteriors to a higher degree of accuracy than MultiNest. Pliny's algorithm also scales more effectively over multiple processors for computationally fast likelihood functions, being able to effectively utilize higher numbers of processors on both distributed– and shared–memory supercomputing architectures and achieving superior peak performance before becoming communication bound. While we are unable to explain an apparent bias in MultiNest to overestimate the evidence values of some functions, we show that it leads to Pliny being able to sample new live points with a lower rejection rate in higher dimensional problems, improving the time–to–solution vs evidence accuracy trade–off for computationally expensive likelihood functions and allowing it to evaluate accurate evidence integrals up to an order of magnitude faster than MultiNest for a ten dimensional Gaussian likelihood. Although MultiNest features multimodal decomposition to increase the sampling efficiency for degenerate and multi–peaked posterior distributions, we claim the fact that

it is based on an inherently biased sampling mechanism makes it unreliable. In general, nested sampling does not adequately address the curse of dimensionality since although it casts the evidence integral into a one–dimensional form, samples are still drawn from the prior space in a manner that is limited by the ability of algorithms to reconstruct a sparsely sampled critical likelihood contour in multidimensional space and currently can only be overcome with artificially low sampling efficiencies.

Chapter 3

Probing the Physical State of Diffuse Gas Towards G10.6–0.4 with Nested Sampling

The work presented in this chapter is based on the paper by Rollins et al., submitted to MNRAS, in collaboration with J. M. C. Rawlings.

Observations of the chemical composition of diffuse interstellar gas represent a unique and powerful probe of its physical state and consequently its contribution to the total gas budget of the interstellar medium as well as the cycling of gas and dust by dynamical processes in galaxies. Indeed, recent observations by the *Herschel Space Observatory* have revealed chemically rich diffuse gas at intermediate densities of $1 - 10\,\mathrm{cm}^{-3}$ and temperatures of 100 K between the extremes of molecular clouds and the hot ionized medium in a range of chemically rich environments (Gerin & PRISMAS Consortium 2010, Gerin et al. 2010, Godard et al. 2010, Neufeld et al. 2010*a*, Neufeld et al. 2010*b*). The gravitational collapse and cooling of diffuse clouds to form clumps, molecular clouds and ultimately stars, as triggered by stellar outflows and radiation pressure (Hollenbach et al. 1976, Koenig et al. 2008), is a process that is intrinsically coupled to the molecular content of the gas (Hartquist et al. 1993). Models suggest that the cycling process into and out of dense

molecular clouds can itself lead to sensitive chemical enrichment signatures of the dynamics at work (Price et al. 2003). Insights into the physical, dynamical and chemical state of diffuse gas are therefore important in understanding the formation of stars and evolution of the interstellar medium on galactic scales.

However, recent observations of certain molecular species raise questions about a range of processes at work in the gas for which chemical models have no simple answer. An example is the large observed column densities of the nitrogen hydrides (ammonia, NH_2 and NH, Persson et al. 2010, Persson et al. 2012) where the primary formation mechanism is still debated. On one hand, Persson et al. (2014) argue from the results of their gas–grain chemical modelling that both progressive hydrogenation on dust mantles and non–thermal desorption are necessary to form the observed relative abundances of nitrogen hydrides. In contrast, the gas–phase chemical model of Faure et al. (2013) considering the ortho and para spin states of the nitrogen hydrides suggests that the observed ortho–to–para ratio (OPR) of ammonia (Persson et al. 2012) can be accounted for by the dissociative recombination of NH_4^+ formed from para–rich molecular hydrogen in dense gas at high extinctions. Furthermore, Persson et al. (2010) modelled the gas towards G10.6–0.4 as a photon–dominated region (PDR) in which NH_4^+ is formed by successive condensation reactions of H_2 with photoionized NH. Although their model yielded correct relative abundance ratios between the nitrogen hydrides, the absolute abundances were all less than those observed. The overabundance of CH^+ and SH^+ with respect to chemical models also presents a challenge. In hot gas, their primary formation is by the ion–neutral reactions of ionized carbon or silicon with molecular hydrogen. However, these reactions are highly endothermic with activation temperatures of 4640 K and 9860 K respectively (McElroy et al. 2013). Since diffuse gas clouds are significantly cooler than those temperatures, it is often argued that higher effective temperatures are needed, such as those found in magnetically turbulent or shocked gas where a significant drift velocity between ionic and neutral species occurs (Federman et al. 1996, Godard et al. 2014, Chieze et al. 1998, Flower & Pineau des Forêts 1998, Lesaffre et al. 2013). Similarly, oxygen chemistry and the formation of OH^+, H_2O^+ and H_3O^+ is believed to be initiated by charge exchange of atomic oxygen with cosmic ray ionized hydrogen followed by progressive reactions with molecular hydrogen. The chemistry is therefore sensitive to the local cosmic ray ionization rate in initiating this chain of reactions (Hollenbach et al. 2012, Indriolo et al. 2015). Similar to CH^+ and SH^+, the temperature dependence of the charge exchange between oxygen and H^+ means

that any relative motions between the ion and neutral fluids would also promote oxygen chemistry. While specialist models specific to each of these chemical mechanisms exist and are individually successful at reproducing observations of individual species, the goal of a global model capable of simultaneously reproducing all observations presents a computational challenge with a large parameter space that must be explored to determine which chemical mechanisms are truly at work.

When making a model of any scientific phenomenon, one of the main goals is to find combinations of the model's free parameters that can best reproduce data from observations or experiments. For astrochemical modelling, the data may be spectroscopic molecular line profiles, column densities or relative molecular abundances, while the physical state of the source, such as its density and temperature, is unknown and represents the input parameters of the model. In many previous works, the range and combination of those parameters considered was chosen by the author of the work, typically a uniform grid of points spanning the parameter space. The best fitting model is then often identified based on a maximum likelihood criterion, and although more detailed studies are typically constrained by available computational resources, the results of such a naive approach are often unsatisfactory. Firstly, the parameter search is biased to those values chosen by the author of the work. This becomes particularly problematic where there are many free parameters or the model output is sensitive to one or more of those parameters, as in these cases it becomes less likely that the true maximum likelihood solution lies close to any of the chosen parameter combinations. Secondly, the optimal solution is often not unique and a range of degenerate and possibly multimodal parameters can all provide equally good representations of the data. A poorly designed grid combined with the selection of only the single maximum likelihood point will fail to identify such detail. Finally, although the authors have control over the parameters considered, there is no further way to incorporate prior knowledge from previous studies into the results under the maximum likelihood selection procedure.

A superior approach is to employ an algorithm that can search the parameter space for you in an unbiased way. Many such algorithms exist that can do such a search within the framework of Bayesian inference, including the Metropolis–Hastings (Metropolis et al. 1953, Hastings 1970) and simulated annealing (Kirkpatrick et al. 1983, Černý 1985) Markov chain Monte Carlo (MCMC) algorithms. The framework allows for the incorporation of prior knowledge and the determination of uncertainties in the inferred parameters and statistics

for the comparison with other models. Such methods are well established and have found many applications in the field of astrophysics (e.g. astrophysical radiative transfer, Keto et al. 2004, absorption line analysis, Schilke et al. 2014, cosmic microwave background cosmology, Hawken et al. 2012, exoplanetary atmospheres, Waldmann et al. 2015). Of particular note, the study by Makrymallis & Viti (2014) demonstrates the application of such methods to the study of interstellar ices, allowing the authors to make inferences about the physical state of the gas in dark molecular clouds using the UCL_CHEM chemical model. In this work, we have developed a fast, time–dependent chemical model for both the quiescent and active phases in diffuse gas and featuring approximations to a range of chemical processes that have been proposed as important in the literature. We then coupled the model with the nested sampling MCMC algorithm in order to match observations of a range of molecules in diffuse gas towards G10.6–0.4 (W31C) and statistically infer a range of physical parameters associated with that gas, as well as the relative importance of the various chemical mechanisms. The nested sampling algorithm is summarised in Section 3.1, while the chemical observations and model are described in Section 3.2 and the results and analysis are presented in Sections 3.3 and 3.4 respectively. We finally offer our conclusions and proposals for further work in Section 3.5.

3.1 Nested Sampling

In the domain of Bayesian statistics, the main objective is frequently to characterize the posterior probability distribution of a vector of free parameters $\boldsymbol{\theta}$ given a model $\mathcal{M}$ and a dataset $\mathcal{D}$, denoted $P(\boldsymbol{\theta}|\mathcal{D},\mathcal{M})$. The posterior can be expressed through Bayes' theorem as a function of the likelihood of data for a given model, $\mathcal{L} = P(\mathcal{D}|\boldsymbol{\theta},\mathcal{M})$, the prior probability on the set of parameters, $\pi = P(\boldsymbol{\theta}|\mathcal{M})$, and the Bayesian evidence of the data, $\mathcal{Z} = P(\mathcal{D}|\mathcal{M})$:

$$P(\boldsymbol{\theta}|\mathcal{D},\mathcal{M}) = \frac{P(\mathcal{D}|\boldsymbol{\theta},\mathcal{M})P(\boldsymbol{\theta}|\mathcal{M})}{P(\mathcal{D}|\mathcal{M})} = \frac{\mathcal{L}(\boldsymbol{\theta})\pi(\boldsymbol{\theta})}{\mathcal{Z}}. \tag{3.1}$$

The evidence is an important statistical quantity since it arises in the calculation of Bayes

factors for the comparison of two competing models $\mathcal{M}^{\mathrm{A}}$ and $\mathcal{M}^{\mathrm{B}}$ of the same dataset:

$$\frac{P(\mathcal{M}^{\mathrm{A}}|\mathcal{D})}{P(\mathcal{M}^{\mathrm{B}}|\mathcal{D})} = \frac{P(\mathcal{D}|\mathcal{M}^{\mathrm{A}})P(\mathcal{M}^{\mathrm{A}})}{P(\mathcal{D}|\mathcal{M}^{\mathrm{B}})P(\mathcal{M}^{\mathrm{B}})} = \frac{\mathcal{Z}_1}{\mathcal{Z}_0}\frac{P(\mathcal{M}^{\mathrm{A}})}{P(\mathcal{M}^{\mathrm{B}})}. \tag{3.2}$$

Since it is independent of $\boldsymbol{\theta}$, $\mathcal{Z}$ can be interpreted as the normalization factor for the posterior probability distribution:

$$\mathcal{Z} = \int \mathcal{L}(\boldsymbol{\theta})\pi(\boldsymbol{\theta})\,\mathrm{d}\boldsymbol{\theta}. \tag{3.3}$$

This is a multi–dimensional integral over all of the free parameters making up the vector $\boldsymbol{\theta}$. Such integrals can be numerically challenging, suffering from the "curse of dimensionality" as the number of free parameters gets large, and are traditionally best tackled with MCMC integration algorithms. The nested sampling algorithm (Skilling 2004) was devised primarily to tackle the calculation of these evidence integrals within the framework of Bayesian statistics. The algorithm is initialized by drawing a random sample of N "live points" from the prior distribution $\pi(\boldsymbol{\theta})$ and calculating their likelihoods. The minimum of all likelihoods from the set of live points is found and defined as the likelihood threshold $\mathcal{L}^*$. Further points are then drawn sequentially and their likelihoods are calculated until a point is found with likelihood $\mathcal{L} > \mathcal{L}^*$. The new point is added to the set of live points while the lowest likelihood point is removed and used to start an ordered MCMC chain of points, $\mathcal{C}$. This process of determining the likelihood threshold, sampling a replacement point and moving the threshold point to the chain is then repeated until convergence. Because the expected shrinkage of the prior volume with each iteration is known statistically, the prior weight w_{i} and likelihood $\mathcal{L}_{\mathrm{i}}$ associated with the i'th point in $\mathcal{C}$ are known and the evidence calculation is approximated by a weighted sum over likelihoods:

$$\mathcal{Z} = \sum_{\mathrm{i}\in\mathcal{C}} w_{\mathrm{i}}\mathcal{L}_{\mathrm{i}}. \tag{3.4}$$

Furthermore, the output chain of points has the property that it is a weighted sample from the posterior distribution. Therefore, for any property $Q(\boldsymbol{\theta})$ derived from the input parameters we can calculate $\bar{Q}$, the expectation value of the posterior distribution of Q as

a weighted sum over the chain:

$$\bar{Q} = \sum_{\mathrm{i}\in\mathcal{C}} \frac{w_\mathrm{i}\mathcal{L}_\mathrm{i}}{\mathcal{Z}} Q(\boldsymbol{\theta}_\mathrm{i}). \tag{3.5}$$

One of the key challenges in implementing the algorithm is finding an efficient way to sample points satisfying $\mathcal{L} > \mathcal{L}^*$, especially for problems with a high number of free parameters. Mukherjee et al. (2006) fit an ellipsoid to the live points as an approximation to the likelihood threshold and sampled new points from within it. Feroz et al. (2009) extended the idea by clustering the live points and fitting multiple ellipsoids to deal with multimodal posterior distributions and non-linear degeneracies in their implementation MultiNest. For our work we choose to use the Pliny implementation of Rollins et al. (submitted). Their code overcomes the relatively poor acceptance rate of candidate points above the likelihood threshold in many nested sampling implementations with an accurate minimum volume covering ellipsoid algorithm (Sun & Freund 2004) and an efficient parallelization scheme to scale the calculation over distributed memory computer architectures. To achieve a workable balance between computational speed and accuracy, we sample using 200 live points, a target acceptance rate of 70% and an evidence convergence tolerance of 1%.

3.2 Model for the gas towards G10.6–0.4

G10.6–0.4 is an extended HII region of high–mass star formation with a number of molecular clouds along its line of sight (Corbel & Eikenberry 2004). It was a target object for the *Herschel* PRISMAS (PRobing InterStellar Molecules with Absorption line Studies) key time program and is used extensively for absorption line studies of diffuse gas in the interstellar medium (Godard et al. 2012, Persson et al. 2014, Indriolo et al. 2015). In this section, we discuss the selection of G10.6–0.4 as a suitable target for our study followed by a detailed description of our model which consists of three components. First, we identify a dataset $\mathcal{D}$ of molecular column densities with errors that we wish to try and match, allowing us to write a likelihood function $\mathcal{L}$ for our model. Secondly, we create a parametrized description of the physical state of the diffuse gas cloud; this defines the parameter space that the nested sampler will explore and determine the posterior distribution

over. Finally, we build a model $\mathcal{N}$ that takes those parameters as input and calculates the chemical state of the gas cloud, outputting modelled counterparts for the observed values in $\mathcal{D}$. We describe these three components in Sections 3.2.1, 3.2.2 and 3.2.3 below.

3.2.1 Observations

Observations made by the HIFI instrument on board *Herschel* provide a rich source of data describing the chemical state of galactic diffuse clouds. In particular, the PRISMAS key program (Gerin & PRISMAS Consortium 2010) aimed to characterize a range of molecular hydrides along eight lines of sight in absorption against background star–forming regions. This wealth of data is ideal for constraining the input parameters of a chemical model and a Bayesian technique such as nested sampling allows us to do it in an efficient and unbiased way. Many of the absorption line profiles have been converted to column densities and relative abundances, meaning that it is sufficient to create a model for the chemical structure of the gas and computationally costly radiative transfer modelling of a synthetic signal is not necessary.

However, not all such data in the literature are suitable for our study. Critically, evaluating the likelihood function $\mathcal{L}$ in the nested sampling algorithm requires that the errors on the observational data are known. While a range of papers conduct thorough analyses to determine the column densities of a range of molecular species in diffuse clouds, a significant proportion fail to make any attempt at quantifying their uncertainties leaving their results unsuitable for this work. Furthermore, HIFI's excellent spectral resolution allows us to see that these diffuse clouds typically comprise multiple independent gas components spanning a wide range of velocities (e.g. Godard et al. 2012, Persson et al. 2012). To remain confident that published data from different studies does in fact relate to the same gas, we further restrict ourselves to only using observations where the molecular column densities are resolved into velocity bins and all of our target species have firm detections across our selected velocity range. It could be argued that selecting which data we will fit to in this way introduces a potential bias into the final solution of our model. However, the restrictions of our chemical model (single–zone) and the likelihood function (requires firm detections) mean that using only a restricted velocity range for a single object where all of our target species are detected together is essential for the working of our model. Furthermore, we need to maximise the number of observed chemical species in order to probe as wide a range of chemical mechanisms as possible and hence provide

Table 3.1. Observed velocity integrated column densities and their fractional errors for a range of species in absorption towards G10.6–0.4 between $15\,\mathrm{km\,s^{-1}}$ and $46\,\mathrm{km\,s^{-1}}$. For the undetected species NH^+ an upper limit on the column density is presented but not used in any calculations. References: *a.* Persson et al. (2012), *b.* Godard et al. (2012), *c.* Indriolo et al. (2015).

Species	Column Densities, $\mathcal{D}$	Fractional Error	Reference
NH	$1.9 \times 10^{14}\,\mathrm{cm^{-2}}$	0.20	*a*
NH_2	$1.1 \times 10^{14}\,\mathrm{cm^{-2}}$	0.25	*a*
NH_3	$9.2 \times 10^{13}\,\mathrm{cm^{-2}}$	0.30	*a*
SH^+	$1.5 \times 10^{13}\,\mathrm{cm^{-2}}$	0.15	*b*
CH^+	$8.0 \times 10^{14}\,\mathrm{cm^{-2}}$	0.15	*b*
OH^+	$2.6 \times 10^{14}\,\mathrm{cm^{-2}}$	0.10	*c*
H_2O^+	$6.0 \times 10^{13}\,\mathrm{cm^{-2}}$	0.20	*c*
H_3O^+	$1.1 \times 10^{13}\,\mathrm{cm^{-2}}$	0.50	*c*
NH^+	$< 1.6 \times 10^{12}\,\mathrm{cm^{-2}}$	–	*a*

tight constraints on all of the physical parameters describing the gas in our model.

Given these conditions, we identified the diffuse gas towards G10.6–0.4 (W31C) in the local standard of rest velocity range $15\,\mathrm{km\,s^{-1}}$ to $46\,\mathrm{km\,s^{-1}}$ as a suitable target object for our model. Column densities are available in the literature for the species NH, ortho–NH_2, ortho–NH_3 (Persson et al. 2012), SH^+, CH^+ (Godard et al. 2012), OH^+, ortho–H_2O^+ and H_3O^+ (Indriolo et al. 2015). These eight species represent the complete set of constraints on our model. Upper limits on the column density of the undetected species NH^+ have also been determined by Persson et al. (2014) which although we cannot use in our likelihood calculation can still be checked against our model outputs for verification. We sum the observed velocity resolved column densities across the specified velocity range to get total column densities for each species which we summarise in Table 3.1. For ammonia, we combined the quoted column densities for the ortho states and the OPR in Persson et al. (2012) of 0.6 ± 0.2 to give a total column density for both the ortho and para states since our chemical model does not distinguish spin states. Similarly for NH_2 and H_2O^+, we take the observed column densities of the ortho states and predicted OPRs of 2.3 (Faure et al. 2013) and 3.0 (Indriolo et al. 2015) respectively to determine their total column densities of ortho and para states. Our uncertainties for those species reflect the uncertainties on both the column densities of the ortho states and the OPRs. As we shall discuss in Section 3.2.3, the physical processes leading to the formation of the eight species are distinct and depend on different properties of the gas. We therefore expect that matching our model

to all eight of them simultaneously will allow for a good constraint on the physical state of the diffuse clouds.

In all cases we assume that the observed column densities are independently log–normal distributed. The likelihood function, $\mathcal{L}$, is then a product over the individual Gaussian errors:

$$\mathcal{L}(\boldsymbol{\theta}) = \prod_{\mathrm{i}\in\mathcal{S}} \frac{1}{\sqrt{2\pi\sigma_{\mathrm{i}}^2}} \exp\left(-\frac{1}{2}\left[\frac{log(\mathcal{N}_{\mathrm{i}}(\boldsymbol{\theta})) - log(\mathcal{D}_{\mathrm{i}})}{\sigma_{\mathrm{i}}}\right]^2\right) \tag{3.6}$$

where $\mathcal{S}$ represents the set of observed chemical species, $\mathcal{N}_{\mathrm{i}}(\boldsymbol{\theta})$ is the modelled column density of species i for the vector of input parameters $\boldsymbol{\theta}$, $\mathcal{D}_{\mathrm{i}}$ is the corresponding observed column density and σ_{i} is its fractional error.

3.2.2 Physical Model

We model the gas towards G10.6–0.4 as being in one of two phases; quiescent or active. These are comparable to the ambient and active phases in the three phase models of Godard et al. (2009) and Godard et al. (2014). Each phase is modelled by a single point representative of all the gas in that phase. In both phases the gas density, cosmic ray ionization rate, strength of the incident ultraviolet (UV) radiation field relative to the standard interstellar radiation field of Draine (1978) and the local effective extinction from the incident radiation field (n, ζ, G_0 and $A_{\mathrm{v}}^{\mathrm{eff}}$ respectively) are the same. In the quiescent phase there is no ion–neutral decoupling, whereas in the active phase there is a relative drift between ionic and neutral species with velocity u_{D}, presumed to be due to magnetohydrodynamic (MHD) turbulent vortices or shocks, leading to additional heating and enhanced effective temperatures for ion–neutral reactions as detailed in Section 3.2.3. The fraction of gas in the active phase is denoted by f_{A} and the total extinction of gas (as a proxy for total column density) is $A_{\mathrm{v}}^{\mathrm{Tot}}$. Together, these seven parameters comprise a vector $\boldsymbol{\theta} = (n, A_{\mathrm{v}}^{\mathrm{eff}}, G_0, \zeta, u_{\mathrm{D}}, f_{\mathrm{A}}, A_{\mathrm{v}}^{\mathrm{Tot}})$ that provides a complete description of the physical state of the gas towards G10.6–0.4. This simple geometry implicitly neglects any three–dimensional structure in the gas, including clumping and velocity structure, assuming instead that all gas in either the quiescent or active phase is identical. However, such approximations are important to ensure that the computation time for an individual model is relatively quick, allowing us to run many thousands of realizations (as is required

by the nested sampling algorithm) in a reasonable time. We also note that the presence of all eight species across the entire velocity range with similar line profiles supports the simplified single point geometry as a representation of many identical low extinction clouds with gas in two phases making up the complete signal.

The nested sampling algorithm demands that we also specify $\pi(\boldsymbol{\theta})$, the probability distribution representing our prior knowledge of the parameters based on previous studies. Unfortunately, our knowledge of the physical state of the diffuse gas towards G10.6–0.4 and indeed diffuse clouds in general is somewhat uncertain. Modest densities of the order $10 - 100\,\mathrm{cm}^{-3}$, warm temperatures of approximately 100 K and extinctions in the range 0.1–0.4 are typically ascribed to diffuse interstellar gas, although models exist where these values vary by up to an order of magnitude (e.g. Godard et al. 2014). In spite of the low extinctions of diffuse clouds, it has been shown that the total extinction of diffuse gas towards G10.6–0.4 is approximately 15 magnitudes (Corbel & Eikenberry 2004, Neufeld et al. 2010*b*). A number of works have suggested that the absorption profiles of various molecular species are made up of between 5 and 12 individual gas components with different systemic velocities, depending on the species observed (e.g. Godard et al. 2012, Persson et al. 2012). Such velocity features also appear frequently in other clouds e.g. towards W49N (Neufeld et al. 2010*a*). Attempts have been made to model the chemical composition of diffuse gas based on UV radiation fields up to 100 times stronger than the Draine radiation field (Persson et al. 2010, Neufeld et al. 2010*a*, Hollenbach et al. 2012) with some success. From observations of OH^+, Gerin et al. (2010) were able to place a conservative lower bound on the cosmic ray ionization rate to density ratio in G10.6–0.4 of $\zeta/n = 1.8 \times 10^{-19}\,\mathrm{cm}^3\,\mathrm{s}^{-1}$. More recently Indriolo et al. (2015) determined a somewhat higher cosmic ray ionization rate in G10.6–0.4 of $\zeta = 3.4 \times 10^{-16}\,\mathrm{s}^{-1}$ and Neufeld et al. (2010*a*) found a similarly enhanced rate of $\zeta = 2.4 \times 10^{-16}\,\mathrm{s}^{-1}$ in diffuse gas towards W49N, both more than an order of magnitude higher than typically assumed interstellar values. Recent models of turbulent vortices in the diffuse interstellar medium implied a turbulent gas temperature of up to 800 K (Godard et al. 2009) and constrained the ion–neutral drift velocity to less than $3.5\,\mathrm{km\,s}^{-1}$ (Godard et al. 2014). By comparison, models of C–shocks in the diffuse interstellar medium have been modelled with a range of shock velocities as high as $8 - 10\,\mathrm{km\,s}^{-1}$ and corresponding ion–neutral drift velocities of $1 - 2\,\mathrm{km\,s}^{-1}$ (Gredel et al. 2002, Flower & Pineau des Forêts 1998) as well as elevated temperatures in excess of 1000 K due to dissipative heating (Chieze et al. 1998). Godard et al. (2009)

Table 3.2. The ranges and distributional forms of the prior functions on each of the free model parameters. Note that ζ_0 is taken to be the value $1.6 \times 10^{-17}\,\mathrm{s}^{-1}$. Uniform means that the prior on the parameter is uniformly distributed between the minimum and maximum values. Logarithmic means that the prior on the logarithm of the parameter is uniformly distributed between the minimum and maximum values.

Parameter	Minimum	Maximum	Prior
$n\,/\,\mathrm{cm}^{-3}$	10^0	10^3	Logarithmic
$A_\mathrm{V}^\mathrm{eff}$	0.1	1.0	Uniform
G_0	10^{-1}	10^2	Logarithmic
$\zeta\,/\,\zeta_0$	10^{-1}	10^2	Logarithmic
$u_\mathrm{D}\,/\,\mathrm{km\,s}^{-1}$	0.0	5.0	Uniform
f_A	10^{-3}	10^{-1}	Logarithmic
$A_\mathrm{V}^\mathrm{Tot}$	10	30	Uniform

also considered the fraction of gas contained in turbulent vortices in the range $0.3 - 2\%$ as motivated by simulations (e.g. Pety & Falgarone 2000, Pan & Padoan 2009). We reflect these uncertainties in our choice of prior function we use with the nested sampling algorithm. The ranges and functional forms of the prior function for the seven parameters are given in Table 3.2.

3.2.3 Chemical Model

As we have discussed, a wide range of chemical mechanisms and pathways are expected to be important in the chemistry of the eight modelled species. This variety is important in being able to constrain all of the parameters describing the physical state of the gas outlined in Section 3.2.2. However, it also demands a sufficiently complex chemical model to be able to describe all of the processes at work. In order that our chemical model is fast enough to be able to be run many thousands of times through the nested sampling algorithm to thoroughly explore the physical parameter space, we necessarily make many analytical approximations to the chemistry which we outline below.

Our chemical model traces the time evolution of the chemical abundances relative to hydrogen nuclei of the 98 species including polycyclic aromatic hydrocarbons (PAHs) given in Table 3.3 at a single point representative of the whole cloud. The gas begins in primarily atomic form, with hydrogen fully molecular, carbon fully ionized and a small abundance

Table 3.3. List of chemical species included in the chemical network used by our model. Species of the form GX represent the molecule X frozen onto a dust grain mantle.

Chemical Species
H, H^+, H^-, H_2, H_2^+, H_3^+
He, He^+, Na, Na^+, e^-, PAH, PAH^+, PAH^-
C, C^+, C^-, CH, CH^+, CH_2, CH_2^+
CH_3, CH_3^+, CH_4, CH_4^+, CH_5^+
O, O^+, OH, OH^+, H_2O, H_2O^+, H_3O^+
O_2, O_2^+, CO, CO^+, HCO, HCO^+
H_2CO, H_2CO^+, CO_2, CO_2^+, HCO_2^+
N, N^+, NH, NH^+, NH_2, NH_2^+, NH_3, NH_3^+, NH_4^+
N_2, N_2^+, N_2H^+, CN, CN^+
HCN, HCN^+, HNC, NO, NO^+
HNO, HNO^+, $HCNH^+$, H_2NC^+, $HNCO^+$, H_2NO^+
S, S^+, HS, HS^+, H_2S, H_2S^+, H_3S^+
CS, CS^+, HCS, HCS^+, H_2CS^+, C_2S, C_2S^+, HC_2S^+
GCH_4, GO, GOH, GH_2O, GO_2, GCO, GH_2CO
GCO_2, GNH_3, GN_2, GHCN, GHNC
GHNO, GH_2S, GHCS, GC_2S, GNA

Table 3.4. Total elemental abundances, X, and initial fractional abundances, Y_0, of gas–phase species relative to the total number of hydrogen nucleons (H_{Tot}). Elemental abundances are taken from Hollenbach et al. (2012) except for nitrogen which was absent from their model and was taken from Godard et al. (2014). All elements start the simulation in purely atomic form except for hydrogen which is all molecular, carbon which is all ionized and a small initial abundance of CO as specified below.

Parameter	Definition	Value
X(He)	He/H_{Tot}	0.1
X(C)	C/H_{Tot}	1.4×10^{-4}
X(O)	O/H_{Tot}	3.2×10^{-4}
X(N)	N/H_{Tot}	7.9×10^{-5}
X(S)	S/H_{Tot}	2.8×10^{-5}
X(Mg)	Mg/H_{Tot}	1.1×10^{-6}
X(PAHs)	PAHs/H_{Tot}	2.0×10^{-7}
$Y_0(H_2)$	H_2/H_{Tot}	0.5
$Y_0(C^+)$	C^+/H_{Tot}	1.4×10^{-4}
Y_0(CO)	CO/H_{Tot}	1.0×10^{-6}

of carbon monoxide as outlined in Table 3.4. The chemistry consists of a network of 1348 chemical reactions with rate equations taken from the UMIST RATE12 database (McElroy et al. 2013) and solved in a time–dependent manner using the LSODE ordinary differential equation software package (Hindmarsh & Petzold 1995) for the two distinct phases. In the initial quiescent phase, we integrate the system for 1 Gyr with $u_{\rm D}$ set to zero, representing the quiescent bulk of the gas being allowed to reach chemical equilibrium in the absence of turbulent dissipation or shocks. The chemical abundances from the end of the quiescent phase are then fed as the input values for the active phase where $u_{\rm D}$ is set to its input value and the chemistry is integrated for a further 1 kyr; a typical lifetime for a turbulently dissipating vortex (Godard et al. 2014) to which we find the results are not particularly sensitive. The model output column density $\mathcal{N}_{\rm i}$ for species i is then calculated similarly to Godard et al. (2009) and Godard et al. (2014) as a weighted sum over the two phases:

$$\mathcal{N}_{\rm i} = 1.6 \times 10^{21}\, A_{\rm V}^{\rm Tot}\, [{\rm Y}_{\rm i}^{\rm Q}(1 - f_{\rm A}) + {\rm Y}_{\rm i}^{\rm A} f_{\rm A}]\, {\rm cm}^{-2}, \tag{3.7}$$

where $\mathrm{Y}_{\rm i}^{\rm Q}$ and $\mathrm{Y}_{\rm i}^{\rm A}$ are the abundances relative to hydrogen of the species i at the end of the quiescent and active phases respectively and we are implicitly assuming a standard value for the galactic gas to dust ratio.

The temperature of the gas is calculated self–consistently with the chemistry and in a time–dependent manner from various heating and cooling rates. We take the prescription of Bakes & Tielens (1994) for the heating of PAHs by UV photons as a function of the density, gas temperature and radiation field strength and that of Tielens & Hollenbach (1985) for the heating due to the ionization of molecular hydrogen by cosmic rays as a function of the cosmic ray ionization rate and the density of H_2. As seen in Godard et al. (2009), for most cases and especially at densities of $100\,{\rm cm}^{-3}$ or greater, the main mechanism of heating in a turbulently dissipating vortex is friction due to a relative drift between the ions and neutrals (as opposed to viscous dissipation). On the assumption that the ion fluid is predominantly ionized carbon and the neutral fluid is molecular hydrogen,

we determine the heating rate due to ion–neutral friction to be approximately:

$$\Gamma_{\rm in} = m_{\rm H} X \langle \sigma v \rangle_{\rm in} n^2 {u_{\rm D}}^2, \tag{3.8}$$

where $m_{\rm H}$ is the mass of atomic hydrogen, X is the ionization fraction, $\langle \sigma v \rangle_{\rm in} = 2.2 \times 10^{-9}\,{\rm cm^3\,s^{-1}}$ is the rate coefficient for momentum transfer (Flower & Pineau des Forêts 1995), n is the gas density and $u_{\rm D}$ is the ion–neutral drift velocity. In diffuse gas it has been noted that three of the main coolants are the [CII] and [OI] fine structure lines and H_2 rovibrational lines (Richings et al. 2014*a*). [SiII] and [FeII] lines have also been identified as contributing in some regimes but we choose to neglect silicon and iron from our model for the sake of chemical simplicity. Emission and collision coefficients for [CII] excited by H_2, H and e^- and for [OI] by H_2, H, H^+ and e^- are taken from Glover & Jappsen (2007) to calculate their level occupations and cooling functions using the escape probability (Sobolev 1960) formalism assuming a homogeneous slab of size $A_{\rm V}^{\rm eff}$:

$$\beta_{\rm ij} = \frac{1 - exp(-3\tau_{\rm ij})}{3\tau_{\rm ij}} \tag{3.9}$$

$$\tau_{\rm ij} = 1.6 \times 10^{21} n A_{\rm V}^{\rm eff} \sqrt{\frac{\pi m_{\rm H}}{8 k_{\rm B} T}} \frac{A_{\rm ij} \lambda_{\rm ij}^3}{8\pi} \left[\frac{n_{\rm j} g_{\rm i}}{n_{\rm i} g_{\rm j}} - 1 \right], \tag{3.10}$$

where $\beta_{\rm ij}$ is the escape probability, $k_{\rm B}$ is Boltzman's constant, T is the gas temperature, $A_{\rm ij}$ and $\lambda_{\rm ij}$ are the Einstein coefficient and wavelength of the emission process respectively and $n_{\rm i}$ and $g_{\rm i}$ represent the density and degeneracy of the level i. Similarly, we approximate the rotational cooling function for H_2, Λ_{H_2}, as a function of the low density and local thermodynamic equilibrium (LTE) limits given by Hollenbach & McKee (1979) and Galli & Palla (1998):

$$\Lambda_{H_2} = \frac{\Lambda_{H_2}({\rm LTE})}{1 + [\Lambda_{H_2}({\rm LTE}) / \Lambda_{H_2}(n \to 0)]}. \tag{3.11}$$

We calculate the LTE limit, $\Lambda_{H_2}(\mathrm{LTE})$, from the parametrized cooling function due to collisions with H, H_2, He, H^+ and e^- given by Glover & Abel (2008). The low–density limit, $\Lambda_{H_2}(n \to 0)$, is calculated from the $J = 0 - 5, v = 0$ rovibrational Einstein coefficients of the ground electronic state tabulated by Wolniewicz et al. (1998). In both cases a standard OPR of 3 for molecular hydrogen is assumed and any effects of deuteration are neglected.

Contributions to the photoionization and photodissociation of chemical species from both the UV radiation field incident on the cloud (described above) and the secondary UV photons originating from the ionization of hydrogen by cosmic rays are included. We use updated direct photorates and cosmic–ray–induced photoionization and photodissociation efficiencies from Rollins & Rawlings (2012) and where neither they nor the RATE12 database give a value for the cosmic–ray–induced photoefficiency of a reaction, we adopt a representative average value of 200. In order to avoid computationally expensive radiative transfer calculations for key photoreactions, we make use of the analytic fits to the multiplicative shielding functions shown in the PDR code benchmark study of Röllig et al. (2007) for gas of density $10^3\,\mathrm{cm}^{-3}$ and a radiation field of 10 Draine as shown in their Figure 8. For the self–shielding function of the direct photodissociation of H_2 as a function of visual extinction we use a two component power law analytic fit to their shielding function:

$$\Theta(A_V^{\mathrm{eff}}) = \begin{cases} 5.0 \times 10^{-6}\, {A_V^{\mathrm{eff}}}^{-1.2} & 0.01 \leq A_V^{\mathrm{eff}} \leq 0.30 \\ 4.0 \times 10^{-7}\, {A_V^{\mathrm{eff}}}^{-3.3} & 0.30 < A_V^{\mathrm{eff}} \leq 1.00 \end{cases} \tag{3.12}$$

Since our choice of prior function in Table 3.2 places a lower limit on the local effective extinction of 0.01 we do not consider a fit to the function below this value. Similarly, the multiplicative shielding functions for the direct photoionization of C^+ and the direct photodissociation of CO are both approximated by the analytic function given by Rawlings, Keto and Caselli (submitted):

$$\Theta(A_V^{\mathrm{eff}}) = \frac{1.4}{1 + \exp(4.9(A_V^{\mathrm{eff}} - 0.18))}. \tag{3.13}$$

While these shielding functions will likely vary for gas of lower density and weaker incident radiation fields, we expect this effect to be relatively small based on the variations seen at higher densities and radiation fields in Röllig et al. (2007). Röllig et al. (2007) do

however use a fixed line–width of $1\,\mathrm{km\,s^{-1}}$ which will clearly be a poor approximation in the second, active phase of our model where ion–neutral drift velocities are allowed to reach up to $5\,\mathrm{km\,s^{-1}}$.

We use the prescriptions of Rawlings et al. (1992), Roberts et al. (2007) and Rawlings, Keto and Caselli (submitted) for treating the gas–grain freeze–out and non–thermal desorption processes respectively, assuming a dust surface area per hydrogen nucleon of $8.0 \times 10^{-21}\,\mathrm{cm^2}$, surface density of binding sites of $10^{15}\,\mathrm{cm^{-2}}$, average grain radius of $0.0083\,\mu\mathrm{m}$ and grain albedo of 0.5. All species that freeze out are assumed to react with hydrogen on the ice mantle and fully hydrogenate on time–scales faster than the desorption time–scale. The exceptions are O, O^+ and OH^+ where it is assumed that 10% of the frozen oxygen reacts with CO in the dust mantle to form CO_2 ice while the remaining 90% forms H_2O ice. The desorption processes modelled are H_2 formation-induced, cosmic ray heating, photodesorption and cosmic–ray–induced photodesorption. Although the ices in our model are typically in the mono–layer regime, we follow Rawlings, Keto and Caselli (submitted) and allow for the possibility of the capping of water ice by CO, CO_2 and CH_4 ices resulting in a decrease of the H_2O photodesorption rate if a multi–layer ice forms. At the start of the active phase, all ice chemistry is assumed to cease and all species frozen onto dust grains are instantly desorbed due to the elevated temperatures from the ion–neutral friction.

Ion–neutral reactions in the gas phase are assumed to operate at a raised effective temperature due to the relative drift between the ionized and neutral fluids (Federman et al. 1996):

$$T_{\mathrm{eff}} = T + 40\,\mathrm{K}\left(\frac{m_{\mathrm{i}}\,m_{\mathrm{n}}}{m_{\mathrm{i}} + m_{\mathrm{n}}}\right)\left(\frac{u_{\mathrm{D}}}{1\,\mathrm{km\,s^{-1}}}\right)^2, \tag{3.14}$$

where m_{i} and m_{n} are the molecular masses (relative to atomic hydrogen) of the ion and neutral reactants respectively and we have assumed that the ion and neutral fluids are at the same temperature T and that MHD waves in the gas can be characterized by a single RMS amplitude u_{D} throughout the whole cloud. To model PAH chemistry we adopt the treatment of Hollenbach et al. (2012), considering a single–sized population of PAHs with 100 carbon atoms each and including photoreactions, charge transfer with H^+ and C^+ and recombination reactions in the chemical network.

Table 3.5. Statistical evidence, $\mathcal{Z}$, physical input parameters and a range of derived model parameters for the model $\mathcal{N}$ of diffuse gas towards G10.6–0.4 between $15\,\mathrm{km\,s^{-1}}$ and $46\,\mathrm{km\,s^{-1}}$. The values quoted for the physical and derived parameters are the posterior distribution expectation values with their errors, calculated using Equation 3.5. The derived parameters are the quiescent gas temperature T_Q, the active phase gas temperature T_A, the local ion–neutral heating rate in the active phase Γ_in, the global ion–neutral heating rate $\bar{\epsilon} = f_\mathrm{A}\Gamma_\mathrm{in}$ and the molecular hydrogen fraction f_{H_2} and ionization fraction X in the quiescent equilibrium.

Parameter	Expectation Value
$\ln(\mathcal{Z})$	-109.2 ± 0.7
$n\,/\,\mathrm{cm^{-3}}$	300 ± 100
$A_\mathrm{V}^\mathrm{eff}$	0.0107 ± 0.0005
G_0	0.7 ± 0.3
$\zeta\,/\,\zeta_0$	30 ± 10
$u_\mathrm{D}\,/\,\mathrm{km\,s^{-1}}$	2.2 ± 0.3
f_A	0.08 ± 0.01
$A_\mathrm{V}^\mathrm{Tot}$	29.2 ± 0.5
$T_\mathrm{Q}\,/\,\mathrm{K}$	49 ± 1
$T_\mathrm{A}\,/\,\mathrm{K}$	2100 ± 100
$\Gamma_\mathrm{in}\,/\,\mathrm{erg\,cm^3\,s^{-1}}$	$(3.0 \pm 1.0) \times 10^{-21}$
$\bar{\epsilon}\,/\,\mathrm{erg\,cm^3\,s^{-1}}$	$(2.0 \pm 1.0) \times 10^{-22}$
f_{H_2}	0.033 ± 0.003
X	$(2.0 \pm 0.2) \times 10^{-4}$

3.3 Results

The main results of our modelling for the diffuse gas towards G10.6–0.4 are presented in Tables 3.5 and 3.6. A key result is the value of the evidence statistic $\mathcal{Z}$ measuring the probability of our model for the dataset given in Table 3.5, along with its uncertainty due to the Poisson variability in the weighting of each nested sampling step (Skilling 2004). While larger evidence values are better, they can never be considered in isolation and rather must be compared with the evidence values for other models of the same data against an heuristic scale such as given by Jeffreys (1961) or Kass & Raftery (1995). It is noted that the maximum likelihood (i.e. frequentist) solution of our model has a value for the chi–squared statistic of $\chi^2 = 85$. While this may seem high, in the Bayesian statistical framework it is not significant on its own since it contains no information of the prior probability for the parameters and again it is stressed that when comparing two models it is the evidence statistic that should be considered.

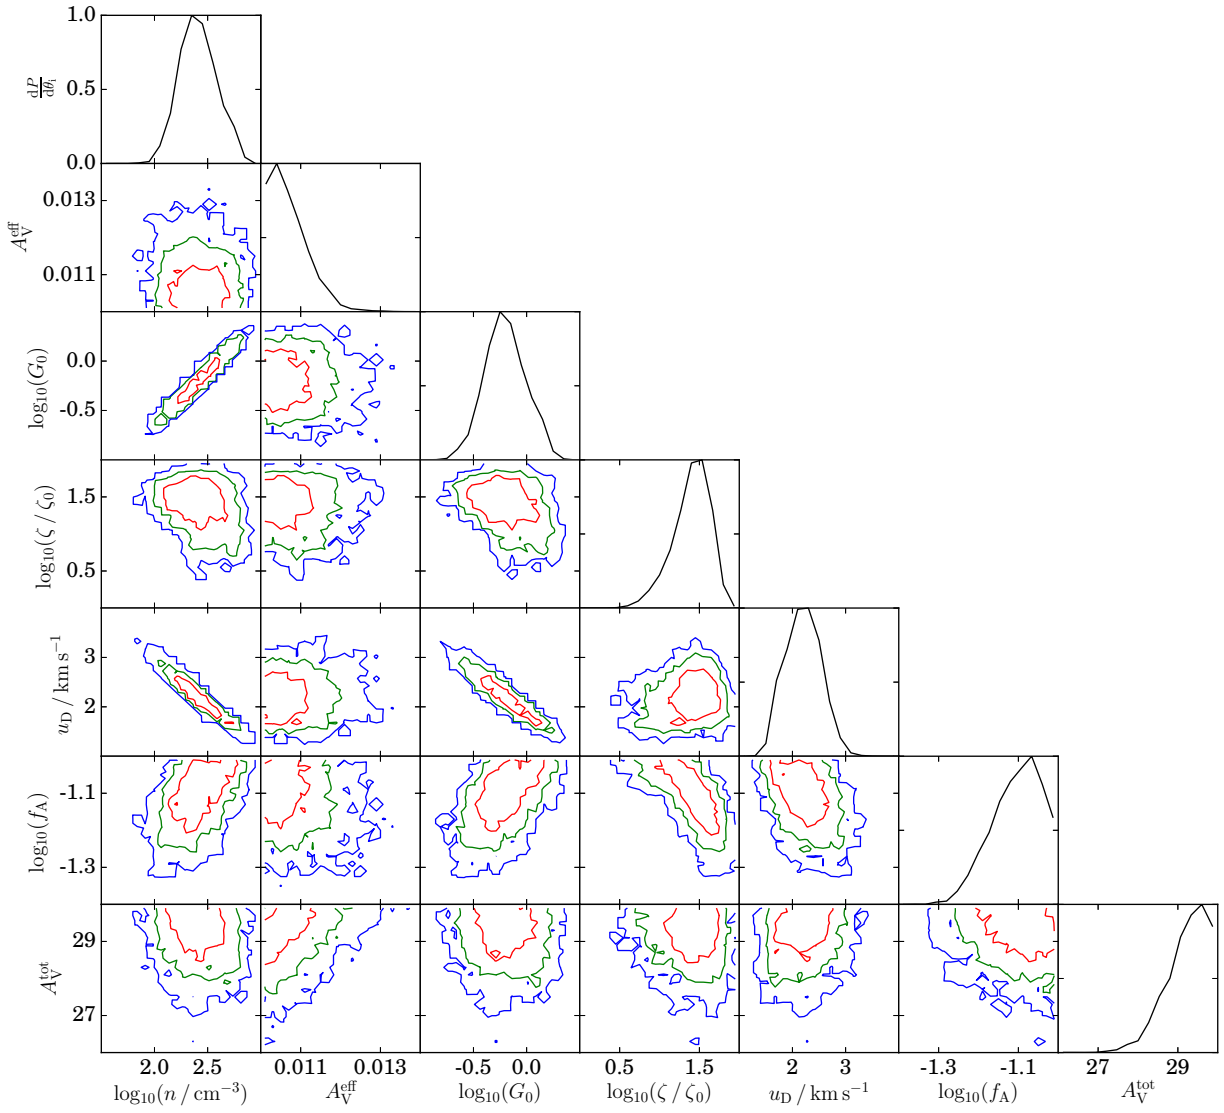

Figure 3.1. Major diagonal: marginal posterior distributions $P(\boldsymbol{\theta}_\mathrm{i}|\mathcal{D},\mathcal{N})$ for the physical parameters $\boldsymbol{\theta}_\mathrm{i}$ of the model $\mathcal{N}$, normalized to a maximum value of one. Off–diagonal: marginal joint–posterior distributions $P(\boldsymbol{\theta}_\mathrm{i},\boldsymbol{\theta}_\mathrm{j}|\mathcal{D},\mathcal{N})$ over all pairs of physical parameters i and j for the same model. Red, green and blue contours represent the one, two and three sigma confidence intervals respectively. Although the ranges plotted do not cover the entire prior space, the parameters are all well constrained to the unimodal solution shown.

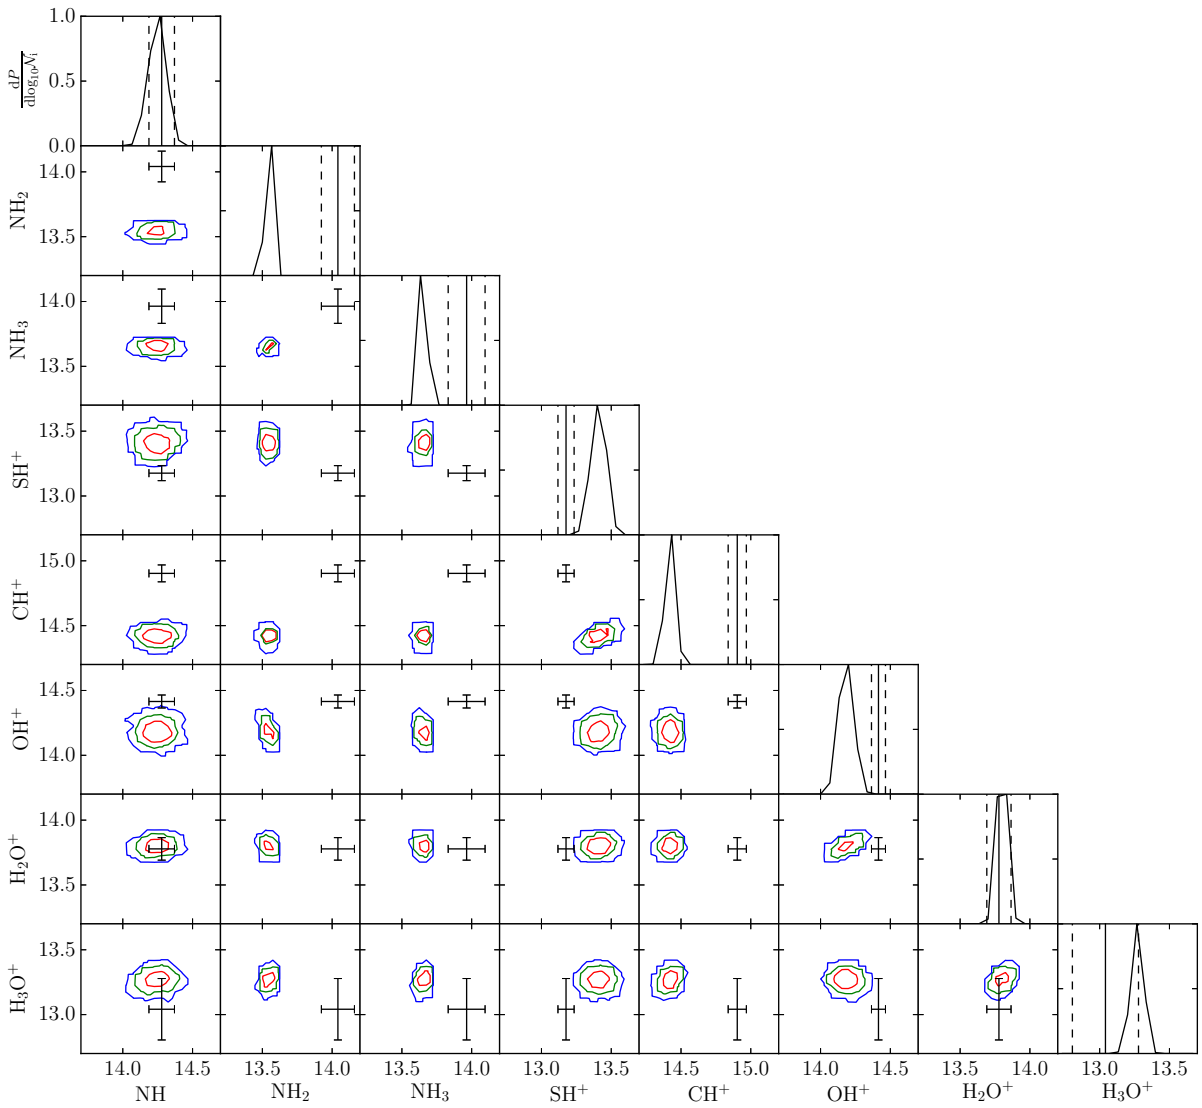

Figure 3.2. Major diagonal: marginal posterior distribution $P(\mathcal{N}_\mathrm{i}|\mathcal{D},\mathcal{N})$ for the output column densities $\mathcal{N}_\mathrm{i}$ of the model $\mathcal{N}$, normalized to a maximum value of one. Vertical lines represent the observed column densities $\mathcal{D}$ and dashed lines represent their one sigma observational errors. Off–diagonal: marginal joint–posterior distributions $P(\mathcal{N}_\mathrm{i},\mathcal{N}_\mathrm{j}|\mathcal{D},\mathcal{N})$ for the output column densities of all pairs of chemical species i and j from the same model. Red, green and blue contours represent the one, two and three sigma confidence intervals respectively. Black crosses represent the observed column densities with one–sigma errors. All distributions are plotted as a function of the logarithms $\log_{10}(\mathcal{N}_\mathrm{i}\,/\,\mathrm{cm}^{-2})$ or $\log_{10}(\mathcal{D}_\mathrm{i}\,/\,\mathrm{cm}^{-2})$ of the modelled or observed column densities of species i respectively. Note that the scales for the column densities on each set of axes are restricted to one order of magnitude in each case to show the detail.

Table 3.6. Observed and modelled column densities for the dataset $\mathcal{D}$ and model $\mathcal{N}$ of diffuse gas towards G10.6–0.4 between $15\,\mathrm{km\,s^{-1}}$ and $46\,\mathrm{km\,s^{-1}}$. The values quoted for the model column densities are the posterior distribution expectation values with their errors, calculated using Equation 3.5. Also given are the fractions of the column densities of individual species found in the active phase, f_{AS}.

Species	$\mathcal{D}\,/\,\mathrm{cm^{-3}}$	$\mathcal{N}\,/\,\mathrm{cm^{-2}}$	f_{AS}
NH	$(1.9 \pm 0.4) \times 10^{14}$	$(1.8 \pm 0.2) \times 10^{14}$	$(48 \pm 7)\,\%$
NH_2	$(1.1 \pm 0.3) \times 10^{14}$	$(3.5 \pm 0.1) \times 10^{13}$	$(0.6 \pm 0.2)\,\%$
NH_3	$(9.2 \pm 2.8) \times 10^{13}$	$(4.5 \pm 0.2) \times 10^{13}$	$(0.002 \pm 0.005)\,\%$
SH^+	$(1.5 \pm 0.2) \times 10^{13}$	$(2.6 \pm 0.3) \times 10^{13}$	$(99.994 \pm 0.001)\,\%$
CH^+	$(8.0 \pm 1.2) \times 10^{14}$	$(2.7 \pm 0.2) \times 10^{14}$	$(99.57 \pm 0.04)\,\%$
OH^+	$(2.6 \pm 0.3) \times 10^{14}$	$(1.5 \pm 0.2) \times 10^{14}$	$(97 \pm 1)\,\%$
H_2O^+	$(6.0 \pm 1.2) \times 10^{13}$	$(6.3 \pm 0.5) \times 10^{13}$	$(98.8 \pm 0.4)\,\%$
H_3O^+	$(1.1 \pm 0.6) \times 10^{13}$	$(1.9 \pm 0.2) \times 10^{13}$	$(99.8 \pm 0.07)\,\%$
NH^+	$< 1.6 \times 10^{12}$	$(1.2 \pm 0.5) \times 10^{12}$	$(27 \pm 9)\,\%$

The expectation values for the input parameters $\boldsymbol{\theta}$ are also given in Table 3.5 while their marginalised posterior and joint–posterior distributions are shown in Figure 3.1. We see immediately that the nested sampling algorithm has successfully constrained a number of the parameters including the gas density, incident radiation field, cosmic ray ionization rate and ion–neutral drift velocity to reasonable values for a diffuse gas cloud giving a unimodal posterior distribution. Importantly, the gas temperatures are consistent with the assumption of complete thermal desorption of ammonia between the quiescent and active phases (Evans et al. submitted) while the high density validates the approximation of ion–neutral friction dominating the gas heating over viscous dissipation (Godard et al. 2009). The effective local extinction and the total extinction of gas however converged to the very edge of the prior domain. It is unclear if the same is true of the gas fraction in the active phase f_{A} or if it has in fact converged to a value that is simply close to the edge of the prior. There are two possible explanations for this, either that the model provides an incomplete or inaccurate description of the diffuse gas towards G10.6–0.4 and as such inferences of these parameters are unreliable, or the previous works on which prior function was based were themselves incomplete or inaccurate. It is not possible to say definitively at this stage which of those cases is true without further observational or modelling efforts. Table 3.5 also gives the expectation values for a number of derived parameters from the model including gas temperatures, ion–neutral heating rates and molecular hydrogen and ionization fractions which we discuss in more detail in Section 3.4.

The joint–posterior distribution plots in Figure 3.1 also make clear a number of intuitive

correlations between the parameters. For example, there is a definite positive correlation between the gas density and the radiation field strength. We believe this is because as ammonia and the other nitrogen hydrides are photoionized at higher radiation fields, a higher density promotes two–body reactions with hydrogen to form the ammonium ion NH_4^+ followed by dissociative recombination back to ammonia, regulating the ammonia abundance. Similarly, there is an apparent negative correlation between the ion–neutral drift velocity and the radiation field since with higher radiation fields there would be more C^+ and S^+ ions in the gas to form CH^+ and SH^+ respectively and as such a higher ion–neutral drift velocity would be less important to raise the (effective) temperature above the energy barriers for forming those molecules. Some other correlations are less intuitive, such as the negative correlation between the cosmic ray ionization rate and the gas fraction in the active phase. We discuss these correlation further in Section 3.4, but note that such degeneracies could possibly be broken in future studies with extra observations leading to tighter constraints on the input parameters.

In Table 3.6 and Figure 3.2 we give the expectation values and marginalized and joint posterior distributions for the column densities of the eight species modelled. All species are fit to within half an order of magnitude of the observations which is reasonable given the significant uncertainties on the chemical rate equations that were not considered in the statistical analysis. The spread in the distributions of output column densities is also small, especially compared with the uncertainties on the input parameters, suggesting that either the chemistry is relatively insensitive to the physical state of the gas or that the nested sampling algorithm creates such an aggressive fit to the observations that any models with slightly different outputs from those of the posterior–peak solution are statistically much less favourable. There are also no correlations between species, even those whose chemistries are intrinsically linked such as the three nitrogen hydrides, which indicates that each species is individually constraining the input parameters and that the addition of further observed species to the likelihood calculation would lead to further constraints.

3.4 Discussion

Our study is able to tell us a lot about the physical state and chemical processes in the diffuse gas towards G10.6–0.4. For example, the posterior distribution on the radiation field is consistent with the intensity of the Draine interstellar radiation field (Draine 1978).

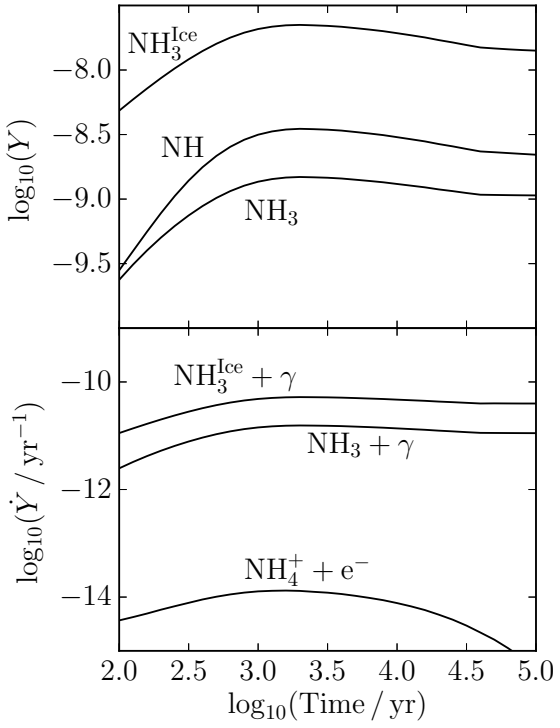

Figure 3.3. Top: Abundances of key nitrogen bearing species as a function of chemical integration time during the initial quiescent phase. NH_3^{Ice} formation is clearly seen to precede the formation of ammonia and NH and is between two and eight orders of magnitude more abundant than NH_4^+ (off the bottom of the plot) suggesting a gas–grain chemistry origin for NH_3. Bottom: Rates for the two fastest formation reactions of ammonia as a function of chemical time in the initial quiescent phase. The relative rates for the non–thermal desorption of ammonia ice from grains ($NH_3^{Ice} + \gamma$) and dissociative recombination of ammonium $NH_4^+ + e^-$ suggests a grain–chemistry origin for ammonia. The rate for the photodissociation to NH ($NH_3 + \gamma$) is also shown suggesting a top–down formation for the simpler nitrogen hydrides.

It is clearly important that the radiation field is not so high as to destroy the observed nitrogen hydride species by photodestruction and photoionization reactions. However, we also see that the constrained radiation field is not sufficient to synthesise those species by the photoionization of NH followed by successive gas–phase hydrogenation and charge–transfer or recombination reactions as outlined by Sternberg & Dalgarno (1995) for dense gas and modelled by Persson et al. (2010) for diffuse gas. Instead, we see clear evidence of those species being produced by the hydrogenation of atomic nitrogen to ammonia on grain surfaces followed by non–thermal desorption and photodissociation to the smaller species NH and NH_2. In Figure 3.3 we show that the formation of ammonia ice precedes that of gas–phase ammonia and NH while the NH_4^+ ion abundance is between two and

eight orders of magnitude less than that of ammonia. Similarly, we see that the rate for ammonia formation via non–thermal desorption from ammonia ice is also typically three to five orders of magnitude faster than any reaction from the ammonium ion (with dissociative recombination being the fastest channel). Furthermore, the rate for the photodissociation of ammonia to NH demonstrates the top–down formation of nitrogen hydrides and explains the tight coupling between the ammonia and NH fractional abundances. We finally note that almost half of the NH column density is in the active phase while nearly all of the NH_2 and ammonia resides in the quiescent phase which could explain why some models have been unsuccessful when using gas–phase or grain–surface chemistry only to model the nitrogen hydride abundances. We also note that the inferred column density of NH^+ is close to the current observational limit making a case for further, more sensitive observations to help constrain the model.

Our model also constrains the velocity of the ion–neutral drift to $2.2 \pm 0.3\,\mathrm{km\,s^{-1}}$. This velocity agrees qualitatively with the line widths in the observed spectra (Persson et al. 2012, Godard et al. 2012, Indriolo et al. 2015) although it is somewhat lower than the HI velocity dispersion seen in the cold neutral medium (Haud & Kalberla 2007). The value is also higher than the $1\,\mathrm{km\,s^{-1}}$ value for many weak C–shocks throughout the gas as modelled by Gredel et al. (2002) ruling out such a scenario. More significantly, the value is in agreement with both the upper limit of $3.5\,\mathrm{km\,s^{-1}}$ for the ion–neutral drift velocity in the turbulent dissipation region model for diffuse gas of Godard et al. (2014) and the value for a single strong C–shock with drift velocity of $2\,\mathrm{km\,s^{-1}}$ as modelled by Flower & Pineau des Forêts (1998). Our constraint on the fraction of gas in the active phase $f_\mathrm{A} = 0.08$ seems more likely for the case of turbulent vortices in the gas (Godard et al. 2009) along with the fact that the species in the shocked gas are seen across the entire velocity range. However, the gas temperature in the active phase was constrained to 2100 K which would seem more consistent with models of shocked gas (Chieze et al. 1998, Flower & Pineau des Forêts 1998) and the inferred global ion–neutral heating rate $\bar{\epsilon}$ is significantly higher than the observed global turbulent heating rate of Hily-Blant et al. (2008). One significant shortcoming of our model is that it considers only a single density for both the quiescent and active phases while a shock model would be expected to have a significant density contrast between the two phases. On balance, we cannot currently distinguish between the two scenarios. Discriminating between turbulent vortices and shocks as mechanisms of chemical enrichment in diffuse gas represents an interesting future work where the evidence

statistic could be used to discriminate between models for the two cases with more accurate geometrical and physical considerations than in this work.

The cosmic ray ionization rate is constrained to $\zeta = (5 \pm 2) \times 10^{-16}\,\mathrm{s}^{-1}$ which is in agreement with the inferred value found by Indriolo et al. (2015) for the same gas. We expected to be able to constrain ζ due to the importance of cosmic rays in initiating ion–neutral oxygen chemistry and hence explaining the observed columns of OH^+, H_2O^+ and H_3O^+. However, Table 3.6 shows that for all three of those species more than 95% of their total column density resides in the active phase of the gas in spite of the relatively short chemical time–scale of 1 kyr in that phase. The reaction $O^+ + H_2 \rightarrow OH^+ + H$ is still orders of magnitude faster than any other formation pathway initiating oxygen chemistry, even though gas heating in the active phase could have promoted other reactions with endothermic energy barriers such as $O + H_2 \rightarrow OH + H$. Therefore, while traditionally used as tracers of cosmic rays, we suggest that these species are equally important tracers of turbulent or shocked diffuse gas. This was already apparent in the model of Godard et al. (2009) where up to 89% of the species H_3O^+ was seen to be in either the active or relaxation phase of their turbulent dissipation model for diffuse gas. In light of this we can also explain the negative correlation between ζ and the fraction of gas f_A in the active phase seen in Figure 3.1, since if f_A is allowed to fall an increased cosmic ray ionization rate would be needed to initiate a more rapid gas–phase oxygen chemistry and increase the relative abundances of OH^+, H_2O^+ and H_3O^+ in the active phase.

One interesting feature of the posterior distribution shown in Figures 3.1 and 3.2 is that it is unimodal. While perhaps unremarkable on its own, there has been active discussion on the phenomenon of bistability of networks of astrochemical rate equations (Le Bourlot et al. 1993, Pineau Des Forêts & Roueff 2000) where there is more than one steady–state solution between which the chemical abundances can vary by orders of magnitude. While it is recognised as an intrinsic mathematical property of the rate equations, it has been demonstrated that certain gas–grain chemical processes such as depletion (Lee et al. 1998) and surface hydrogenation (Le Bourlot et al. 1995) could significantly change or reduce the regions of the parameter space where the phenomena occurs. In the model of Le Bourlot et al. (1995), the two solutions are seen to correspond with low– and high–ionization phases of the gas. Our inferred value in Table 3.5 for the ionization fraction $X = 2 \times 10^{-4}$ places the gas towards G10.6–0.4 in the high–ionization phase. We see no evidence in our model for bistable solutions; either there is no bistability, the Bayesian fit to the observational data

only admits one of the two solutions or the second solution is suppressed by hydrogenation reactions on dust grain surfaces.

The fact that our model has been able to fit the observed column densities and provide statistical constraints on some of the physical parameters in the diffuse gas towards G10.6–0.4 as well as it has suggests that it captures many of the essential physical and chemical processes in spite of the significant approximations. However it is clearly not complete. Although the modelled column densities are not perfect fits to the observations, all are accurate to within half an order of magnitude or better. These discrepancies can be largely accounted for by the uncertainties in the chemical rate coefficients. We also note that when repeating the model and fitting to only a single species at a time, we were able to get perfect fits between the posterior distributions and the observations for each species individually with reasonable (although varying) physical parameters for diffuse gas. Clearly the model contains sufficient chemical pathways to form each of the species in isolation and only "'fails" when trying to fit multiple species at once. One aspect of the modelling that was not successful was constraining the extinction properties of the gas, both the total and local effective extinctions. In both cases they did not agree with observations (Neufeld et al. 2010*b*) or models (Godard et al. 2014) and the fits were to values at the edge of the domain of the prior function. The success in modelling individual species yet failure to explain the extinction properties suggests that our initial assumption of a single point geometry due to the presence of all eight species across the whole velocity range was incorrect. Further chemical stratification such as non–homogeneous clouds or variations between clouds appears likely and would demand a more complex and detailed geometry for the model.

3.5 Conclusions

We have developed a fast and approximate chemical model for the diffuse gas clouds along the line of sight towards G10.6–0.4. By using the nested sampling algorithm to compare its outputs with observed molecular column densities we were able to statistically constrain many of the physical properties of the gas. The gas is relatively dense, cool and features significant velocity decoupling between the ionic and neutral species but is not chemically dominated by the presence of UV photons or many weak C–shocks. The column densities of NH, NH_2, NH_3, NH^+, SH^+, CH^+, OH^+, H_2O^+ and H_3O^+ were all modelled to within

half an order of magnitude of their observed values through a combination of grain surface chemistry, ion–neutral reactions and cosmic ray ionization processes.

This paper presents a novel method for astrochemical modelling that overcomes the somewhat arbitrary nature of defining the constraints in systems containing many free parameters and demonstrates how the field can transition to higher precision science to meet the demands of future observational programs. The statistical technique that we applied is completely general and as such can be used to analyse refinements to our model or completely new models based on different physics and/or chemistry since evidence values are directly comparable when the same datasets are used. This opens up the possibility of many improvements to the current work. As identified in Section 3.4, a key improvement in helping to distinguish between shocked and turbulently dissipating gas would be to move towards a fully three dimensional model with a realistic geometry and models for the dynamics of those processes. For example, we could model individual velocity components rather than the total velocity integrated column densities, accounting for spatial variations by applying realistic density profiles and assigning each gas component its own independent physical properties including a systemic velocity. Such an approach would also allow us to carry out radiative transfer modelling of the system and match the output directly to the observed absorption spectrum, although this would add significantly to the computational complexity. However, it would also lead to increased dimensionality of the parameter space along with possible degeneracy and multimodality that would significantly reduce the efficiency of the nested sampling algorithm.

Nested sampling is also able to address many other interesting issues. We could, for example, compare different databases of chemical reaction data to investigate systematic differences, or use more complex models for chemical processes. Of particular interest for the specific astrochemical problem that we have addressed in this paper, the addition of spin–dependent chemistry for the nitrogen hydrides (e.g. Faure et al. 2013, Le Gal et al. 2014) would be advantageous. This could be used to see if the "anomalous" OPR of ammonia and the predicted OPR of NH_2 could be reproduced via grain surface chemistry and what constraints that might place on, for example, the OPR of molecular hydrogen in diffuse gas. Alternatively, parameters that we have assumed to be constant, such as the durations of the quiescent and active phases, could be allowed to vary and ultimately constrained by the nested sampling algorithm. As discussed, the parametrized rates for each reaction in our chemical network are assumed constant yet many are likely to have

large errors associated with them and allowing the coefficients for a small number of key reactions to vary may shed light on uncertainties in their adopted values (e.g. Woods et al. 2012).

Additional data might also help to better constrain the physical conditions of the diffuse gas. For example, with appropriate additions to the network of chemical reactions, *Herschel* observations of HF (Neufeld et al. 2010*b*), HCl (Monje et al. 2013), H_2Cl^+ (Neufeld et al. 2012) and ArH^+ (Schilke et al. 2014) towards G10.6–0.4 could be analysed as part of our dataset. Similarly, new data in the era of the *Atacama Large Millimeter Array* is likely to lead to more accurate measurements of existing absorption spectra, the discovery of new species and resolved morphology which again may help to discriminate between the cases of shocked and turbulently dissipating gas. Furthermore, better constraints on some or all of the physical parameters from other studies could also be incorporated into the nested sampling calculation as informative priors to further constrain the results of the statistical analysis.

Chapter 4

The Chemical Effects of Mutual Shielding in Photon Dominated Regions

The work presented in this chapter is based on the paper by Rollins & Rawlings (2012) in collaboration with J. M. C. Rawlings.

Photodissociation or Photon Dominated Regions (PDRs) are ubiquitous components of the interstellar medium, being the interfaces between the partially ionized and the neutral, dense, molecular components. Examples of PDRs include the boundaries of molecular clouds subjected to the interstellar radiation field, protoplanetary discs irradiated by accreting protostars and the interstellar medium of starburst galaxies illuminated by clusters of massive star formation. These interfaces are chemically rich and they often dominate the molecular emission from the host regions due to the relatively high excitation of the molecular gas. Moreover, as earlier, simplistic assumptions about the geometries of these regions give way to more realistic and complex morphologies, it would seem that a larger proportion of interstellar gas is contained within PDRs.

However, the physical and chemical structures of PDRs are complex and only partially described by existing models. The thermal balance in PDRs is determined by heating

by ambient and stellar far–ultraviolet (FUV) radiation ($6 < h\nu < 13.6$eV) and cooling by atomic and molecular line emission and continuum emission by dust (Hollenbach & Tielens 1999, Sternberg 2004). This in turn depends on the chemistry. The transition from atomic to molecular gas is defined by the radiative transfer coupled to the photochemistry and entails complex self– and mutual shielding effects.

It is usually reasonable to assume that H_2, CO and its isotopologues are the only molecular species that have large enough abundances to result in significant absorption of the incident radiation field. The primary consequence of this is the increasing self–shielding of those photoreactions as the transitions become more optically thick. In addition, where there is overlap between the absorption lines of different species, mutual shielding will occur. Also, one must consider the contributions to the FUV opacity from the carbon ionization continuum (for $\lambda < 1102\,\text{Å}$) and absorption by dust. This leads to a range of possible mutual and self–shielding effects as listed below:

1. H_2 self–shielding,
2. CO self–shielding,
3. H_2, CO and CO isotopologue mutual shielding,
4. Mutual shielding by coincident Lyman transition lines of HI and Lyman and Werner transition lines of H_2,
5. Shielding of all species by dust absorption,
6. Shielding of H_2 and CO by the carbon ionization continuum, and
7. Shielding of *other* species by all of the above processes.

Of these, the cross–section for attenuation by the dust continuum is approximately constant over the wavelength ranges of individual photoreaction channels and so is separable from the other shielding terms. In practice, the theoretical studies are somewhat limited by the paucity of laboratory and theoretical data for molecules other than H_2 and CO.

Tielens & Hollenbach (1985) developed a PDR model to study the chemistry and thermal balance of gas with density $10^3\,\text{cm}^{-3} < n < 10^6\,\text{cm}^{-3}$ illuminated by FUV fluxes $\chi = 10^3 - 10^6$ times more intense than the interstellar radiation field. They identified the presence of several physically distinct regions within PDRs, including an HI $\rightarrow$ H_2

transition zone at $A_{\rm v} < 1$ and a complex, warm $C^+ \rightarrow C \rightarrow CO$ transition zone at $A_{\rm v} \sim 3 - 5$. Their model, which included carbon continuum self–shielding, yielded CI column densities of $4.5 - 16.0 \times 10^{18} {\rm cm}^{-2}$ at which levels the continuum is very optically thick. Similarly, Rawlings (1988) developed a photochemical model of the (dust–free) PDR in the vicinity of a nova, which concentrated on the mutual shielding between H_2, CO and CI.

Subsequent studies such as the PDR benchmarking study of Röllig et al. (2007), although more sophisticated than the earlier models and employing accurate photoreaction data for CO and H_2, essentially confirmed the basic structure with the $C^+ \rightarrow C \rightarrow CO$ transition occurring at $A_{\rm v} \sim 2 - 5$ for $n \sim 10^3 {\rm cm}^{-3}$ and $\chi = 10$. These various models take careful account of the overlap between the CO and H_2 absorption lines and include the dust continuum, but do not make specific reference to shielding by the CI ionization continuum.

With the availability of reliable wavelength and oscillator strength data for the photodissociation of CO, a much more detailed and accurate treatment of the self– and mutual shielding of H_2 and CO has been possible. van Dishoeck & Black (1988) considered the photodissociation of CO and its main isotopologues together with H and H_2, including self– and mutual shielding effects. They also included shielding by dust and referred to the attenuation by the carbon continuum. As shown in that paper (e.g. Figure 4), when the CO, HI and H_2 absorption lines become very optically thick, they can effectively blanket a significant fraction of the continuum in the $912\,{\rm \AA} < \lambda < 1079\,{\rm \AA}$ wavelength range, with implications for the photodissociation rates of other species.

Visser et al. (2009) included self–shielding, mutual shielding and shielding by H and H_2 to calculate the photodissociation rates of CO and its isotopologues, including rarer variants. For these species only, the photodissociation rate coefficient can be written as:

$$k = \chi k_0 \Theta \left[N({\rm H}_2), N({\rm CO}), N({\rm H}), \phi \right] \exp(-\gamma A_{\rm v}), \tag{4.1}$$

where k_0 is the unattenuated rate, Θ is the shielding function, $\exp(-\gamma A_{\rm v})$ is the (separable) dust extinction term and ϕ represents other factors such as the Doppler widths of the absorption lines, the excitation temperatures for CO and H_2 and the ratios of the column densities of the CO isotopologues. As with previous studies the line and continuum attenuation components were treated as separable. However, as pointed out by van

Dishoeck (1988) the $A_{\rm v}$ dependence does *not* include the attenuation at $\lambda < 1102\,$Å by H_2 line/continuum and CI continuum absorption which will be particularly important for species like CN and N_2 that photodissociate at the shortest wavelengths.

Thus, the earliest studies of PDRs (e.g. Tielens & Hollenbach 1985) – although lacking the details of more recent models – attempted to deal with most of the shielding effects listed above. In subsequent studies all discussion of the role of the carbon ionization continuum is seemingly absent. This is presumably on the assumption – which may or may not be valid – that dust extinction dominates. However, the cross–section for carbon photoionization is $1.6 \times 10^{-17}{\rm cm}^2$. It is fairly reasonable to assume that column densities of atomic carbon in excess of $10^{17}{\rm cm}^{-2}$ will pertain deep within PDRs (e.g. Figure 11 of Visser et al. 2009) and so this opacity source must be considered.

In addition, the more recent PDR models have tended to concentrate on the carbon budget, while the effects of mutual shielding by H_2 and CO on *other* molecules have been largely overlooked. In this study we consider two of these effects; the mutual shielding by:

1. molecular hydrogen lines, and

2. the carbon ionization continuum.

Critically, there is extensive overlap of both of these with the photodissociation cross–sections of CN and N_2. The mutual shielding effects therefore have the potential to promote nitrogen-based chemistry at extinctions much shallower than currently seen in simulations. This paper presents shielding factors for a range of species due to these effects and investigates their influence on the chemical structure. Our calculation of the degree of mutual shielding by H_2 dissociation and carbon ionization are presented in Section 4.1. These shielding factors are then implemented in a simple PDR model, described in Section 4.2 with results presented in Section 4.3. We close with a discussion of our results and concluding remarks in Section 4.4.

4.1 Photoreaction Rates

In PDRs, the shielding of the incident radiation field by molecules is typically a combination of line and continuum processes leading to a complex radiative transfer problem. Most models include self–shielding by CO and H_2 by adding multiplicative shielding factors to their photorates as described above. We extend this procedure to investigate mutual

shielding by the H_2 line and CI ionization continuum absorption. In this section we detail our method for calculating photoreaction rates and shielding factors for these species.

4.1.1 Direct Photorates

It is usually assumed that the line and continuum contributions to the photorates are separable (van Dishoeck 1988). Unshielded rate coefficients, $k_{\rm pd}$, can be calculated using the following equations for line and continuum photoprocesses (e.g. van Dishoeck 1988):

$$k_{\rm pd}^{\rm cont} = \int_{912\,\text{Å}}^{\infty} \sigma(\lambda) I(\lambda)\, d\lambda\, s^{-1}, \tag{4.2}$$

$$k_{\rm pd}^{\rm line} = \frac{\pi e^2}{mc^2} \lambda_{\rm ul}^2 f_{\rm ul} \eta_{\rm u} x_{\rm l} I(\lambda_{\rm ul})\, s^{-1}, \tag{4.3}$$

where I is the specific photon flux of the radiation field, σ is the cross–section of the continuum transition and the line transition is at a wavelength $\lambda_{\rm ul}$ from a lower level l to an upper level u with oscillator strength $f_{\rm ul}$, dissociation efficiency $\eta_{\rm u}$ for the upper level and occupation fraction $x_{\rm l}$ for the lower level. The lower limit of the integral (912 Å) is the Lyman cutoff wavelength for the unshielded interstellar radiation field. From these expressions we have (re)calculated the unshielded interstellar photorates using cross–section and oscillator strength data primarily from van Dishoeck (1988), referring to other sources (e.g. Rawlings et al. 1993) for species and reaction channels not covered in their database. Where we have one set of cross–section data for a molecule but multiple dissociation channels are listed in the UMIST Database for Astrochemistry (UMIST06, Woodall et al. 2007), our calculated rate is divided between those channels in the same ratio as the original UMIST06 rates. We note that the high–resolution N_2 dissociation line data from Carter (1972) has been re–binned to a lower wavelength resolution in the database of van Dishoeck (1988). We therefore elect to use the mid-points of the dissociation wavelength intervals from Carter (1972) as the representative wavelengths of the line transitions. Since no cross–section data is available for HNC we assign the same values as derived for HCN. For the radiation field we adopt the standard interstellar radiation field of Draine (1978) with the extension to $\lambda > 2000$ Å given by van Dishoeck & Black (1982). The calculated photorates are

Table 4.1. Photoionization rates, carbon continuum shielding factors, cosmic–ray–induced photoreaction efficiencies and cross–section data references; van Dishoeck 1988 (`http://home.strw.leidenuniv.nl/~ewine/photo/`, vD88), Branscomb 1962 (B62), Barsuhn 1977 (B77), Black & Dalgarno 1977 (BD77) and Rawlings et al. 1993 (RDB93).

Species	Photorate / s^{-1}	Maximal Carbon Shielding Factor, S_i	CR–Induced Photoreaction Efficiency, p_M	Reference
C	3.15×10^{-10}	0.000	264.8	vD88
C^-	3.31×10^{-08}	0.994	1037.8	B62
CH	7.63×10^{-10}	0.142	547.7	vD88
CH_2	3.58×10^{-10}	0.449	300.0	B77
CH_3	4.93×10^{-10}	0.600	509.5	BD77
NH	1.50×10^{-11}	0.000	22.0	B77
NH_2	1.73×10^{-10}	0.000	149.3	RDB93
NH_3	2.82×10^{-10}	0.298	230.5	vD88
O_2	7.65×10^{-11}	0.000	57.7	vD88
OH	2.00×10^{-11}	0.000	26.6	B77
H_2O	3.44×10^{-11}	0.000	24.5	vD88
Na	1.54×10^{-11}	0.829	9.1	vD88
HCO	4.81×10^{-10}	0.590	499.2	RDB93
H_2CO	4.82×10^{-10}	0.215	379.5	vD88
NO	2.61×10^{-10}	0.328	218.9	vD88
S	6.06×10^{-10}	0.343	480.2	vD88
H_2S	7.40×10^{-10}	0.344	594.4	vD88

presented in Tables 4.1 and 4.2.

The newly calculated direct photorates are in agreement with the values in UMIST06 to within a factor of two for all but 10 of the 61 reactions for which we had data available. Four show rates differing by more than one order of magnitude; our calculations for the dissociation of H_3^+ to either H_2^+ + H or H_2 + H^+ are each larger by a factor of ~ 400, while the ionization of OH is larger by a factor of ~ 12 and the dissociation of OH^+ to O^+ + H is larger by a factor of ~ 19.

4.1.2 Shielding of N_2 and CN by H_2

The photodissociation of H_2 is effected by the absorption of Lyman and Werner band photons followed by decay into the dissociation continuum (Stecher & Williams 1967). To investigate the possible effects of mutual shielding by these lines we note that the range of wavelengths for H_2 dissociation ($844.8 - 1108.5$ Å) fully overlaps that for the dissociation lines of N_2 ($914 - 980$ Å, Carter 1972, van Dishoeck 1988) and nearly the whole

Table 4.2. Photodissociation rates, carbon continuum shielding factors, cosmic–ray–induced photoreaction efficiencies and cross–section data references; van Dishoeck 1988 (`http://home.strw.leidenuniv.nl/~ewine/photo/`, vD88), Kulander & Bottcher 1978 (KB78) and Carter 1972 (C72).

Species	Photorate / s^{-1}	Maximal Carbon Shielding Factor, S_i	CR–Induced Photoreaction Efficiency, p_M	Reference
$H_2^+ \rightarrow H^+ + H$	5.73×10^{-10}	0.778	427.7	vD88
$H_3^+ \rightarrow H_2^+ + H$	2.16×10^{-12}	0.000	0.15	KB78
$H_3^+ \rightarrow H_2 + H^+$	2.16×10^{-12}	0.000	0.15	KB78
$CO \rightarrow O + C$	2.59×10^{-10}	0.000	105.0, see text	vD88
$CO^+ \rightarrow C^+ + O$	1.01×10^{-10}	0.954	45.7	vD88
$CH \rightarrow C + H$	9.19×10^{-10}	0.975	465.4	vD88
$CH^+ \rightarrow C + H^+$	3.27×10^{-10}	0.036	254.7	vD88
$CH_2 \rightarrow CH + H$	5.81×10^{-10}	1.000	200.4	vD88
$CH_2^+ \rightarrow CH^+ + H$	4.53×10^{-11}	0.461	66.1	vD88
$CH_2^+ \rightarrow CH + H^+$	4.53×10^{-11}	0.461	66.1	vD88
$CH_2^+ \rightarrow C^+ + H_2$	4.53×10^{-11}	0.461	66.1	vD88
$CH_3 \rightarrow CH_2 + H$	1.38×10^{-10}	1.000	76.1	vD88
$CH_3 \rightarrow CH + H_2$	1.38×10^{-10}	1.000	76.1	vD88
$CH_4 \rightarrow CH_3 + H$	1.89×10^{-10}	0.703	182.8	vD88
$CH_4 \rightarrow CH_2 + H_2$	8.42×10^{-10}	0.703	182.8	vD88
$CH_4 \rightarrow CH + H_2 + H$	1.89×10^{-10}	0.703	182.8	vD88
$CH_4^+ \rightarrow CH_2^+ + H_2$	2.23×10^{-10}	0.405	139.3	vD88
$CH_4^+ \rightarrow CH_3^+ + H$	5.23×10^{-11}	0.405	139.3	vD88
$NH \rightarrow N + H$	5.02×10^{-10}	0.794	219.2	vD88
$NH_2 \rightarrow NH + H$	7.45×10^{-10}	0.783	152.8	vD88
$NH_3 \rightarrow NH + H_2$	4.83×10^{-10}	0.780	159.2	vD88
$NH_3 \rightarrow NH_2 + H$	6.82×10^{-10}	0.891	148.1	vD88
$N_2 \rightarrow N + N$	2.28×10^{-10}	0.000	184.3	vD88, C72
$O_2 \rightarrow O + O$	7.90×10^{-10}	0.890	361.3	vD88
$O_2^+ \rightarrow O^+ + O$	3.46×10^{-11}	1.000	3.0	vD88
$OH \rightarrow O + H$	3.76×10^{-10}	0.752	279.3	vD88
$OH^+ \rightarrow O^+ + H$	1.30×10^{-11}	0.008	10.7	vD88
$H_2O \rightarrow OH + H$	7.54×10^{-10}	0.746	527.0	vD88
$HCO \rightarrow CO + H$	1.11×10^{-09}	1.000	319.8	vD88
$HCO^+ \rightarrow CO^+ + H$	5.39×10^{-12}	0.000	0.0	vD88
$H_2CO \rightarrow CO + H + H$	5.79×10^{-10}	0.844	132.1	vD88
$H_2CO \rightarrow CO + H_2$	5.79×10^{-10}	0.844	132.1	vD88
$H_2CO \rightarrow HCO^+ + H + e^-$	1.16×10^{-11}	0.844	132.1	vD88
$CO_2 \rightarrow CO + O$	8.81×10^{-10}	0.329	643.8	vD88
$CN \rightarrow N + C$	2.92×10^{-10}	0.018	229.2	vD88
$HCN \rightarrow CN + H$	1.56×10^{-09}	0.677	436.9	vD88
$HNC \rightarrow CN + H$	1.56×10^{-09}	0.677	436.9	vD88
$NO \rightarrow O + N$	4.73×10^{-10}	0.737	163.5	vD88
$HS \rightarrow S + H$	9.79×10^{-10}	0.947	271.9	vD88
$HS^+ \rightarrow S^+ + H$	2.60×10^{-10}	0.343	171.9	vD88
$H_2S \rightarrow HS + H$	1.55×10^{-09}	0.845	412.6	vD88
$H_2S \rightarrow S + H_2$	1.55×10^{-09}	0.845	412.6	vD88
$CS \rightarrow S + C$	9.75×10^{-10}	0.973	1570.0	vD88

range for CN photodissociation ($912 - 1133$ Å, peaking at 940 Å). We are investigating the photochemistry well within a PDR ($A_{\rm v} > 2$) and the H_2 lines will be extremely optically thick, to the extent that (a) the overlap with individual N_2 lines will be extensive and (b) there will be significant blanketing of the Lyman/Werner continuum. To include these effects we therefore make the extreme assumption that N_2 and CN are both *fully* shielded and that the effective rates of photodissociation can be set to zero. These assumptions are clearly very simplistic and will need addressing using more sophisticated radiative transfer techniques in a future study.

4.1.3 Shielding by Atomic Carbon

To account for shielding by the CI ionization continuum we take a more detailed and accurate approach. First of all, considering the situation where the continuum is very optically thick, the photoreaction rates are recalculated using Equations 4.2 and 4.3 but using a lower limit of 1102 Å for the integral and excluding line contributions for $\lambda <$ 1102 Å. This implicitly assumes that the radiation field intensity is completely attenuated at those wavelengths shortwards of the CI ionization potential; I($\lambda <$ 1102 Å)=0. In Tables 4.1 and 4.2 we give the resulting shielding factors $S_{\rm i}$ which represent the ratios of the photoreaction rates without and with the contributions from the $912 < \lambda <$ 1102 Å wavelength range. Thus, in conditions where the carbon ionization continuum is very optically thick, the rate coefficient is just $S_{\rm i}$ times the unshielded value. These ratios are calculated for all species and reaction channels where cross–section data is known.

We note that the photoionization cross–section of atomic carbon is $\sigma_{\rm C} = 1.6 \times 10^{-17}$ cm^2 and is approximately independent of wavelength for $\lambda <$ 1102 Å (van Dishoeck 1988). CI column densities of up to the order of 10^{19} cm^{-2} (Tielens & Hollenbach 1985) can be expected in PDRs, leading to a wide range of potential optical depths. Taking the first approximation that attenuation by dust and molecules are separable and the fact that the cross–section is uniform with wavelength simplifies the radiative transfer to give shielding functions, $\Theta_{\rm i}$, for each photoreaction i (analogous to those used for CO):

$$\Theta_{\rm i}(N_{\rm C}) = S_{\rm i} + (1 - S_{\rm i}){\rm e}^{-\sigma_{\rm C} N_{\rm C}}, \tag{4.4}$$

where $S_{\rm i}$ is the relevant shielding factor for the optically thick limit, given in Tables 4.1 and

4.2, and $N_{\rm C}$ is the carbon (CI) column density. Strictly speaking, the $A_{\rm v}$-dependence (due to continuum absorption by dust; γ in equation 4.1) will also depend on the CI continuum opacity. However, this is a complicating second-order effect and we follow the practice of previous studies by decoupling the shielding by CI from that by dust absorption, as in equation 4.1.

4.1.4 Cosmic–Ray–Induced Photoreactions

At extinctions greater than a few magnitudes, or if any of the various shielding effects described above become important, cosmic–ray–induced photodissociation and photoionization processes are significant and must be included in PDR models. The radiation field is generated by secondary electrons produced in the cosmic ray ionization of H_2 (Prasad & Tarafdar 1983). The calculation of these rates is as described in Sternberg et al. (1987) and Gredel et al. (1989). On the assumption that the total absorption cross–section is dominated by dust (rather than the molecular component, which in normal circumstances is true, except for H_2 and CO) then the photorates are given by:

$$R_{\rm i} = \frac{\zeta n_{\rm i}}{(1-\omega)} p_{\rm i}\, cm^{-3}\, s^{-1}, \tag{4.5}$$

where ζ is the cosmic ray ionization rate, $n_{\rm i}$ is the density of species i, and ω is the grain albedo. The cosmic–ray–induced photoreaction efficiency $p_{\rm i}$ is given by:

$$p_{\rm i} = \int \frac{\sigma_{\rm i}(\nu) P(\nu)}{2\sigma_{\rm g}} d\nu, \tag{4.6}$$

where $\sigma_{\rm i}$ is the photoreaction cross–section and $\sigma_{\rm g}$ is the grain extinction cross–section per H–nucleon. The factor of 2 in the denominator takes into account the fact that the definition of $\sigma_{\rm g}$ is per H–nucleon, whereas ζ is defined per H_2 molecule (Woodall et al. 2007). We have recalculated the values of $p_{\rm i}$ using the same cross–section data as in Section 4.1.1 and the (high resolution) cosmic–ray–induced H_2 emission spectrum, $P(\nu)$ (Gredel, personal communication, as depicted in Figure 1 of Gredel et al. 1989). This spectrum has been normalized to take account of the various transition probabilities and excitations per cosmic ray ionization. Thus it includes contributions from excitations to

Table 4.3. Range of values for the radiation field strength χ relative to the Draine interstellar radiation field and gas density n used for our three models F1, F2 and F3.

Model	Definition	χ	n / cm^{-3}
F1	Standard	10	10^3
F2	Bright	10^5	10^3
F3	Dense	10	$10^{5.5}$

the $B^1\Sigma_u^+$, $B'^1\Sigma_u^+$, $B''^1\Sigma_u^+$, $C^1\Pi_u$, $D^1\Pi_u$, and $D'^1\Pi_u$ Rydberg states and to the valence $E, F^1\Sigma_g^+$ and $a^3\Sigma_g^+$ states. Also included are excitations into the repulsive $b^3\Sigma_u^+$ state and the vibrational levels of the ground state ($X^1\Sigma_g^+$), together with cascades to the $B^1\Sigma_u^+$ state. The ratio of the H_2 populations in J = 0 to J = 1 is taken to be 1 : 3. Results of these calculations are also given in Tables 4.1 and 4.2. Note that we have not recalculated the cosmic–ray–induced rate for the photodissociation of CO, for which multiple line overlap occurs and the values given in the UMIST06 database are reasonably accurate.

For 24 of the 61 reaction channels, cosmic–ray–induced photoreaction efficiencies are not specified in the UMIST06 database. Of the remaining 36 (excluding CO), all but 15 agree to within a factor of 2. Relative to the UMIST06 database we found the rates for the ionization of NH, the dissociation of CN and the dissociation of H_2CO to H_2 + CO are all lower by factors of $\sim$ 12, 25 and 10 respectively.

4.2 The Chemical Model

To investigate the impact of the shielding mechanisms discussed in Sections 4.1.2 and 4.1.3 on chemical abundances, we implement them in a time– and depth–(A_v) dependent chemical model. Since we are not attempting to replicate any specific astrophysical environment, we adopt physical and chemical parameters similar to the PDR code benchmarking effort of Röllig et al. (2007). These are summarised in Tables 4.3 and 4.4. In particular, we do not attempt to solve the equations of thermal balance and the temperature and density (T and n) are fixed. We also assume that the PDR is dynamically static so that the extinction, A_v, is only a function of position. The incident radiation field used is the standard interstellar radiation field described in Section 4.1.1 scaled by the factor χ.

Our chemistry consists of 7 elements in 81 gas–phase species (Table 4.5) linked by 1129 reactions. For the sake of simplicity, we do not include any gas–grain chemical reactions. In any case, at the assumed temperature of 50 K, gas–grain interactions will have negligible

Table 4.4. Values taken by physical parameters and elemental abundances X relative to hydrogen.

Parameter	Value
Temperature, T	50 K
CR Ionization Rate, ζ	$5 \times 10^{-17}\,\mathrm{s}^{-1}$
Grain Albedo, ω	0.5
X(He)	10^{-1}
X(C)	10^{-4}
X(O)	3×10^{-4}
X(N)	8×10^{-5}
X(Na)	10^{-6}
X(S)	10^{-6}

Table 4.5. List of chemical species included in our model.

Chemical Species
H, H^+, H^-, H_2^+, H_3^+
He, He^+, Na, Na^+, e^-
C, C^+, C^-, CO, CO^+, CH, CH^+, CH_2, CH_2^+
CH_3, CH_3^+, CH_4, CH_4^+, CH_5^+
N, N^+, NH, NH^+, NH_2, NH_2^+, NH_3, NH_3^+, NH_4^+, N_2, N_2^+, N_2H^+
O, O^+, O_2, O_2^+, OH, OH^+, H_2O, H_2O^+, H_3O^+
HCO, HCO^+, H_2CO, H_2CO^+, CO_2, CO_2^+, HCO_2^+
CN, CN^+, HCN, HCN^+, HNC
NO, NO^+, HNO, HNO^+, $HCNH^+$, H_2NC^+, $HNCO^+$, H_2NO^+
S, S^+, HS, HS^+, H_2S, H_2S^+, H_3S^+
CS, CS^+, C_2S, C_2S^+, HC_2S^+, HCS, HCS^+, H_2CS^+

significance. The elemental helium, carbon and oxygen abundances are the same as in Röllig et al. (2007) while for nitrogen, sodium and sulphur, which were absent from their models, we take abundances from Savage et al. (1992) and Asplund et al. (2005). Reaction data and rate coefficients are taken from the UMIST06 database except where we calculated updated photorates as described in Section 4.1. We assume a fixed cosmic ray ionization rate of $5.0 \times 10^{17}\,\mathrm{s}^{-1}$ and a grain albedo of 0.5. The cosmic–ray–induced photodissociation rate for CO is modified by a temperature–dependent factor of $(T/300\,\mathrm{K})^{1.17}$ (Woodall et al. 2007) yielding a factor of eight reduction in the rate at 50K. For species with no cross–section data available we use a default cosmic–ray–induced photoreaction efficiency of 200 broadly representative of the average values for other species.

The incorporation of the various shielding effects uses a number of simplifying approximations to allow for a focused study on the effects of H_2 and CI shielding. As in other studies we treat the attenuation of the radiation field by dust as separable from other

processes, described by the usual exponential dependences on extinction of the photorates. Strictly speaking, these dependences should be modified to take account of the CI absorption, but such variations are relatively minor as compared to the other effects reported in this paper.

For standard PDR models (e.g. model F1 of Röllig et al. 2007), the transition from atomic to molecular hydrogen occurs at $A_v \sim 0.1$, the CII $\rightarrow$ CI transition is at $A_v \sim 1.0$ and carbon is mostly converted to CO for $A_v > 2.0$. Deeper than this we can assume sufficient columns of H_2 and CO have formed that they are completely self–shielded and the effective photodissociation rates for these species are zero. For the same reason, when we investigate the effects of the inclusion of mutual shielding by molecular hydrogen lines, we take the extreme case scenario described in Section 4.1 and set the photodissociation rates of N_2 and CN to zero. We therefore restrict our investigation to $A_v > 2.0$ where shielding by H_2 and CO take this simple form and do not attempt to consider the details of the PDR itself, which we defer to a later study. Within our model we can therefore crudely switch the H_2 line mutual shielding on or off and assign a fixed optical depth for the carbon continuum across the whole extinction range so as to investigate the interplay between the two mechanisms and their effect on the chemistry.

4.3 Results

Examples of the key results obtained from the model are given in Figures 4.1 and 4.2. Our model was tested against the results for C^+, C and CO abundances in model F1 of Röllig et al. (2007). We are able to reproduce abundances of C^+, C and CO at positions well within the PDR, while observing typical sensitivities to the C/O ratio, the temperature dependence of the cosmic–ray–induced CO photodissociation rate and the chemical age. For a direct comparison we have evolved the chemistry to equilibrium. This typically takes $\sim 3 \times 10^6$ years, although we integrate to 100 Myr so as to be consistent with their calculations.

The figures give a clear comparison of the chemical structure that is obtained with and without the implementation of H_2 mutual shielding and carbon continuum shielding. We present relative abundances for observable molecules showing variations of factors greater than 10 as a result of these two mechanisms. The results are given in Figure 4.1 for the models F1, F2 and F3. We also show abundances for C^+, C and CO (PDR carbon transition

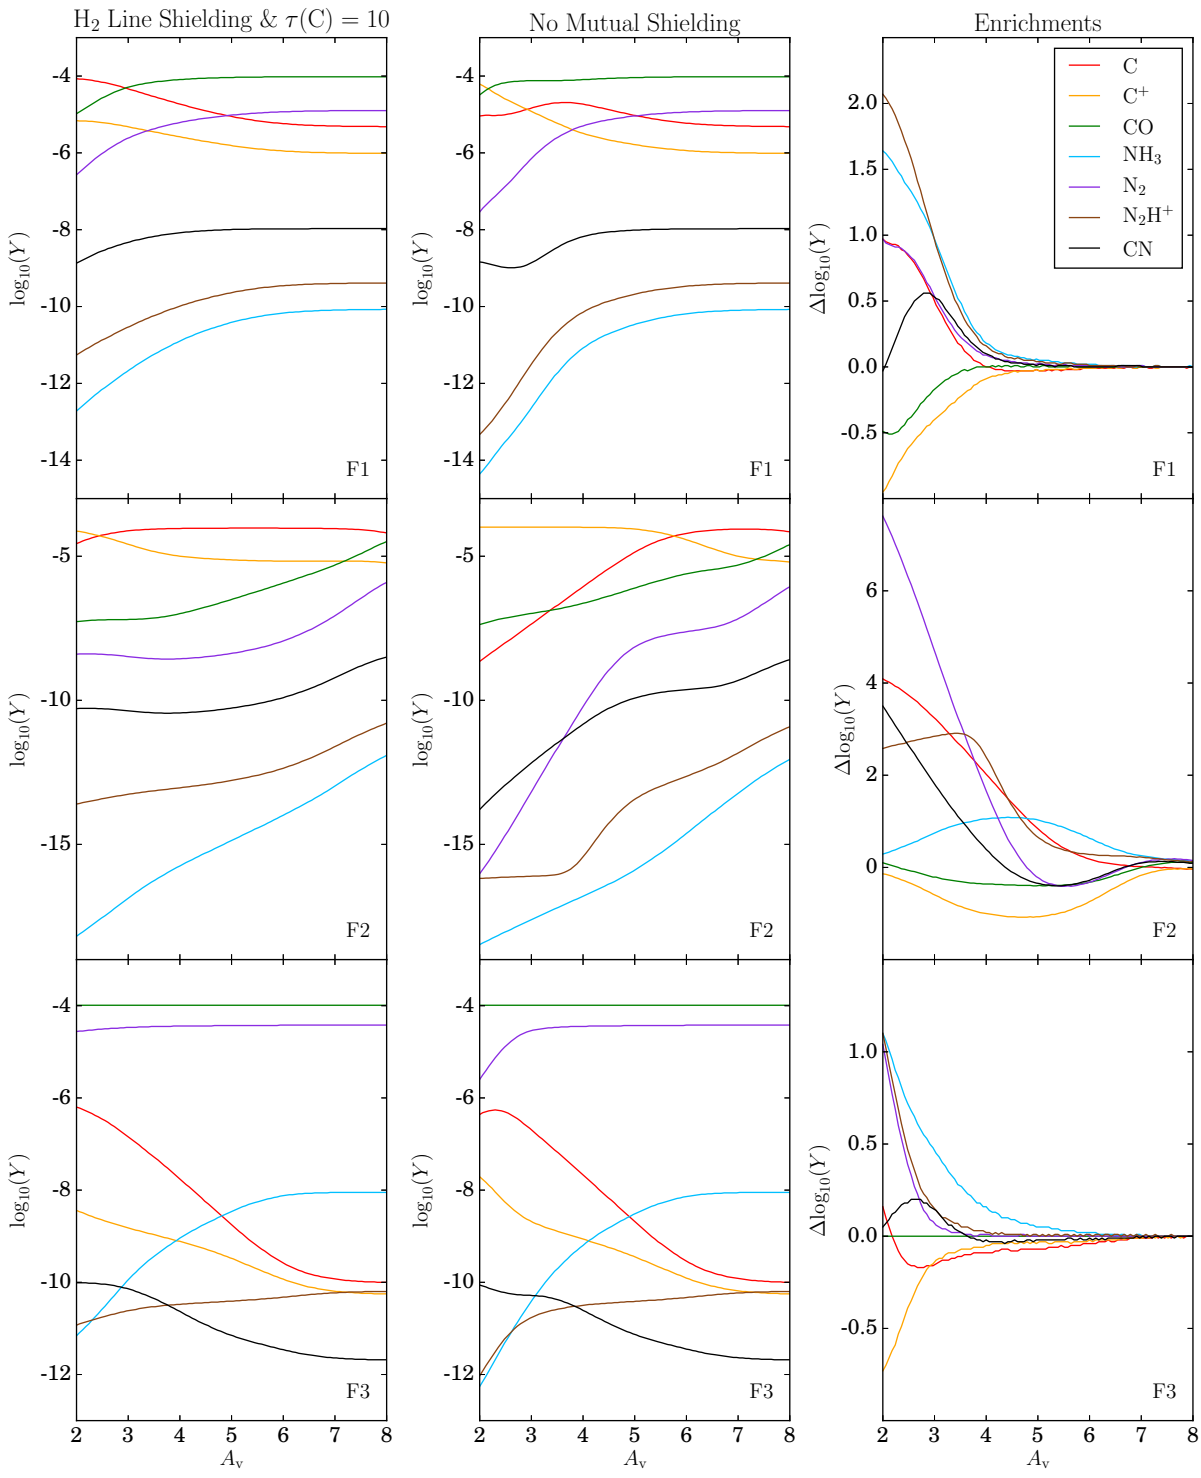

Figure 4.1. Abundances Y for models with H_2 mutual shielding and $\tau(C) = 10$ (left), without H_2 mutual shielding and $\tau(C) = 0$ (centre) and the relative enrichments between the two (right). Top: Standard (F1), Middle: Bright (F2), Bottom: Dense (F3).

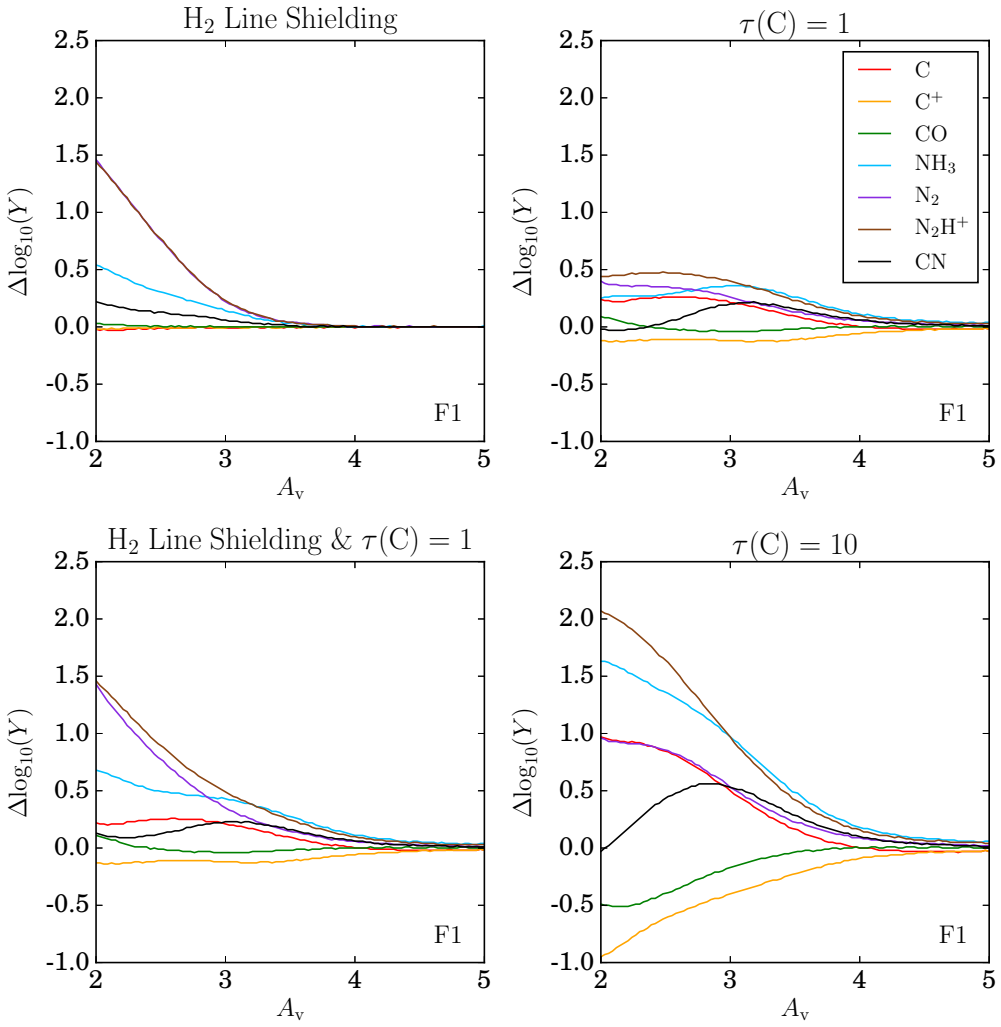

Figure 4.2. Enrichments from model F1 for a range of H_2 and CI shielding regimes, all relative to the case where there is no H_2 mutual shielding and $\tau(C) = 0$. Note that in the first panel, the curves for N_2 and N_2H^+ coincide.

species) as well as N_2 and CN (species shielded by H_2) for reference. Probably the most important effect on the chemistry is the shielding of N_2, for which the photodissociation is completely suppressed by both mechanisms. Thus, significant enrichments are visible in all three environments and particularly in the range of extinctions $A_v \sim 2-4$. This abundance excess obviously drives a more vigorous nitrogen chemistry so that the abundances of other nitrogen bearing species, most notably N_2H^+, NH_3 and CN, are significantly enhanced. The photoionization and photodissociation of NH_3 are both partially shielded by the carbon continuum, while CN is fully shielded by molecular hydrogen. As is usually the case for molecular ions, photodissociation of N_2H^+ is not significant and has been excluded from our chemical network. The presence of significant column densities of nitrogen-bearing

molecules at relatively low extinctions is an important result. However, it is also interesting to note that so *few* species are significantly enhanced by the shielding processes. For example, one might have expected to see large increases in HCN and HNC, but while their enrichments are tightly correlated they are still limited to factors of less than 10. Similarly, despite the strong enrichments in NH_3, the simpler hydrides NH and NH_2 show correlated enrichments but of less than an order of magnitude.

One issue with these particular models is the assumption that the CI continuum is very optically thick; with $\tau(C) = 10$ at $A_\mathrm{v} = 2$. Using a typical relation for $A_v = 6.289 \times 10^{-22} N_{\mathrm{H_{Total}}}$ (Röllig et al. 2007) depending on the total hydrogen column density $N_{\mathrm{H_{Total}}}$ allows us to approximate the atomic carbon column density. For model F1, a column of roughly 10^{16} cm^{-2} is obtained at $A_\mathrm{v} = 2$, rising to $\sim 10^{17}$ cm^{-2} at $A_\mathrm{v} = 4$, implying a CI optical depth of order unity, but somewhat less than $\tau(C) = 10$. The other models yield lower values, with the bright model not achieving a column of 10^{17} cm^{-2} until $A_\mathrm{v} \sim 6$ and the dense model only achieving a maximum column of roughly $\sim 10^{15}$ cm^{-2}. However, the A_v–$N_{\mathrm{H_{total}}}$ relationship depends on the gas to dust ratio, a quantity which is not only poorly constrained but can also vary by orders of magnitude for different astrophysical environments. So whilst the results of Figure 4.1 may only be valid beyond a certain extinction for our given A_v–$N_{\mathrm{H_{total}}}$ relation, the results become valid at lower A_v as the dust–to–gas number ratio falls (since the CI column density would increase for a given A_v). We also note that for model F2 the C $\rightarrow$ CO transition is pushed much deeper to $A_\mathrm{v} \sim 8$ by the more intense radiation field. This probably invalidates our assumption of CO being optically thick and perfectly self–shielded by $A_\mathrm{v} = 2$. However, one effect of this would be to increase the abundance of atomic carbon at intermediate depths and hence the optical depth of the CI continuum shielding CO.

Figure 4.2 shows the range of enrichments achieved when using different levels of shielding in model F1. The final panel shows that as carbon becomes optically thick, the resulting enrichments are very similar with or without H_2 line shielding, even though CN is not fully shielded by the carbon ionization continuum. More importantly, the first panel shows that in environments where carbon column densities are low and only H_2 line shielding occurs, significant enrichments in N_2, N_2H^+ and NH_3 are still possible. Qualitatively similar behaviour is also seen for models F2 and F3. The exception is CN, which in model F1 is seen

to be coupled to the enrichment in atomic carbon through the formation channel:

$$C + NO \rightarrow CN + O.$$

Since the H_2 photodissociation lines become optically thick at low extinctions ($A_v \sim 0.1$), the effects of H_2 mutual shielding will always be significant for $A_v > 2$. We therefore expect the true chemical enrichment for $A_v > 2$ to be some intermediate between these two plots; that is to say, the effects of mutual shielding by the H_2 photodissociation lines will always be present even if the carbon ionization continuum is optically thin.

4.4 Conclusions and Discussion

We have recalculated the interstellar photodissociation and photoionization rates of a range of interstellar molecules, both for direct interstellar and for cosmic–ray–induced radiation fields. We have also calculated direct photorates which include the effect of partial blocking of the incident radiation field by the carbon ionization continuum.

The specific quantitative results that can be drawn from our model are clearly limited by its simplicity. However, using approximations for shielding of the incident radiation field deep in PDRs by molecular hydrogen and atomic carbon, we present the result of significant (more than an order of magnitude) enrichments in the nitrogen-bearing species N_2, N_2H^+, NH_3 and CN. These results are seen for a range of physical parameters, including density, radiation field intensity and carbon column density, but should be most prominent in dust poor environments. These species are all important observationally in inferring various characteristics of the PDRs in which they are detected. For example, Bayet et al. (2009) have used PDR models to show that CN traces the density and relative nitrogen abundance of PDR dominated galaxies. First results from the *Herschel Space Observatory* towards W31C have shown nitrogen hydride (NH_3, NH_2 and NH) abundances that cannot be explained using simple gas–grain or PDR chemical modelling (Persson et al. 2010). Other nitrogen-bearing species such as HCN and HNC, also believed to trace dense molecular gas, show smaller (< 0.5 magnitudes), correlated enrichments. They are often used in both galactic (e.g. Roberts et al. 2011) and extragalactic (e.g. Kamenetzky et al. 2011) environments to classify the gas chemistry as photon dominated or not. It is clear from our work that the shielding by H_2 and C have a large impact on the abundances of such nitrogen-bearing species and both mechanisms need to be included in future PDR models.

Current ongoing studies suggest these mechanisms have only a limited effect on column densities in high–extinction clouds and cannot explain, for example, the high column densities of nitrogen hydrides seen in the interstellar medium towards W31C (Persson et al. 2010). In future work we will apply the shielding factors to more rigorous PDR models with self–consistently calculated column densities to investigate the effects on the chemistry in regions of lower visual extinction ($A_{\rm v} < 2$). This is particularly important given that our observed enrichments appear more pronounced at lower extinctions where molecular shielding by H_2 becomes more significant relative to dust extinction. It would also be useful to do a thorough reanalysis of the cosmic–ray–induced photodissociation rate of CO, which is the dominant source of atomic carbon in regions of high extinction ($A_{\rm v} > 4 - 5$).

Chapter 5

Champagne Flutes and Brandy Snifters: Modelling Protostellar Outflow–Cloud Chemical Interfaces

The work presented in this chapter is based on the paper by Rollins et al. (2014) in collaboration with J. M. C. Rawlings, D. A. Williams and M. P. Redman.

Our understanding of the formation and evolution of protostars has evolved into a highly complex and three dimensional model. Since the one dimensional inside out collapse theory of Shu (1977), we now observe many distinct features such as discs, outflows, jets, HII regions and turbulence to be important in governing these systems. To develop an understanding of such features is to further our understanding of the star formation process on all scales. Of particular interest are the molecular outflows observed in low–mass protostars. They have the potential to liberate accreting material of angular momentum as well as impacting dynamically on the surrounding molecular core. This creates a turbulent mixing zone which can promote a rich and unique chemistry. Spectroscopic observations now show that many molecules are harboured in the outflow–cloud interface of these objects, particularly carbon monoxide (Yıldız et al. 2010), water (Kristensen et al. 2010, Nisini et al. 2010, van Kempen et al. 2010, Herczeg et al. 2012) and HCO^+ (Hara et al. 2013).

Codella et al. (2006) have also identified methanol, deuterated water, OCS and H_2CS in the high–mass analogue CepA–East. An understanding of the dynamics and chemistry of such systems is therefore clearly important for a complete understanding of the star formation process.

Observations of low–mass protostellar outflows exhibit two particularly interesting chemical features that can help diagnose the physical conditions. The first is an apparent observational preference towards wide angled cavities mapped in carbon monoxide within a few thousand astronomical units of the central source (for example, B5 IRS1, Langer et al. 1996, HH 46/47, Arce et al. 2013, Rawlings et al. 2013*a*). We liken the observed morphologies to "brandy snifters" – short hour–glass shapes with wide opening angles at the base and a given height at which the width becomes maximal. By analogy, there are few observed molecular outflows shaped like "champagne flutes" – narrow opening angles at the base, tall and narrow all the way up. The apparent absence of "champagne flute" shaped outflows from observations is an issue that has not been adequately addressed by current models. Many attempts have been made to understand the carbon monoxide morphologies in terms of the dynamics of the outflow–cloud system. Initially, Gueth et al. (1996) presented a model for a precessing cavity and argued that the observed CO maps were a combined result of those dynamics with radiative transfer effects. More recently, Li et al. (2013) modelled the cavity shape as a result of turbulent entrainment and produced synthetic observations using radiative transfer for the high–mass source G240.31+0.07 to be compared with observations by Qiu et al. (2009). Their results suggest that "champagne flute" shaped outflows should in fact be observable. While both studies produced qualitatively accurate synthetic observations for individual objects, they made the simplifying assumption of a fixed CO abundance with polar angle along the interface. No attempt was made to consider three dimensional variations in the carbon monoxide chemistry throughout the interface as the origin of the observed morphology. Visser et al. (2012) attempted such a model; a static cavity with heating by ultraviolet (UV) radiation and local C–shocks driving the chemistry. However, by considering only fits to observed objects those authors were unable to address a potential chemical origin of the morphological bias.

Secondly, abundances of HCO^+ in the outflows are observed to be significantly higher than estimated from modelling of dark clouds (e.g. L1527, Hogerheijde et al. 1998). In a previous attempt to model HCO^+ formation, Rawlings et al. (2000) proposed that the high abundance was a result of photoionized carbon reacting with water liberated from dust

grain surfaces by shocks. Follow–up work showed that chemical variations along the cavity wall could reproduce HCO^+ observations (Rawlings et al. 2004). However, the increased abundance is only temporary as the dissociative recombination of HCO^+ is rapid (10–100 years) under typical conditions, especially at temperatures below 100 K. The temporary nature of its enhancement would suggest that in a decollimating outflow, it would only be present near the central source, and requires that cavities steadily grow over time, otherwise the enhancement disappears from the system on a 100 year time–scale.

Photochemistry is certainly not the only possible process driving the chemistry in such systems. In this paper we develop and investigate a model for the chemistry in these cavity walls driven by hot ions from the protostellar ejecta mixing with molecular gas, with the aim of understanding the observations of CO and HCO^+ summarised above. In particular, we speculate that the degree of interaction between the outflow and the interface (and hence the dynamically-induced molecular enrichment) is a function of the impact angle between them. This could help explain the observed bias towards certain morphologies.

In Section 5.1 we describe a parametric model for the dynamics of outflows to quantify the ion injection rate and also the key dynamical time–scales. In Section 5.2, we present the chemistry used in our model, discussing assumptions on the initial conditions and chemical time–scales. The results of our combined chemo–dynamical model are presented in Section 5.3 and discussion with respect to observations and free parameters in Section 5.4. Our final conclusions are presented in Section 5.5 along with brief proposals for further work.

5.1 Cavity Injection Model

In order to quantify the rate at which material may be injected into the outflow–cloud interface, we must first be able to characterize its dynamical properties, critically the shape and velocity and how these may vary as a function of age. A number of possible origins for the structure of such cavity walls have been discussed, most notably magnetohydrodynamic (Bürzle et al. 2011, Bate et al. 2014) and turbulent entrainment (Cunningham et al. 2006, Li et al. 2013). However, for the sake of our model, it is not important to understand the physical origin of the cavities but rather to be able to parametrize their typical observed shapes. An elegant example of this approach was presented by Cantó et al. (2008) for the cavity wall of Barnard 5 IRS 1. Their dynamical model is of an outflow coming from the central source with constant total mass loss rate ($\dot{M}_0$) and velocity (v_0) and steadily

Table 5.1. Physical parameters for the cavity in B5 IRS 1 (Cantó et al. 2008) used for our dynamical model.

Cavity Parameters	Values
Scale Length, r_0	$350\,\mathrm{au}$
Mass Loss Rate, $\dot{M}_0$	$2.9 \times 10^{-6}\,\mathrm{M_\odot\,yr^{-1}}$
Outflow Velocity, v_0	$100\,\mathrm{km\,s^{-1}}$
Decollimation Rate, ϵ	$1.05 \times 10^{-4}\,\mathrm{rad\,yr^{-1}}$
Isothermal Coefficient, κ	$6.7 \times 10^{14}\,\mathrm{g\,cm^{-1}}$

decollimating at a constant angular rate (ϵ, the rate at which the outflow opening angle increases in radians per unit time). The outflow impacts on an isothermal sphere of density $\rho(R) = \kappa R^{-2}$ where R is the radius and κ is a scaling coefficient treated as a free parameter. Those authors were able to reproduce the observed shape and inferred dynamical properties whose parameters are summarised in Table 5.1. We choose to adopt this model since it not only gives the parametrized shape of a typical wide–angled source, but also defines cavity shapes as a function of dynamical age; younger ages correspond to narrower cavities, and vice versa. This allows us to directly test the chemical origin of the observational bias towards wider opening angles.

Given the radial shape equation $R_\mathrm{S}(\theta, t_\mathrm{dyn})$ for the cavity as a function of polar angle θ and cavity age t_dyn, we assume an interface forms between the cavity and gas envelope with thickness equal to 10% of the dynamical radius (a value suggested by observations e.g. Langer et al. 1996). Here, the radially outflowing material impacts on the interface at an angle ψ (as derived in Appendix A and shown in Figure 5.1) given by the equation:

$$\psi = \arctan\left(\frac{R'_\mathrm{S} - R_\mathrm{S}\tan\theta}{R'_\mathrm{S}\tan\theta + R_\mathrm{S}}\right) + \theta - \frac{\pi}{2} \tag{5.1}$$

where R'_S is the derivative $\mathrm{d}R_\mathrm{S}/\mathrm{d}\theta$. This angle is then used with the other parameters to characterize the rate at which material from the outflow mixes turbulently at each point into the interface, as derived in Appendix B:

$$\dot{n}_\mathrm{inj} = \frac{5\dot{M}_0}{\pi m_\mathrm{H}[R_\mathrm{s}(\theta)]^3}\frac{\sin\psi}{1-\cos[\epsilon(t_\mathrm{dyn} - \frac{R_\mathrm{S}(\theta)}{v_0})]}\,\mathrm{cm^{-3}\,s^{-1}} \tag{5.2}$$

where m_H is the atomic mass of hydrogen. The factor $\sin\psi$ is a nominal mixing efficiency

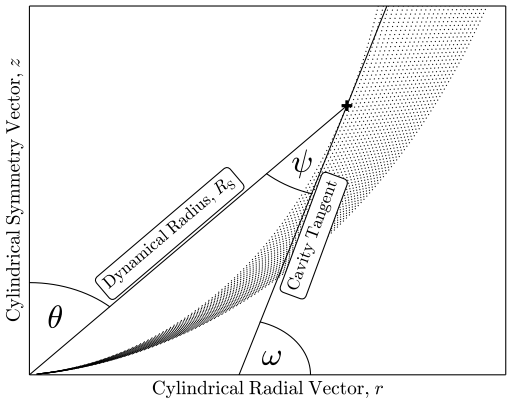

Figure 5.1. Cross–section of the model geometry for an example cavity–cloud interface. The shaded area shows the interface with its thickness exaggerated to be 50% of the dynamical radius for clarity. The tangent to the cavity wall is shown for a given point along with the corresponding polar angle, θ, tangent gradient, ω, and turbulent mixing angle, ψ.

chosen on the assumption that only the component of the outflow normal to the cavity wall will be able to mix turbulently into the interface. A key prediction of this work is that this mixing angle is the key parameter in determining the local injection rate at each point in the interface. We expect that only those regions where ψ takes a large value are able to form significant abundances of CO and HCO^+ due to the enhanced chemical interaction between the gas from the envelope and the injected material. The choice of function for this mixing efficiency and its impact on the model are considered in Section 5.4.

There are three key time–scales associated with such a system. One is the time–scale on which the cavity grows, which is simply $R_S/(\mathrm{d}R_S/\mathrm{d}t_{dyn})$. We note that this tends towards approximately one fifth of the dynamical age of the cavity on all angular scales. The second is the time–scale on which the interface density increases and its composition becomes dominated by the material injected from the outflow into the interface, given by $n_0/\dot{n}_{inj}$ where n_0 is the initial particle density in the interface. While the interface density is a free parameter of our model, typical values for the density and cavity age give a time–scale many orders of magnitude longer than the dynamical age, so that the composition of the interface is never dominated by the injected material. However, even if the composition of the cavity is not dominated by hot ions from the outflow, their presence can still drive the dominant chemical processes.

Instead, it is the time–scale on which the injection rate varies, $\tau_{inj} = \dot{n}_{inj}/(\mathrm{d}\dot{n}_{inj}/\mathrm{d}t_{dyn})$, that is important to our problem and visualized in Figure 5.2. Again, we see that it

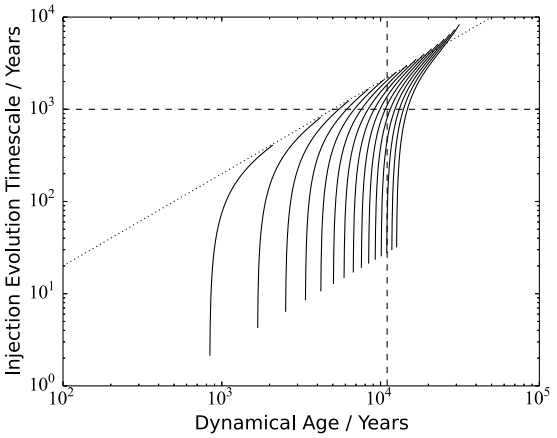

Figure 5.2. Iso–angle tracks for the time–scale on which the injection rate (Equation 5.2) changes. Each solid line represents one fixed value of θ between 4 and 60 degrees and shows the time–scale evolving with the dynamical age of the growing cavity at that polar angle. The dotted line $\tau_{\rm inj} = t_{\rm dyn}/5$ is an approximate asymptote and shows that this time–scale approaches approximately 20% of the dynamical age for all angles. The dashed vertical line shows the dynamical age of B5 IRS1 (11 kyrs) inferred by Cantó et al. (2008). The dashed horizontal line demonstrates that for most angles, by a cavity age of 11 kyrs, the time–scale on which the injection rate changes is more than 1 kyr. As discussed in Section 5.2, the chemistry of the interface reaches a quasi–equilibrium within this time. 1 kyr is therefore a suitable chemical integration time for our decoupled chemical model at all but the angles close to the opening angle. For younger cavities, shorter chemical time–scales will be suitable.

approaches approximately a fifth of the dynamical age of the cavity for all angles. As a result, any chemistry driven purely by the injecta that can reach equilibrium on time–scales an order of magnitude shorter than the dynamical age can be modelled independently of the dynamics since the injection rate is effectively constant. However, if the chemical time–scales are significantly longer than the dynamical age then it is necessary to integrate the dynamics and chemistry coupled in time. As we shall further argue in Section 5.2, the much simpler first situation holds for our model, with chemical equilibrium times of order 1 kyr being typical.

5.2 Chemical Model

Given that the outflowing material interacts dynamically with the envelope to form the observed interface, we now wish to model the chemical interactions between the two gases as they mix turbulently. Our chemical model borrows heavily from the work of Rawlings et al. (2000) in modelling outflow interface chemistry. In their work, HCO^+ is formed as a

product of the reaction between C^+ and water:

$$C^+ + H_2O \rightarrow HCO^+ + H.$$

The available C^+ was presumed to have been formed from the photoionization of atomic carbon by UV radiation given off in the bow shock associated with the expansion of the cavity. At the same time, the shock was assumed to liberate the icy mantles of dust grains, giving an initial gas–phase chemical composition equal to typical dark cloud dust mantle compositions, including the necessary water for processing. Indeed, Bachiller & Pérez Gutiérrez (1997) observed a rich array of species produced in the shocked regions of L1157, including some believed to be processed from interstellar ices. More specifically, Nisini et al. (2010) see emission from water that is spatially correlated with the outflow shock of L1157. While this reaction pathway could generate high fractional abundances of HCO^+ ($>10^{-7}$), it was only a temporary feature as the dissociative recombination reaction:

$$HCO^+ + e^- \rightarrow CO + H$$

was rapid at typical conditions, and especially below 100 K due to its strong temperature dependence. Hence, the high abundances in Rawlings et al. (2000) were only realized for 10–100 years.

Our chemical model includes the species shown in Table 5.2. We consider similar initial conditions to those adopted by Rawlings et al. (2000), with the initial gas–phase fractional chemical abundances, Y_0, coming from observationally inferred values for the ice composition of dark clouds. These values are given, as a function of the total elemental abundances, X, in Table 5.3. The gas is thus assumed to be fully molecular, apart from an atomic hydrogen abundance of $1\,cm^{-3}$, established by the balance between cosmic–ray–induced destruction and grain surface formation of H_2.

The abundances of the chemical species are allowed to evolve with the gas–phase reactions given in the RATE06 database (Woodall et al. 2007). In addition, H^+, C^+ and O^+ are continuously injected at a rate given by Equation 5.2 multiplied by the appropriate

Table 5.2. List of chemical species included in our model.

Chemical Species
H, H^+, H_2, H_2^+, H_3^+, H^-
He, He^+, Na, Na^+, e^-
C, C^+, CH, CH^+, CH_2, CH_2^+
CH_3, CH_3^+, CH_4^+, CH_4, CH_5^+
O, O^+, O_2, O_2^+, OH, OH^+, H_2O, H_2O^+, H_3O^+
CO, CO^+, HCO, HCO^+, H_2CO, H_2CO^+

Table 5.3. Total elemental abundances and initial fractional abundances of gas–phase species, relative to hydrogen nucleons.

Parameter	Definition	Value
X(He)	He/H	0.0763
X(C)	C/H	2.0×10^{-4}
X(O)	O/H	3.02×10^{-4}
X(Na)	Na/H	5.2×10^{-6}
$Y_0(H_2O)$	H_2O/H	0.9 X(O)
Y_0(CO)	CO/H	0.1 X(O)
$Y_0(CH_4)$	CH_4/H	0.84 X(C)
Y_0(C)	C/H	8.5×10^{-3} X(C)

elemental abundance and using the parameters in Table 5.1. Helium is also injected and, following previous studies (e.g. Rawlings et al. 2000), is assumed to be only 10% ionized. Sufficient electrons to maintain overall charge neutrality are also added to the interface. The cosmic ray ionization rate is set to a standard value of $\zeta = 10^{-17}\,\mathrm{s}^{-1}$. We set the UV radiation field in our model interface to zero throughout this paper, considering only cosmic–ray–induced UV photoprocesses. This allows us to compare results of our hot ion chemistry model with those of the photochemistry model (Rawlings et al. 2000). Rawlings et al. (2000) argue that UV photons are created in the bow shocks of the outflow. In fact, the leading initial bow shock of a jet should have propagated a long way from the regions near the central source we are examining. Instead, such jets usually have internal working surfaces that are weaker shocks caused by variations in the jet mass loss rate and speed. These weaker shocks and working surfaces should form a reasonably constant source of high temperature and local UV radiation that can ionize the outflow material near the star but with negligible effect on the interface. van Kempen et al. (2009) argue that for the source HH 46 the radiation field reaching the cavity wall could be as high as 600 Draine. However, factors including weaker shocks in different sources, photoionization of outflowing gas within the cavity and geometric variations along the cavity wall could

all significantly reduce this, especially close to the central source. We use this argument to justify our assumption that the outflow material is ionized while there is negligible UV radiation in the interface itself.

The density and temperature of the interface are poorly understood and are treated as free parameters. A number of works have attempted to infer these properties using observations, models and simulations and can guide us in our choice. Cunningham et al. (2006) simulate turbulent interfaces with densities of 10^3–$10^5\,\mathrm{cm}^{-3}$, while the magnetohydrodynamic simulations of Bürzle et al. (2011) and Bate et al. (2014) are capable of much higher densities (although at earlier times). Temperatures of order 100 K are also seen in the last of these works. Observational studies such as van Kempen et al. (2009) suggest relatively low temperatures in the range of 80–150 K and densities of a few times $10^4\,\mathrm{cm}^{-3}$ for the HH46 outflow, although there are uncertainties since the observed CO 3–2/6–5 line ratio is not expected to change significantly as the kinetic temperature increases above 100 K (van der Tak et al. 2007). These contrast with the predictions of entrainment models where outflow temperatures of over 1000 K are possible (Cunningham et al. 2006, Li et al. 2013).

In Section 5.3 we choose to consider a grid of densities from 3×10^4 to $10^6\,\mathrm{cm}^{-3}$ and temperatures from 100 to 150 K for all reactions. These ranges represent a suitable intermediate between the cold, dense molecular envelope and the hotter, more diffuse outflowing material as they mix turbulently with one another. While high–J CO (Yıldız et al. 2010) and HCO^+ (Kama et al. 2013) observations would seem to demand a much hotter interface, it is possible that such transitions could be collisionally excited by energetic electrons from the injecta due to the typically lower critical densities. In fact, our typical equilibrium ionization fractions of a few times 10^{-4} are close to the prediction by Jiménez-Serra et al. (2006) of 10^{-3} in the magnetic precursor of an outflow C–shock from measurements of the excitation of $H^{13}CO^+$, among other species. We also note that if temperatures were significantly less than 100 K, freeze out of water onto grains would remove the main reagent for the formation of HCO^+ and a non–thermal desorption mechanism would be required to liberate the water ice. While other models may appeal to photodesorption by direct UV radiation, it may also be possible to initiate our model with secondary UV photons due to ionization by cosmic rays in such a case.

Figure 5.3 shows chemical abundance profiles as a function of time for a typical parameter set. As with the photochemical pathway presented in Rawlings et al. (2000), our injection model can give an enhancement in the HCO^+ abundance to $\approx 10^{-8}$, but only on

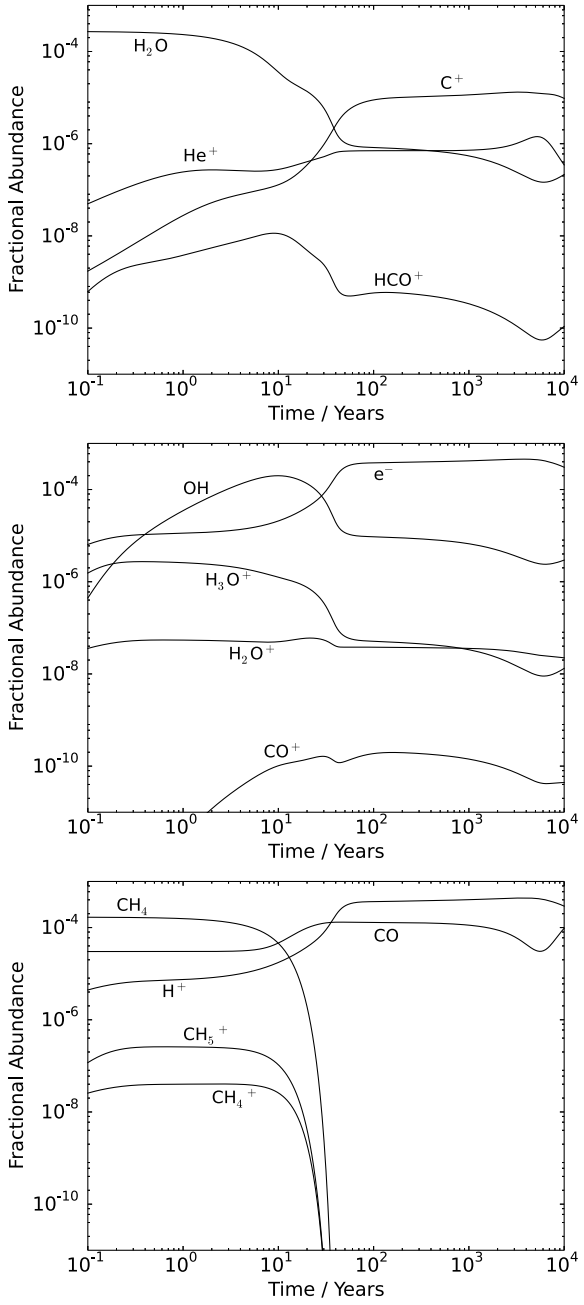

Figure 5.3. Modelled fractional chemical abundances relative to hydrogen near the base of a cavity of age 11 kyrs as a function of chemical integration time. The interface temperature is 150 K and the density is $10^5\,\mathrm{cm}^{-3}$. The abundance of HCO^+ is seen to correlate with a variety of parent molecules including H_2O, OH and CH_4 before they are all destroyed. A quasi–equilibrium is reached after approximately 1 kyr.

time–scales of 10–100 years. Beyond that time, dissociative recombination becomes significant and a lower abundance remains. An approximate chemical equilibrium is reached by the system on time–scales of the order 100 years for all species, roughly independent of the physical parameters. With reference to Figure 5.2, for a cavity of age over 10 kyrs such as B5 IRS1, the injection evolution time–scale is more than 1 kyr on all angular scales and so the interface can reach this equilibrium. We are therefore justified in using a chemical integration time of 1 kyr to represent the equilibrium state of older cavities like B5 IRS1. For younger cavities, this injection evolution time–scale becomes shorter and so a shorter chemical integration time will be necessary. Note that small abundance variations are seen at approximately 10 kyrs due to reactions with He^+, but these are only temporary fluctuations lasting a few kyrs.

As in the work of Rawlings et al. (2000), the main pathway to HCO^+ formation on time–scales of approximately 10 years is the reaction of water from the initial chemical composition of the gas with ionized carbon. The material injected into the interface is a much weaker source of ionized carbon than photoionization, meaning this reaction alone does not produce as much HCO^+ in our model. However, two other routes to HCO^+ are initiated by hydrogen ions present in the injected material in our model but typically absent in photochemical reaction networks. The first involves the formation of the CH_5^+ ion which can then react with carbon monoxide:

$$CH_4 + H^+ \rightarrow CH_4^+ + H$$

$$CH_4^+ + H_2 \rightarrow CH_5^+ + H$$

$$CH_5^+ + CO \rightarrow HCO^+ + CH_4.$$

The other pathway involves the formation of CO^+ ions, either by the cosmic ray ionization of carbon monoxide or the reaction of OH with injected carbon ions:

$$H_2O + H^+ \rightarrow H_2O^+ + H$$

$$H_2O^+ + H_2 \rightarrow H_3O^+ + H$$

$$H_3O^+ + e^- \rightarrow OH + H + H$$

$$OH + C^+ \rightarrow CO^+ + H.$$

CO^+ then readily reacts with available molecular hydrogen to produce HCO^+:

$$CO^+ + H_2 \rightarrow HCO^+ + H.$$

These three reaction pathways together lead to abundances of HCO^+ that are comparable to the results of Rawlings et al. (2000); typically between 10^{-7} and 10^{-9} relative to hydrogen in the first 100 years of chemical integration.

5.3 Results

In this section we present the grid of outputs from our model for a range of parameters. First, we took four cavity dynamical ages ($t_{\rm dyn}$ = 2, 6, 11, 14 kyrs) and calculated the turbulent injection rate (Equation 5.2) as a function of position. We then integrated the resulting chemistry for a chemical age $t_{\rm chem}$ = 10 kyrs at all positions and using a grid of densities ($n = 3 \times 10^4$ to $10^6\,{\rm cm}^{-3}$) and temperatures (100 K, 125 K and 150 K). The spatial abundance profiles of the interfaces in CO and HCO^+ are presented in Figures 5.4 and 5.5 respectively. The 2 kyrs interface is plotted after a chemical integration time of 100 years, the 6 kyrs interface after 300 years and the older 11 and 14 kyrs interfaces after a chemical age of 1 kyr; motivated by the results of Figure 5.2 and Section 5.2.

A number of features are apparent and allow us to constrain the physical conditions of the interface. Over the modelled range of the parameters, the chemical composition is more strongly dependent on density than on temperature. Below densities of $10^5\,{\rm cm}^{-3}$, carbon monoxide is unable to saturate due to the dissociation by injected He^+:

$$He^+ + CO \rightarrow C^+ + O + He.$$

At $10^5\,{\rm cm}^{-3}$ a transition to CO saturation is seen, with the regions of the interface at larger radii (and hence lower ion injection rates) able to sustain a significant amount of carbon monoxide. With higher densities, the injected material has a lower impact on

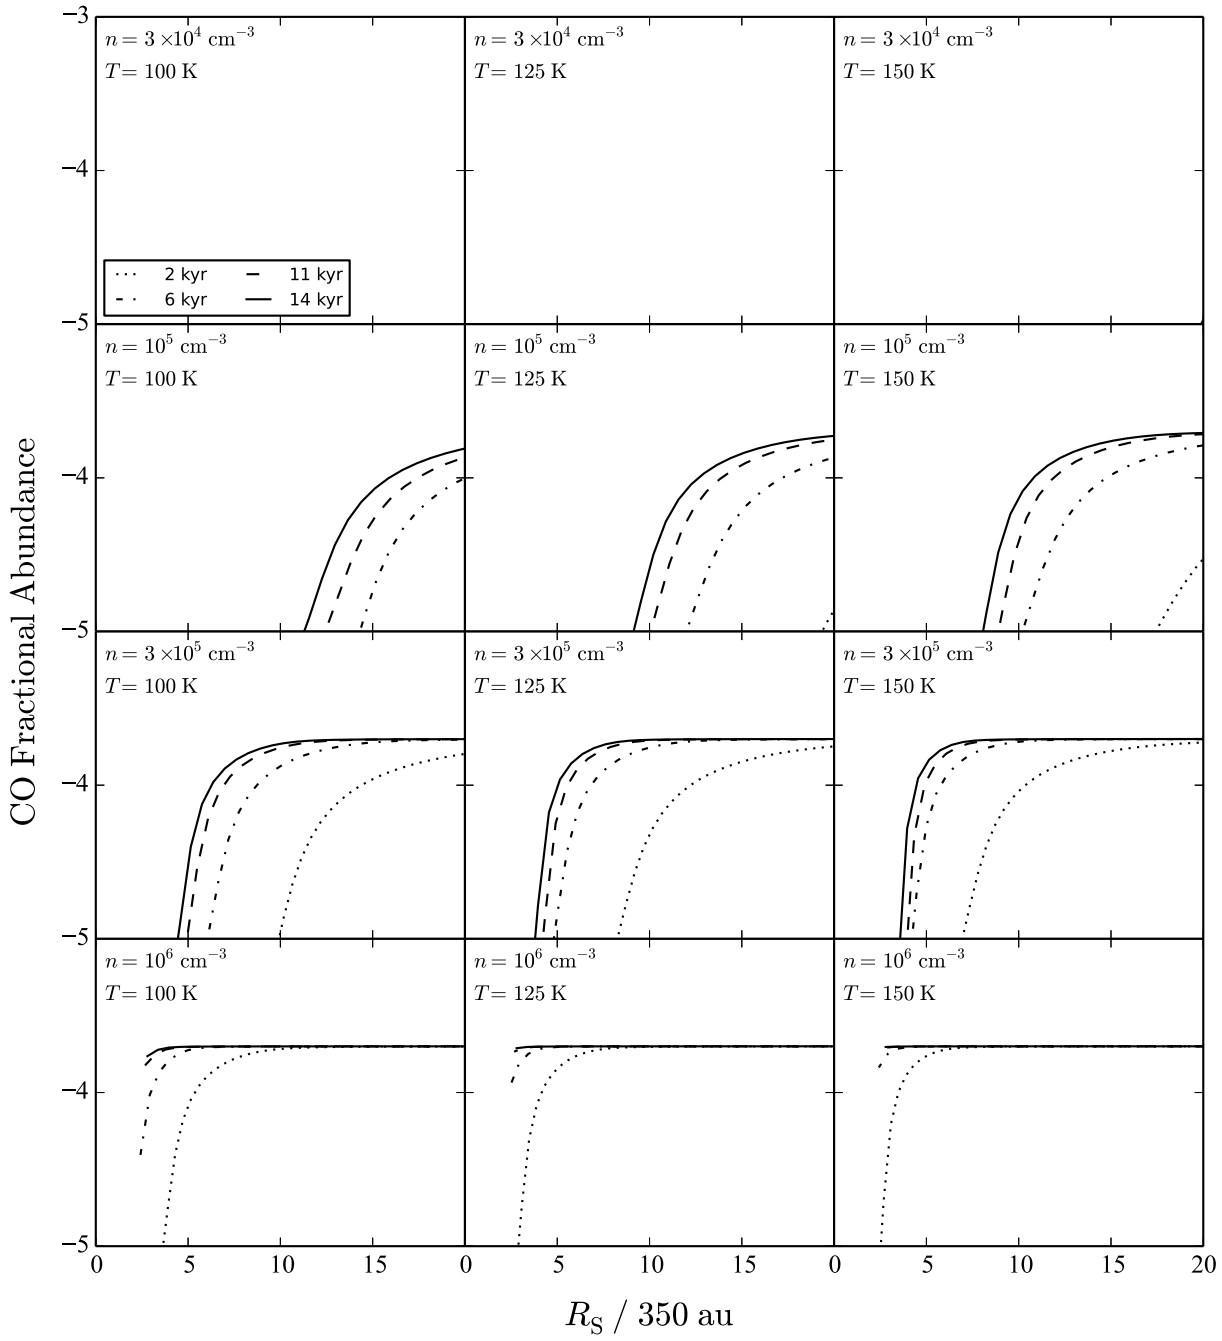

Figure 5.4. Fractional abundances of CO relative to hydrogen for a cavity of dynamical age 2 (dot), 6 (dot–dash), 11 (dash) and 14 (solid) kyrs as a function of distance from the central source. The temperatures used are 100 K (left), 125 K (centre) and 150 K (right). Densities are fixed at $3 \times 10^4\,\mathrm{cm}^{-3}$ (top) through $10^5\,\mathrm{cm}^{-3}$, $3 \times 10^5\,\mathrm{cm}^{-3}$ and $10^6\,\mathrm{cm}^{-3}$ (bottom). The chemical integration times at which the abundances are plotted are 100, 300, 1000 and 1000 years for the 2, 6, 11 and 14 kyrs cavities respectively.

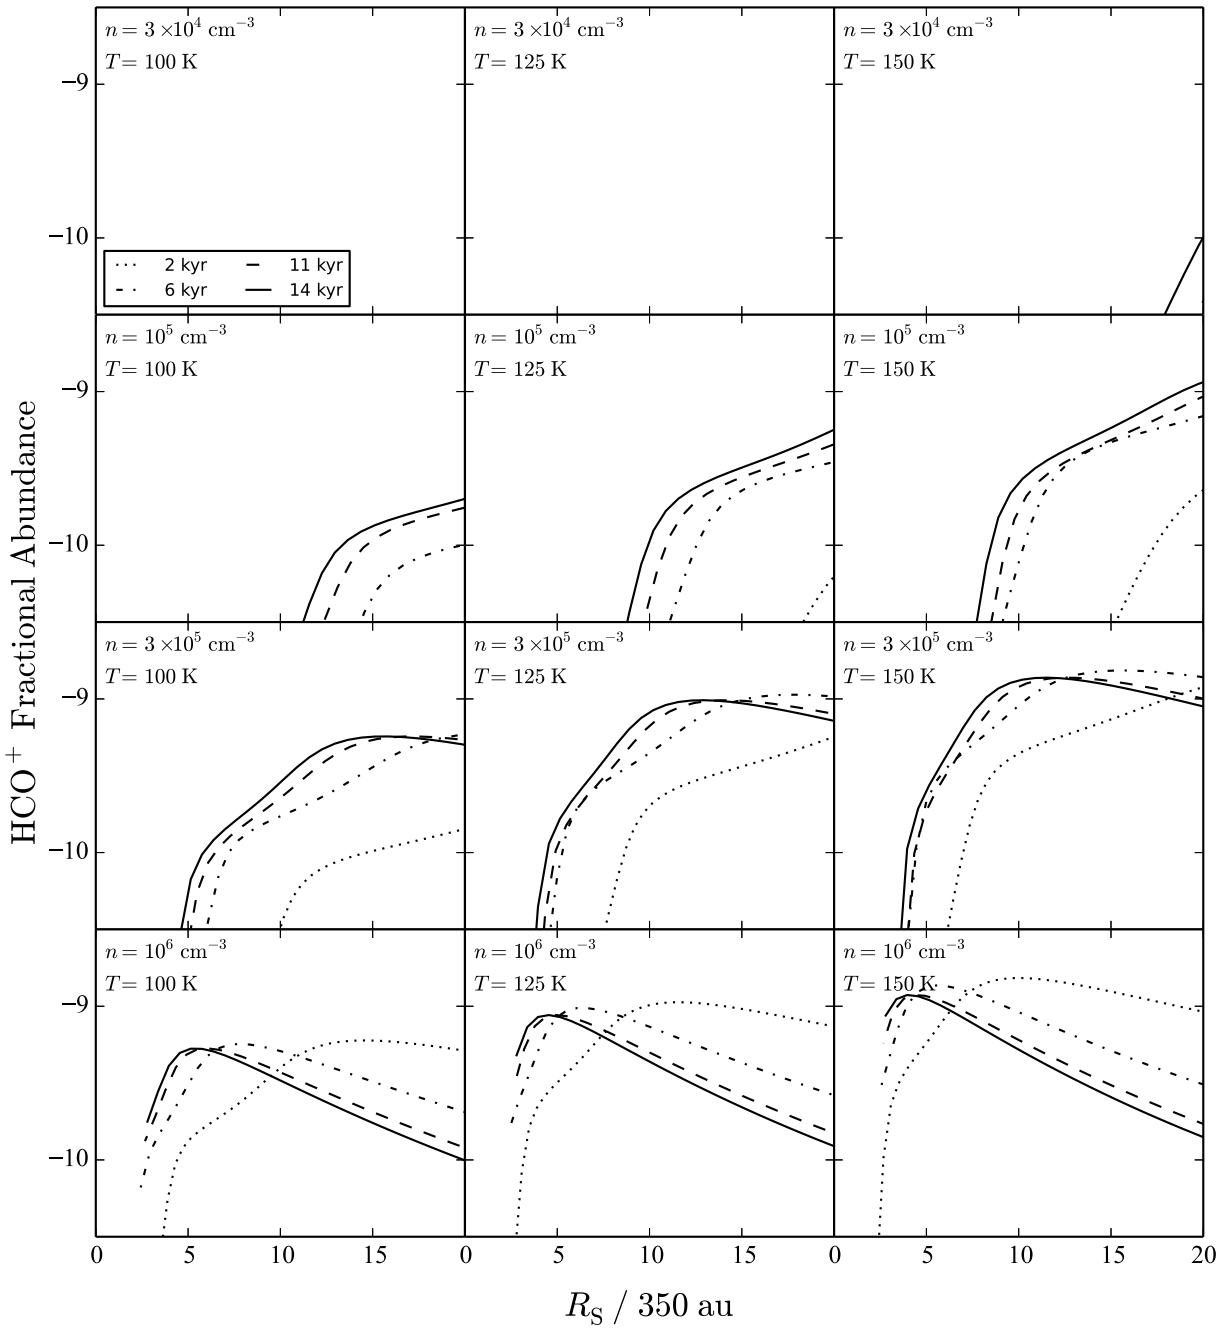

Figure 5.5. The same as Figure 5.4 except showing the fractional abundances of HCO^+ relative to hydrogen.

molecular abundances and saturation is achieved for all but the regions closest to the central source. The CO dependence on temperature is minimal over the modelled range. What is of most interest is that at the density of $10^5\,\mathrm{cm}^{-3}$ there is a marked variation in CO abundance patterns over the age of the cavity. The youngest, narrowest interface modelled has a severely diminished carbon monoxide fractional abundance of less than 10^{-6} over all angles. By comparison, the 6 kyrs cavity exhibits a small band of CO saturation at radii greater than approximately 4000 au. The two older and wider angle interfaces show a significantly increased band of CO saturation in to closer radii. We argue in Section 5.4 that this variation is not sensitive to the choice of chemical time–scales at each cavity age.

This feature of more CO at wider angles in the model can be compared with the observational preference towards wide angled outflows observed in CO. It raises the possibility that if the dynamics of protostellar outflows cause a decollimation effect then there could be an observational bias against the youngest of these objects due to the lower abundance of carbon monoxide. This would reproduce the apparent observational selection of wide angled "brandy snifter" shaped outflows in preference to narrower "champagne flute" shapes noted previously. However, at the smallest radii, the model yields negligible CO abundances at all ages, failing to reproduce the observations that continue all the way in to the central object. Since our injection function (Equation 5.2) scales as R_{S}^{-3}, the density of hot ions is greatest nearest the central source and as such, helium ions more readily dissociate any carbon monoxide present in this region. Potential resolutions to this issue are considered in Section 5.4.

The modelled HCO^+ abundance distributions exhibit similar properties to CO. Below the density of $10^5\,\mathrm{cm}^{-3}$ almost none is able to form, while fractional abundances of up to $\approx 3{\times}10^{-9}$ can form above this threshold. HCO^+ abundances are also marginally higher at higher temperatures due to the reduced dissociative recombination rate. Furthermore, HCO^+ is more prominent at larger radii due to the reduced electron injection rate leading to a slower rate of dissociative recombination. At densities of $3 \times 10^5\,\mathrm{cm}^{-3}$ and $10^6\,\mathrm{cm}^{-3}$ there is also evidence for a preferential distance (or equivalently polar angle) at which the abundance of HCO^+ peaks in equilibrium. Whether or not such a feature would be observable is unclear, but it does show the fine balance between needing enough H^+ from the injecta to form a large fractional abundance of HCO^+ in equilibrium and not adding too many electrons to destroy it all.

The features of the model discussed above suggest that the physical conditions leading

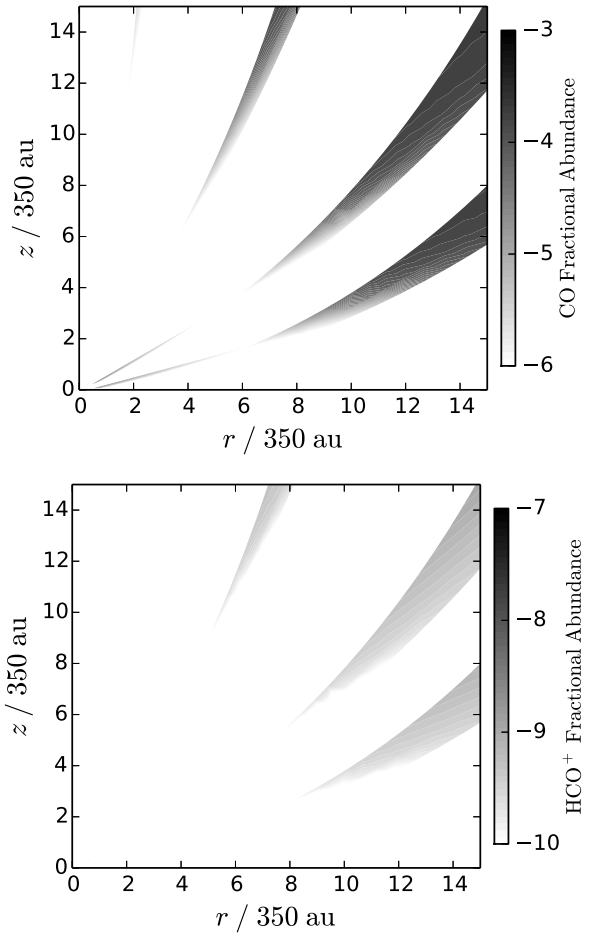

Figure 5.6. Cross–sectional fractional chemical abundance maps of CO (top) and HCO^+ (bottom) for cavities of dynamical age 2 (narrowest, not visible in either plot due to low abundances), 6, 11 and 14 (widest) kyrs. The interface density is $3 \times 10^5\,\mathrm{cm}^{-3}$ and the temperature is $T = 150\,\mathrm{K}$ giving the best qualitative fit to the observed chemical properties. The chemical integration times at which the abundances are plotted are 100, 300, 1000 and 1000 years for the 2, 6, 11 and 14 kyrs cavities respectively. The interface thickness has been exaggerated to 50% of the dynamical radius for clarity. Note that the youngest cavity is not seen in either plot due to the fractional abundances of both species being off the bottom of their respective scales.

to the qualitative observations of abundant carbon monoxide and HCO^+ are quite finely tuned. The combination of density $n = 3{\times}10^5\,\mathrm{cm}^{-3}$ and temperature $T = 150\,\mathrm{K}$ appears to give the best qualitative match. Such parameters are in broad agreement with the various observations, models and simulations discussed in Section 5.2. The abundance cross–sectional maps of CO and HCO^+ for that combination of parameters are shown in Figure 5.6. Most strikingly, the youngest 2 kyr cavity fails to reach CO saturation and would therefore likely be unobservable, in agreement with the hypothesis that the morphological bias towards wide–angled outflows is chemical in origin. While the equilibrium abundances of HCO^+ shown in the figure are somewhat lower than observed, it is still true that a short lived burst in fractional abundance is seen on time–scales of approximately 10 years as shown previously in Figure 5.3 and the same as in the previous photochemical model (Rawlings et al. 2000).

5.4 Discussion

We have produced a model for the chemistry in protostellar outflow interfaces driven by injected ions that is able to broadly reproduce some of the quantitative and qualitative features of those systems. For the right choice of density, bands of carbon monoxide are observed in the edges of cavities with opening angles larger than approximately 50°. For a cavity with a symmetry axis inclined in the plane of the sky, this would yield column densities of the order a few times $10^{18}\,\mathrm{cm}^{-2}$ which is easily observable. Furthermore, the critical density for CO(2–1) observations as taken from the Leiden Atomic and Molecular Database (Schöier et al. 2005) is approximately $10^5\,\mathrm{cm}^{-3}$. This value of critical density supports the view that interfaces with this density should be readily observable and, according to the model presented here, should show a wide–angled morphology. HCO^+ is also produced in our model, although not in quite the quantities inferred from observations. However, temporary enhancements to 10^{-8} occur on time–scales of 10–100 years as in Rawlings et al. (2000).

One key shortcoming is the inability to form carbon monoxide all the way into the central source. The cubic radial dependence $\propto R_\mathrm{S}^{-3}$ was the main factor in determining the local ion injection rate and not $\sin\psi$ as predicted. Due to this strong dependence on radius, CO is rapidly destroyed close to the opening angle due to too much He^+ being injected. There are, however, a number of potential refinements that may correct this deficiency.

The first point of note is that while the ejected material is presumed to be uniform with solid angle in our model, it could be argued that the presence of any magnetic field (for example) may collimate it, with more material ejected towards the poles, so that the injection rate at wide angles is reduced. This could take the form of a further geometric factor in the injection rate equation, e.g. $\propto \cos(\theta/\theta_{\rm m})$; maximal near the pole, zero at the opening angle and normalized over solid angle. Another consideration is that the cavity wall thickness may not be proportional to the dynamical radius. In their simulations, Cunningham et al. (2006) observe interfaces that are not only approximately constant thickness as a function of angle, but in fact grow with time. While the way in which this growth may be parametrized is not clear, it is evident that the thickness of the interface in our model could have been underestimated close to the central source. Increasing it would have the effect of diluting the injected helium ion abundance and hence reducing the rate of CO destruction. Finally, while reasonably argued from geometric considerations, the turbulent mixing efficiency $\sin\psi$ is surely uncertain. If it were a stronger function of the mixing angle (for example, $\sin^2\psi$), this would have the effect of again reducing the injection rate near the central source where the mixing angle $\psi \to 0$ and could yield an enhanced CO abundance.

The model features mentioned above are not considered explicitly in this exploratory paper, rather reserved for future work. However, variations in some of the parameters that were considered fixed in previous sections can shed some light on the response of the model that can be expected. The simplest of these is to allow a global multiplicative pre–factor to scale the injection rate in Equation 5.2. The results of varying this pre–factor are summarised in Figure 5.7. Reassuringly, on the time–scales we considered before, our initial conditions are in equilibrium in the absence of injection, while scaling the injection rate up by a factor of 10 removes both CO and HCO^+ from the system on all angular scales of interest. However, reducing the injection rate (as given by Equation 5.2) by a factor of ten demonstrates quite how sensitive the system is to the rate of turbulent injection. The interface then contains somewhat more HCO^+ on all angular scales and CO is observed in to much closer radii, although the discrimination between CO abundances in older and younger cavities is lost. All of the issues discussed above – interface thickness, mixing efficiency and collimation of the ejecta – could be realized in terms of a reduction in the injection efficiency, particularly close to the central source.

One other choice of parameter that was discussed thoroughly in Sections 5.1 and 5.2 is

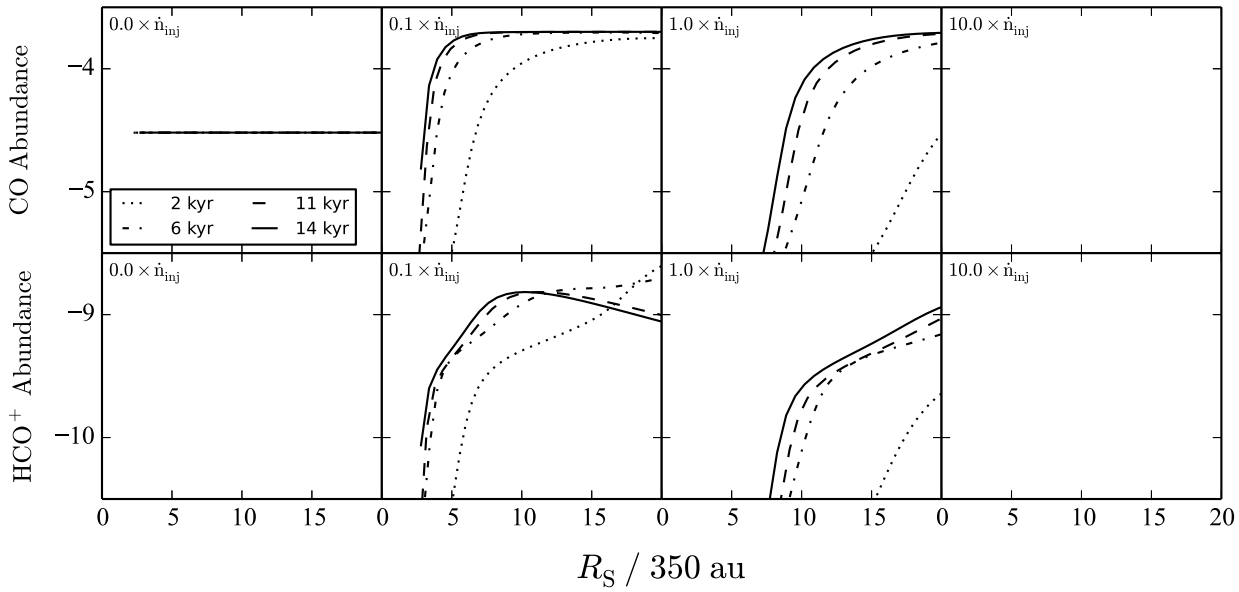

Figure 5.7. Fractional abundances of CO (top) and HCO^+ (bottom) relative to hydrogen for a cavity of dynamical age 2 (dot), 6 (dot–dash), 11 (dash) & 14 (solid) kyrs. The temperature is fixed at 150 K and the density at $10^5\,cm^{-3}$. A multiplicative prefactor to the injection rate of Equation 5.2 is used, ranging from 0 (left) through 0.1, 1 and 10 (right). The chemical integration times at which the abundances are plotted are 100, 300, 1000 and 1000 years for the 2, 6, 11 and 14 kyrs cavities respectively.

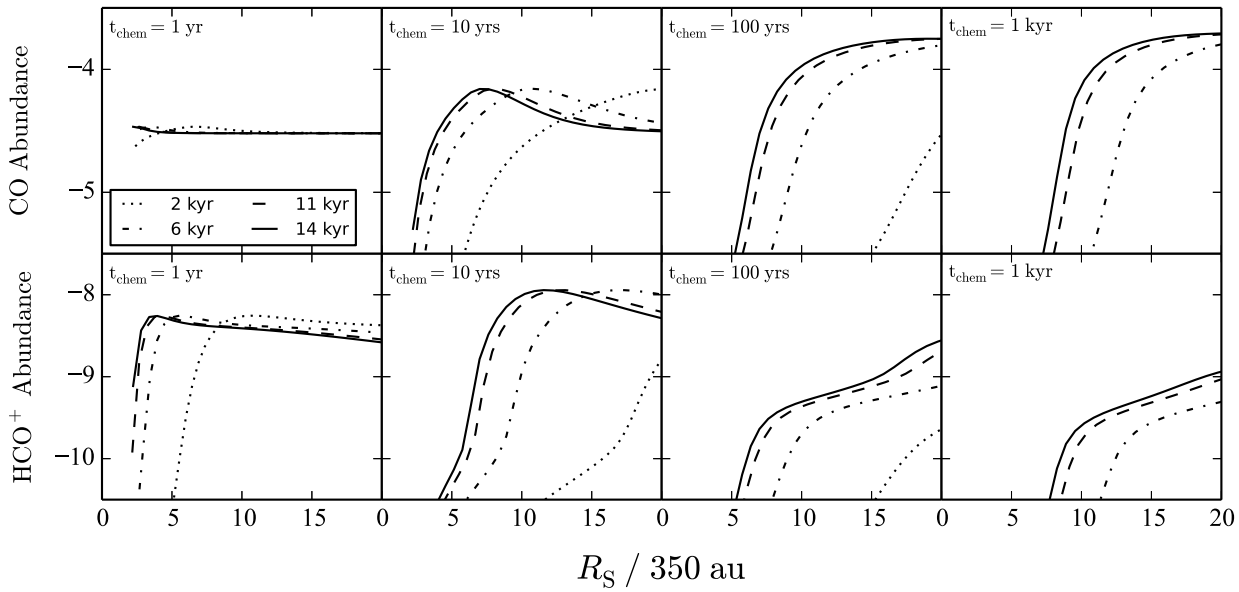

Figure 5.8. Fractional abundances of CO (top) and HCO^+ (bottom) relative to hydrogen for a cavity of age 2 (dot), 6 (dot–dash), 11 (dash) & 14 (solid) kyrs. The temperature is fixed at 150 K and the density at $10^5\,cm^{-3}$. The chemical integration time is varied from 1 year (left) through 10, 100 and 1000 years (right).

the chemical integration time, $t_{\rm chem}$. This choice was made to be shorter than the typical time–scale on which the injection rate changes for its age. However, as shown in Figure 5.2 the time–scales near the cavity base are significantly shorter than further out; this was not considered in Section 5.3. Figure 5.8 shows the evolution of the interface carbon monoxide and HCO^+ fractional abundances with chemical integration time. The relative stability of CO and the flare up of HCO^+ between 1 and 10 years of chemical integration is evident. However, between 100–1000 years, the profiles of both remain roughly constant, with only a small amount lost to the usual reaction channels. Figure 5.2 shows that the time–scale on which the injection rate changes is greater than 100 years for all angles and cavity ages. Therefore, we argue that the choice of chemical age between 100–1000 years from Section 5.3 has limited effect on the overall results. From Figure 5.3 we can see that the abundances of both CO and HCO^+ dip at $t \sim 6\,$kyrs. This *could* inhibit emission from some older outflow cavities, suggesting that, rather than just a minimum, there may be a range of cavity opening angles that are conducive to observable emission. However, the abundances recover at later times, so the effect may be marginal.

As a part of their work to model the abundance of CO^+ in outflow systems, Bruderer et al. (2009) considered a similar mixing mechanism driving the chemistry. While their model injected only ionized atomic hydrogen and electrons, they concluded it could only produce column densities $N(CO^+) \approx 10^8\,{\rm cm}^{-2}$; some two orders of magnitude less than their two dimensional photochemical model that matches the observational data. While our chemical scheme produces comparable peak abundances ($Y(CO^+) \approx 10^{-10}$), in some regions this abundance is sustained for as long as $10^4\,$yrs by comparison with less than one year in their model. We can therefore assume in our model that CO^+ is present through the full width of our cavity wall giving column densities for a cavity in the plane of the sky as high as $N(CO^+) \approx 10^{12}\,{\rm cm}^{-2}$. Although this maximal value may be larger than the inferred column density, it is possible that our full injection model can provide an adequate explanation of the observed CO^+ signal. The relatively high column densities as well as the strong correlation with HCO^+ over the chemical integration time make it an interesting species for comparison with other models. Another chemical signature of our model is the high abundance of OH seen in Figure 5.3. While previous studies have not provided a prediction on the presence of the OH radical, it would be expected in photochemical models following the photodissociation of water.

A number of potential chemical mechanisms were neglected in our modelling that could

have an effect on the numerical results of our study. Importantly, the presence of small dust grains and polycyclic aromatic hydrocarbons (PAHs) is believed to reduce the abundance of free electrons by combining with them rapidly to form negatively charged particles. This could significantly reduce the rate for dissociative recombination of HCO^+ by electrons; the dominant destruction pathway. While the negatively charged PAHs can themselves react to dissociate HCO^+, they move significantly slower than electrons and so the rate coefficient for that process will be significantly lower than with electrons, leading to an expected increase in the HCO^+ abundance in the presence of PAHs. This argument would suggest that the HCO^+ abundances in our model are therefore underestimates of the true column densities. On the other hand, negatively charged PAHs provide further chemical pathways for recombination, such as mutual neutralization reactions with C^+ ions (Bakes & Tielens 1998). This could reduce the formation rate of HCO^+ meaning that our modelled fractional abundances are actually overestimates. Since the model of Bruderer et al. (2009) also contained PAH chemistry, such reactions might also provide an explanation for the significant drop in CO^+ abundances that they found after one year. In the case of high–mass young stellar objects, Stäuber et al. (2005) demonstrated that the presence of X–ray photons from the central source acted to provide an enhancement in the HCO^+ abundance due to more C^+ ions produced by secondary photons from excited molecular hydrogen. Low–mass analogues are expected to have similar X–ray luminosities (Montmerle 2001) and so we expect that the presence of X–rays would complement our model by enhancing the HCO^+ abundance while leaving CO abundances roughly the same.

We also note that the origin of the cavity is not explicitly considered. Whilst we adopt the model of a decollimating outflow of constant total mass loss rate as in Cantó et al. (2008), the physics leading to such a system was never considered; the model was simply fit to the observed shape to constrain the physical parameters of the ejecta and envelope. In fact, the model of Li et al. (2013) is static by comparison, with the ram pressure of the outflow balanced by the turbulence in the interface. It is possible to apply our model to such static systems by considering the full chemical equilibrium as $t_{\rm chem} \rightarrow \infty$. Once again, Figure 5.8 suggests that this would have the effect of reducing the HCO^+ abundance but also removing most CO from the two younger cavities. This then fits even better with the observed bias, with only more powerful outflows that lead to wider opening angles being observable. Furthermore, our ignorance of the dynamics of the mixing in the cavity wall means we are unable to comment on the observed broad line profiles of CO (Yıldız et al.

2012) and HCO^+ (Kama et al. 2013) at high J. While a shock origin has been proposed, a more detailed dynamical model of the mixing in the cavity walls may tell us if a turbulent origin may also be valid.

Finally, it is noted that a two dimensional photochemical model could produce a similar morphological bias. However, the main source of UV radiation in the cavity is uncertain. In their model, Visser et al. (2012) considered the central protostellar object as a source of isotropic UV radiation, meaning the local radiation field varies along the cavity wall in a similar manner to our ion injection rate due to the same geometric spreading and impact angle. The resulting chemistry is therefore likely to have a qualitatively similar spatial dependence to our turbulent mixing model. The same is also likely to be true for UV photons originating from the accretion luminosity of the inner protostellar accretion disc (Spaans et al. 1995). However, Visser et al. (2012) also suggest local shocks in the cavity wall as it expands into the surrounding cloud could be a source of UV radiation that is more uniform along the cavity wall leading to a more uniform abundance profile. Furthermore, shocks from the internal working surfaces of the optical jet (Rawlings et al. 2000, van Kempen et al. 2009) could lead to the local radiation field and hence chemical abundances in the cavity wall being enhanced in the vicinity of knots in the jet. The spatial variations in the CO and HCO^+ abundances should be uniquely sensitive to the morphology of the system driving the chemistry, whether driven by photochemistry or turbulent mixing of ions.

5.5 Conclusions

We have developed a model for the chemistry at the interface between the outflow and the molecular envelope for low–mass protostars. The chemistry is driven by turbulent mixing of hot ions from the ejecta with cold molecular gas in the envelope. In particular we observe the following features of the model:

- A preference towards systems with wider opening angles in the modelled CO abundances.
- High abundances of HCO^+ without the need for photochemistry.
- Approximate chemical equilibrium reached on the time–scale at which the injection rate changes, leaving the results independent of the choice of chemical integration

times greater than 100 years.

- The results are acutely sensitive to the interface density, with $10^5\,\mathrm{cm}^{-3}$ providing the best fit to observed CO morphologies.

- Temperatures of 150 K give the maximum possible stable HCO^+ abundance within the range investigated.

- The assumed magnitude of the injection rate is critical in determining if CO and HCO^+ form or are destroyed by injected He^+ and e^- respectively.

These results are quantitatively similar to those of Rawlings et al. (2000) using photochemistry, with the addition that observed CO maps are qualitatively matched by the model abundances. This provides a possible explanation for the observational preference towards wide opening angle outflows: higher injection rates of dissociating ions for narrower opening angles reduce the CO abundance and make it unobservable. A three dimensional radiative transfer treatment would allow us to consider the degree to which a realistic abundance profile affects the inferred CO map, and to determine if such a treatment is necessary to understand the observations. It would also be interesting to consider photochemistry and ion driven chemistry in these systems side by side and in combination, to look for unique tracers that might reveal the true nature of the chemistry at work. We reserve these considerations for future work.

Chapter 6

Conclusions

In this thesis we have designed a range of new models, both chemical and statistical, and made application to a range of problems pertinent to astrochemical research. From the basic theory underlying chemical kinetics combined with recent laboratory and theoretical studies of rate coefficients that underpin all of this work, we have been able to address the physical processes driving chemistry in diffuse gas, photochemical shielding of the nitrogen chemical network by molecular hydrogen and atomic carbon and the origin of the carbon monoxide morphology and HCO^+ abundance in low–mass protostellar outflows. More importantly, in developing the software package Pliny for the statistical analysis of models against data with cutting edge parallel performance, we have provided a general framework to allow more precise and rigorous inferences to be made from astrochemical observations and scientific data in general. We summarise the key achievements and conclusions of each chapter in this thesis below while remarking upon possibilities for follow up work.

In Chapter 2 we demonstrated the performance of our new parallel implementation of the Bayesian nested sampling algorithm, Pliny. In terms of accuracy and reliability of the evidence integral, parallel time–to–solution and sampling efficiency we showed that it represents a significant improvement on the current state–of–the–art software for a range of unimodal posterior distributions. These improvements are down to utilizing a more accurate bounding ellipsoid algorithm to describe the sampled space and an MPI parallelization scheme that duplicates many inexpensive calculations across processors so as to minimize potential communication bottlenecks over low latency networks. The demonstration of

linear parallel scaling is a huge asset and we expect the software to find widespread use on current peta– and future exa–scale supercomputers and in the analysis of high–volume datasets coming from next generation telescopes including the *Square Kilometre Array* and *Euclid*. There are a range of potential performance gains still to be had, from moving to a hybrid MPI + OpenMP parallelism that exploits local shared memory for faster matrix methods, to implementing samplers that can deal with non–Gaussianity in a geometry free manner. Development is ongoing.

In Chapter 3 we made the first scientific application of PLINY to investigate the chemical complexity of diffuse gas clouds. From spectroscopic observations of carbon–, oxygen–, nitrogen– and sulphur–bearing ions and molecules in gas towards G10.6-0.4 made by the *Herschel Space Observatory*, we were able to constrain the physical conditions of the gas from a chemical model in a statistically robust and rigorous manner. As well as describing the physical state of the gas in terms of its density and temperature, we were able to discriminate between possible chemical mechanisms at work in regulating the abundances of those species. Photochemical processes which have been appealed to in previous models were ruled out as a viable source of ammonia with hydrogenation of nitrogen on grain surfaces inferred to be the more probable formation pathway. Furthermore, we confirm the need for a magnetically–induced decoupling of the ionic and neutral fluids; our inferences were consistent with turbulent dissipation and a single strong C–shock while ruling out the presence of many low velocity C–shocks. As identified in Section 3.4, future models with more complex geometries are likely needed in order to provide a more accurate description of the absorption line profiles and to discriminate between mechanisms for the origin of the magnetic decoupling.

In Chapter 4 we presented new calculations for the direct and cosmic–ray–induced photorates for a range of species. In addition we calculated mutual shielding factors for those species by the atomic carbon ionization continuum along with approximations for the mutual shielding by molecular hydrogen of molecular nitrogen and CN, both of which had been notably absent from many recent models of photon dominated regions (PDRs). We were able to model the importance of these mechanisms in promoting nitrogen chemistry, including significant enrichments in the relative abundances of ammonia and N_2H+, for a range of physical conditions. We expect our model for mutual shielding by atomic carbon to be integrated into many future PDR models due to their chemical significance yet low computational cost due to simple parametrized forms. Having demonstrated its

qualitative importance, there is clearly scope for more rigorous and quantitative work on the radiative transfer of molecular hydrogen to investigate the extent to which our assumption of complete mutual shielding of N_2 and CN is true, as well as for the dependence on dust extinction of our newly calculated photorates.

Finally, in Chapter 5 we modelled the origin of carbon monoxide and HCO^+ in the cavity walls of young, low–mass protostellar outflows as a chemical process driven by the turbulent mixing between hot ionized gas ejected from the central source and the cold molecular envelope. We were able to demonstrate quantitative agreement with the observed abundances of those species and with existing PDR models, as well as qualitatively argue that the origin of an observational bias towards outflows with wide opening angles could be due to the geometry–sensitive chemical mixing. Further work is needed to determine if turbulent mixing or photochemistry (or indeed a combination of the two) are the dominant chemical mechanism, which represents a suitable model comparison application for PLINY. Such an analysis would require robust and quantitative tracers to be calculated such as the CO excitation ladder and emission spectrum and column densities of CO^+ for comparison with observations.

Appendix A

Deriving the Interface Mixing Angle

As presented in Cantó et al. (2008), given an arbitrary scale length r_0, the dimensionless shape equation $\bar{R}_{\rm S}(\theta) = R_{\rm S}/r_0$ for a momentum driven wind of constant mass loss rate, steadily decollimating at a constant angular rate and expanding into an isothermal sphere, can be defined implicitly as:

$$\bar{\epsilon}^2\bar{\kappa}\bar{R}_{\rm S}^2 + 4(\bar{\epsilon}\bar{R}_{\rm S} + \theta - \theta_{\rm m})\cot(\theta/2) + 8\ln\left[\frac{\sin[(\theta_{\rm m} - \bar{\epsilon}\bar{R}_{\rm S})/2]}{\sin(\theta/2)}\right] = 0, \tag{A.1}$$

where $\bar{\epsilon}$ and $\bar{\kappa}$ are the dimensionless outflow decollimation rate and isothermal coefficient respectively and $\theta_{\rm m}$ is the cavity opening angle. This equation can be differentiated with respect to θ to give an implicit equation for $\bar{R}'_{\rm S} = {\rm d}\bar{R}_{\rm S}/{\rm d}\theta$:

$$\begin{aligned} 2\bar{\epsilon}^2\bar{\kappa}\bar{R}_{\rm S}\bar{R}'_{\rm S} + 4(\bar{\epsilon}\bar{R}'_{\rm S} + 1)\cot(\theta/2) - 2(\bar{\epsilon}\bar{R}_{\rm S} + \theta - \theta_{\rm m})\mathrm{cosec}^2(\theta/2) \\ - 4\left[\bar{\epsilon}\bar{R}'_{\rm S}\cot\left(\frac{\theta_{\rm m} - \bar{\epsilon}\bar{R}_{\rm S}}{2}\right) + \cot(\theta/2)\right] = 0. \end{aligned} \tag{A.2}$$

Converting back to dimensional quantities and considering the geometry of the problem in cylindrical coordinates (r,ϕ,z), we can calculate ω, the angle between the cavity wall

tangent and the cylindrical radial vector (r) as:

$$\mathrm{d}z = \mathrm{d}R_\mathrm{S}\cos\theta - R_\mathrm{S}\sin\theta\mathrm{d}\theta, \tag{A.3}$$

$$\mathrm{d}r = \mathrm{d}R_\mathrm{S}\sin\theta + R_\mathrm{S}\cos\theta\mathrm{d}\theta, \tag{A.4}$$

$$\tan\omega = \frac{\mathrm{d}z}{\mathrm{d}r} = \frac{R'_\mathrm{S} - R_\mathrm{S}\tan\theta}{R'_\mathrm{S}\tan\theta + R_\mathrm{S}}. \tag{A.5}$$

From triangle geometry we then get a function for the mixing angle, ψ, as a function of R_S, R'_S and θ:

$$\psi = \arctan\left(\frac{R'_\mathrm{S} - R_\mathrm{S}\tan\theta}{R'_\mathrm{S}\tan\theta + R_\mathrm{S}}\right) + \theta - \frac{\pi}{2}. \tag{A.6}$$

Appendix B

Deriving the Density Injection Rate

From the definitions given in Cantó et al. (2008), we define the density injection rate at time $t_{\rm dyn}$ into the interface, $\dot{\rho}_{\rm inj}$, as the mass loss rate per unit solid angle ($\dot{m}_{\Omega_0}$) at a time $R_{\rm S}/v_0$ in the past divided into a box of volume $R_{\rm S}^2 l$ where $l = 0.1R_{\rm S}$ is the assumed thickness of the cavity:

$$\dot{\rho}_{\rm inj} = \frac{\dot{m}_{\Omega_0}}{R_{\rm S}^2 l} \sin\psi. \tag{B.1}$$

Here we have added a nominal mixing efficiency $\sin\psi$ depending on the angle of impact between the outflowing material and the cavity wall, ψ, as derived in Appendix A. The definitions of these variables from Cantó et al. (2008) then give our equation for the rate of injected particle density, $\dot{n}_{\rm inj}$, assuming it is dominated by H^+ ions of mass $m_{\rm H}$:

$$\dot{n}_{\rm inj} = \frac{5\dot{M}_0}{\pi m_{\rm H} R_{\rm S}(\theta)^3} \frac{\sin\psi}{1 - \cos[\epsilon(t_{\rm dyn} - \frac{R_{\rm S}(\theta)}{v_0})]} \,{\rm cm}^{-3}\,{\rm s}^{-1}, \tag{B.2}$$

where $\dot{M}_0$ is the total mass loss rate of the outflow, $m_{\rm H}$ is the hydrogen atomic mass, ϵ is the outflow decollimation rate and v_0 is the outflow velocity.

Bibliography

Altwegg, K., Balsiger, H., Bar-Nun, A., Berthelier, J. J., Bieler, A., Bochsler, P., Briois, C., Calmonte, U., Combi, M., De Keyser, J., Eberhardt, P., Fiethe, B., Fuselier, S., Gasc, S., Gombosi, T. I., Hansen, K. C., Hässig, M., Jäckel, A., Kopp, E., Korth, A., LeRoy, L., Mall, U., Marty, B., Mousis, O., Neefs, E., Owen, T., Rème, H., Rubin, M., Sémon, T., Tzou, C.-Y., Waite, H. & Wurz, P., 2015, *Science*, **347**(27), A387 1

André, P., Men'shchikov, A., Bontemps, S., Könyves, V., Motte, F., Schneider, N., Didelon, P., Minier, V., Saraceno, P., Ward-Thompson, D., di Francesco, J., White, G., Molinari, S., Testi, L., Abergel, A., Griffin, M., Henning, T., Royer, P., Merín, B., Vavrek, R., Attard, M., Arzoumanian, D., Wilson, C. D., Ade, P., Aussel, H., Baluteau, J.-P., Benedettini, M., Bernard, J.-P., Blommaert, J. A. D. L., Cambrésy, L., Cox, P., di Giorgio, A., Hargrave, P., Hennemann, M., Huang, M., Kirk, J., Krause, O., Launhardt, R., Leeks, S., Le Pennec, J., Li, J. Z., Martin, P. G., Maury, A., Olofsson, G., Omont, A., Peretto, N., Pezzuto, S., Prusti, T., Roussel, H., Russeil, D., Sauvage, M., Sibthorpe, B., Sicilia-Aguilar, A., Spinoglio, L., Waelkens, C., Woodcraft, A. & Zavagno, A., 2010, *A&A*, **518**, L102 1.1

Arce, H. G., Mardones, D., Corder, S. A., Garay, G., Noriega-Crespo, A. & Raga, A. C., 2013, *ApJ*, **774**, 39 (document), 1.2, 5

Asplund, M., Grevesse, N. & Sauval, A. J., 2005, in T. G. Barnes, III & F. N. Bash (eds.), *Cosmic Abundances as Records of Stellar Evolution and Nucleosynthesis*, volume 336 of *Astronomical Society of the Pacific Conference Series*, p. 25 4.2

Bachiller, R. & Pérez Gutiérrez, M., 1997, *ApJL*, **487**, L93 5.2

Bakes, E. L. O. & Tielens, A. G. G. M., 1994, *ApJ*, **427**, 822 1.3.3, 3.2.3

Bakes, E. L. O. & Tielens, A. G. G. M., 1998, *ApJ*, **499**, 258 5.4

Barsuhn, J., 1977, *A&AS*, **28**, 453 (document), 4.1

Bate, M. R., Tricco, T. S. & Price, D. J., 2014, *MNRAS*, **437**, 77 5.1, 5.2

Bayet, E., Viti, S., Williams, D. A., Rawlings, J. M. C. & Bell, T., 2009, *ApJ*, **696**, 1466 4.4

Bisbas, T. G., Bell, T. A., Viti, S., Yates, J. & Barlow, M. J., 2012, *MNRAS*, **427**, 2100 1.3.3

Black, J. H., 1987, in D. J. Hollenbach & H. A. Thronson, Jr. (eds.), *Interstellar Processes*, volume 134 of *Astrophysics and Space Science Library*, pp. 731–744 1.3.3

Black, J. H. & Dalgarno, A., 1977, *ApJSS*, **34**, 405 (document), 4.1

Boisse, P., 1990, *A&A*, **228**, 483 1

Bolatto, A. D., Wolfire, M. & Leroy, A. K., 2013, *Annual Review of Astronomy and Astrophysics*, **51**, 207 1.2.3

Branscomb, L. M., 1962, *in Atomic and Molecular Processes*, **ed. D. R. Bates, Academic Press Inc, New York**, p.100 (document), 4.1

Bruderer, S., Benz, A. O., Doty, S. D., van Dishoeck, E. F. & Bourke, T. L., 2009, *ApJ*, **700**, 872 5.4

Buch, V. & Zhang, Q., 1991, *ApJ*, **379**, 647 1.3.2

Bürzle, F., Clark, P. C., Stasyszyn, F., Dolag, K. & Klessen, R. S., 2011, *MNRAS*, **417**, L61 5.1, 5.2

Cantó, J., Raga, A. C. & Williams, D. A., 2008, *RMxAA*, **44**, 293 (document), 5.1, 5.1, 5.2, 5.4, A, B, B

Carter, V. L., 1972, **56**, 4195 (document), 4.1.1, 4.1.2, 4.2

Caselli, P., Keto, E., Pagani, L., Aikawa, Y., Yıldız, U. A., van der Tak, F. F. S., Tafalla, M., Bergin, E. A., Nisini, B., Codella, C., van Dishoeck, E. F., Bachiller, R., Baudry, A., Benedettini, M., Benz, A. O., Bjerkeli, P., Blake, G. A., Bontemps, S., Braine, J., Bruderer, S., Cernicharo, J., Daniel, F., di Giorgio, A. M., Dominik, C., Doty, S. D., Encrenaz, P., Fich, M., Fuente, A., Gaier, T., Giannini, T., Goicoechea, J. R.,

de Graauw, T., Helmich, F., Herczeg, G. J., Herpin, F., Hogerheijde, M. R., Jackson, B., Jacq, T., Javadi, H., Johnstone, D., Jørgensen, J. K., Kester, D., Kristensen, L. E., Laauwen, W., Larsson, B., Lis, D., Liseau, R., Luinge, W., Marseille, M., McCoey, C., Megej, A., Melnick, G., Neufeld, D., Olberg, M., Parise, B., Pearson, J. C., Plume, R., Risacher, C., Santiago-García, J., Saraceno, P., Shipman, R., Siegel, P., van Kempen, T. A., Visser, R., Wampfler, S. F. & Wyrowski, F., 2010, *A&A*, **521**, L29 1.3.2

Cecchi-Pestellini, C., Rawlings, J. M. C., Viti, S. & Williams, D. A., 2010, *ApJ*, **725**, 1581 1.2.1

Černý, V., 1985, *Journal of Optimization Theory and Applications*, **45**, 41 1.4, 2, 3

Chieze, J.-P., Pineau des Forêts, G. & Flower, D. R., 1998, *MNRAS*, **295**, 672 3, 3.2.2, 3.4

Codella, C., Viti, S., Williams, D. A. & Bachiller, R., 2006, *ApJL*, **644**, L41 5

Collings, M. P., Anderson, M. A., Chen, R., Dever, J. W., Viti, S., Williams, D. A. & McCoustra, M. R. S., 2004, *MNRAS*, **354**, 1133 1.3.2

Corbel, S. & Eikenberry, S. S., 2004, *A&A*, **419**, 191 3.2, 3.2.2

Cunningham, A. J., Frank, A., Quillen, A. C. & Blackman, E. G., 2006, *ApJ*, **653**, 416 5.1, 5.2, 5.4

Draine, B. T., 1978, *ApJS*, **36**, 595 1.3.2, 1.3.2, 3.2.2, 3.4, 4.1.1

Draine, B. T., 2011, *Physics of the Interstellar and Intergalactic Medium* (document), 1.3

Duane, S., Kennedy, A., Pendleton, B. & Roweth, D., 1985, *Physics Letters B*, **195**, 216 2

Efstathiou, A., Rowan-Robinson, M. & Siebenmorgen, R., 2000, *MNRAS*, **313**, 734 1

Faure, A., Hily-Blant, P., Le Gal, R., Rist, C. & Pineau des Forêts, G., 2013, *ApJ*, **770**, L2 3, 3.2.1, 3.5

Fayolle, E. C., Öberg, K. I., Cuppen, H. M., Visser, R. & Linnartz, H., 2011, *A&A*, **529**, A74 1.3.2

Federman, S. R., Rawlings, J. M. C., Taylor, S. D. & Williams, D. A., 1996, *MNRAS*, **279**, L41 3, 3

Feroz, F. & Hobson, M. P., 2008, *MNRAS*, **384**, 449 2, 2.1

Feroz, F., Hobson, M. P. & Bridges, M., 2009, *MNRAS*, **398**, 1601 2, 2.1, 3.1

Feroz, F., Hobson, M. P., Cameron, E. & Pettitt, A. N., 2013, *ArXiv e-prints* 2, 2.4

Feroz, F., Hobson, M. P. & Trotta, R., 2010, *ArXiv e-prints* (document), 2.1, 2.2.1

Ferrière, K. M., 2001, *Reviews of Modern Physics*, **73**, 1031 (document), 1.1

Field, G. B., Goldsmith, D. W. & Habing, H. J., 1969, *ApJL*, **155**, L149 1.1

Flower, D. R. & Pineau des Forêts, G., 1995, *MNRAS*, **275**, 1049 3.2.3

Flower, D. R. & Pineau des Forêts, G., 1998, *MNRAS*, **297**, 1182 3, 3.2.2, 3.4

Galli, D. & Palla, F., 1998, *A&A*, **335**, 403 3

Garrod, R. T., 2013, *ApJ*, **778**, 158 1.3.2

Garrod, R. T., Wakelam, V. & Herbst, E., 2007, *A&A*, **467**, 1103 1.3.2

Gerin, M., de Luca, M., Black, J., Goicoechea, J. R., Herbst, E., Neufeld, D. A., Falgarone, E., Godard, B., Pearson, J. C., Lis, D. C., Phillips, T. G., Bell, T. A., Sonnentrucker, P., Boulanger, F., Cernicharo, J., Coutens, A., Dartois, E., Encrenaz, P., Giesen, T., Goldsmith, P. F., Gupta, H., Gry, C., Hennebelle, P., Hily-Blant, P., Joblin, C., Kazmierczak, M., Kolos, R., Krelowski, J., Martin-Pintado, J., Monje, R., Mookerjea, B., Perault, M., Persson, C., Plume, R., Rimmer, P. B., Salez, M., Schmidt, M., Stutzki, J., Teyssier, D., Vastel, C., Yu, S., Contursi, A., Menten, K., Geballe, T., Schlemmer, S., Shipman, R., Tielens, A. G. G. M., Philipp-May, S., Cros, A., Zmuidzinas, J., Samoska, L. A., Klein, K. & Lorenzani, A., 2010, *A&A*, **518**, L110 3, 3.2.2

Gerin, M. & PRISMAS Consortium, 2010, in *38th COSPAR Scientific Assembly*, volume 38 of *COSPAR Meeting*, p. 2477 3, 3.2.1

German, S. & German, S., 1984, *IEEE Transactions on Pattern Analysis and Machine Intelligence*, **6**, 721 2

Glover, S. C. O. & Abel, T., 2008, *MNRAS*, **388**, 1627 3

Glover, S. C. O. & Jappsen, A.-K., 2007, *ApJ*, **666**, 1 3.2.3

Godard, B., Falgarone, E., Gerin, M., Hily-Blant, P. & de Luca, M., 2010, *A&A*, **520**, A20 3

Godard, B., Falgarone, E., Gerin, M., Lis, D. C., De Luca, M., Black, J. H., Goicoechea, J. R., Cernicharo, J., Neufeld, D. A., Menten, K. M. & Emprechtinger, M., 2012, *A&A*, **540**, A87 (document), 3.2, 3.2.1, 3.1, 3.2.1, 3.2.2, 3.4

Godard, B., Falgarone, E. & Pineau Des Forêts, G., 2009, *A&A*, **495**, 847 1.3.3, 3.2.2, 3.2.3, 3.2.3, 3.3, 3.4

Godard, B., Falgarone, E. & Pineau des Forêts, G., 2014, *A&A*, **570**, A27 (document), 1, 3, 3.2.2, 3.4, 3.2.3, 3.4

Goldsmith, P. F. & Langer, W. D., 1978, *ApJ*, **222**, 881 1

Graff, P., Feroz, F., Hobson, M. P. & Lasenby, A., 2012, *MNRAS*, **421**, 169 2

Gredel, R., Lepp, S., Dalgarno, A. & Herbst, E., 1989, *ApJ*, **347**, 289 1.3.2, 4.1.4, 4.1.4

Gredel, R., Pineau des Forêts, G. & Federman, S. R., 2002, *A&A*, **389**, 993 3.2.2, 3.4

Greene, T., 2001, *American Scientist*, **89**, 316 (document), 1.2

Gueth, F., Guilloteau, S. & Bachiller, R., 1996, *A&A*, **307**, 891 5

Habing, H. J., 1968, *Bull. Astr. Inst. Netherlands*, **19**, 421 1.3.2

Han, Y. & Han, Z., 2012, *ApJ*, **749**, 123 2

Hara, C., Shimajiri, Y., Tsukagoshi, T., Kurono, Y., Saigo, K., Nakamura, F., Saito, M., Wilner, D. & Kawabe, R., 2013, *ApJ*, **771**, 128 5

Hartogh, P., Lis, D. C., Bockelée-Morvan, D., de Val-Borro, M., Biver, N., Küppers, M., Emprechtinger, M., Bergin, E. A., Crovisier, J., Rengel, M., Moreno, R., Szutowicz, S. & Blake, G. A., 2011, *Nature*, **478**, 218 1

Hartquist, T. W., Rawlings, J. M. C., Williams, D. A. & Dalgarno, A., 1993, *Quarterly Journal of the Royal Astronomical Society*, **34**, 213 1, 3

Hastings, W., 1970, *Biometrika*, **57**, 97 1.4, 2, 3

Haud, U. & Kalberla, P. M. W., 2007, *A&A*, **466**, 555 3.4

Hawken, A. J., Abdalla, F. B., Hütsi, G. & Lahav, O., 2012, *MNRAS*, **424**, 2 3

Herbst, E., Lee, H. H., Howe, D. A. & Millar, T. J., 1994, *MNRAS*, **268**, 335 1.2.1

Herczeg, G. J., Karska, A., Bruderer, S., Kristensen, L. E., van Dishoeck, E. F., Jørgensen, J. K., Visser, R., Wampfler, S. F., Bergin, E. A., Yıldız, U. A., Pontoppidan, K. M. & Gracia-Carpio, J., 2012, *A&A*, **540**, A84 5

Hily-Blant, P., Falgarone, E. & Pety, J., 2008, *A&A*, **481**, 367 3.4

Hindmarsh, A. C. & Petzold, L. R., 1995, *Computers in Physics*, **9(2)**, 148 1.3.1, 3.2.3

Hogerheijde, M. R., van Dishoeck, E. F., Blake, G. A. & van Langevelde, H. J., 1998, *ApJ*, **502**, 315 5

Hollenbach, D., Chu, S.-I. & McCray, R., 1976, *ApJ*, **208**, 458 1.1, 3

Hollenbach, D., Kaufman, M. J., Neufeld, D., Wolfire, M. & Goicoechea, J. R., 2012, *ApJ*, **754**, 105 (document), 3, 3.2.2, 3.4, 3

Hollenbach, D. & McKee, C. F., 1979, *ApJS*, **41**, 555 3

Hollenbach, D. J. & Tielens, A. G. G. M., 1999, *Reviews of Modern Physics*, **71**, 173 4

Indriolo, N., Geballe, T. R., Oka, T. & McCall, B. J., 2007, *ApJ*, **671**, 1736 1.3.2

Indriolo, N. & McCall, B. J., 2012, *ApJ*, **745**, 91 1.3.2

Indriolo, N., Neufeld, D. A., Gerin, M., Schilke, P., Benz, A. O., Winkel, B., Menten, K. M., Chambers, E. T., Black, J. H., Bruderer, S., Falgarone, E., Godard, B., Goicoechea, J. R., Gupta, H., Lis, D. C., Ossenkopf, V., Persson, C. M., Sonnentrucker, P., van der Tak, F. F. S., van Dishoeck, E. F., Wolfire, M. G. & Wyrowski, F., 2015, *ApJ*, **800**, 40 (document), 3, 3.2, 3.1, 3.2.1, 3.2.2, 3.4

Jeffreys, H., 1961, *The Theory Of Probability, 3rd Edition*, **Oxford Publishing**, 432 2.1, 3.3

Jiménez-Serra, I., Martín-Pintado, J., Viti, S., Martín, S., Rodríguez-Franco, A., Faure, A. & Tennyson, J., 2006, *ApJL*, **650**, L135 5.2

John, F., 1948, *Studies and Essays Presented to R.Courant on his 60th Birthday. January 8, 1948*, **Wiley Interscience, New York**, 187 2.1.1

Jørgensen, J. K., Favre, C., Bisschop, S. E., Bourke, T. L., van Dishoeck, E. F. & Schmalzl, M., 2012, *ApJL*, **757**, L4 1

Kama, M., López-Sepulcre, A., Dominik, C., Ceccarelli, C., Fuente, A., Caux, E., Higgins, R., Tielens, A. G. G. M. & Alonso-Albi, T., 2013, *A&A*, **556**, A57 5.2, 5.4

Kamenetzky, J., Glenn, J., Maloney, P. R., Aguirre, J. E., Bock, J. J., Bradford, C. M., Earle, L., Inami, H., Matsuhara, H., Murphy, E. J., Naylor, B. J., Nguyen, H. T. & Zmuidzinas, J., 2011, *ApJ*, **731**, 83 4.4

Kamp, I. & Bertoldi, F., 2000, *A&A*, **353**, 276 1.3.2

Kass, R. & Raftery, A., 1995, *Journal of the American Statistical Association*, **90**, 773 2.1, 3.3

Kennicutt, Jr., R. C., 1998, *ApJ*, **498**, 541 1.2.2

Keto, E. & Caselli, P., 2008, *ApJ*, **683**, 238 1.3.2

Keto, E., Rybicki, G. B., Bergin, E. A. & Plume, R., 2004, *ApJ*, **613**, 355 1.4, 3

Khachiyan, L., 1996, *Mathematics of Operations Research*, **21**, 307 2.1.1

Kirkpatrick, S., Gelatt Jr, C. D. & Vecchi, M. P., 1983, *Science*, **220**, 671 1.4, 2, 3

Knuth, D., 1969, *The Art of Computer Programming, vol. 2: Seminumerical Algorithms* 2.1.1

Koenig, X. P., Allen, L. E., Gutermuth, R. A., Hora, J. L., Brunt, C. M. & Muzerolle, J., 2008, *ApJ*, **688**, 1142 1.1, 3

Kristensen, L. E., Visser, R., van Dishoeck, E. F., Yıldız, U. A., Doty, S. D., Herczeg, G. J., Liu, F.-C., Parise, B., Jørgensen, J. K., van Kempen, T. A., Brinch, C., Wampfler, S. F., Bruderer, S., Benz, A. O., Hogerheijde, M. R., Deul, E., Bachiller, R., Baudry, A., Benedettini, M., Bergin, E. A., Bjerkeli, P., Blake, G. A., Bontemps, S., Braine, J., Caselli, P., Cernicharo, J., Codella, C., Daniel, F., de Graauw, T., di Giorgio, A. M., Dominik, C., Encrenaz, P., Fich, M., Fuente, A., Giannini, T., Goicoechea, J. R., Helmich, F., Herpin, F., Jacq, T., Johnstone, D., Kaufman, M. J., Larsson, B., Lis, D., Liseau, R., Marseille, M., McCoey, C., Melnick, G., Neufeld, D., Nisini, B., Olberg, M., Pearson, J. C., Plume, R., Risacher, C., Santiago-García, J., Saraceno, P., Shipman, R., Tafalla, M., Tielens, A. G. G. M., van der Tak, F., Wyrowski, F., Beintema, D., de Jonge, A., Dieleman, P., Ossenkopf, V., Roelfsema, P., Stutzki, J. & Whyborn, N., 2010, *A&A*, **521**, L30 5

Kulander, K. C. & Bottcher, C., 1978, **29**, 141 (document), 4.2

Lada, C. J., 1987, in M. Peimbert & J. Jugaku (eds.), *Star Forming Regions*, volume 115 of *IAU Symposium*, pp. 1–17 1.1

Langer, W. D., Velusamy, T. & Xie, T., 1996, *ApJL*, **468**, L41 5, 5.1

Le Bourlot, J., Pineau des Forêts, G., Roueff, E. & Flower, D. R., 1995, *A&A*, **302**, 870 1.3.1, 3.4

Le Bourlot, J., Pineau des Forêts, G., Roueff, E. & Schilke, P., 1993, *ApJL*, **416**, L87 1.3.1, 3.4

Le Gal, R., Hily-Blant, P., Faure, A., Pineau des Forêts, G., Rist, C. & Maret, S., 2014, *A&A*, **562**, A83 3.5

Le Petit, F., Nehmé, C., Le Bourlot, J. & Roueff, E., 2006, *ApJS*, **164**, 506 1.3.3

Lee, H.-H., Herbst, E., Pineau des Forêts, G., Roueff, E. & Le Bourlot, J., 1996, *A&A*, **311**, 690 1.3.2

Lee, H.-H., Roueff, E., Pineau des Forêts, G., Shalabiea, O. M., Terzieva, R. & Herbst, E., 1998, *A&A*, **334**, 1047 1.3.1, 3.4

Lesaffre, P., Pineau des Forêts, G., Godard, B., Guillard, P., Boulanger, F. & Falgarone, E., 2013, *A&A*, **550**, A106 1, 3

Li, G.-X., Qiu, K., Wyrowski, F. & Menten, K., 2013, *A&A*, **559**, A23 5, 5.1, 5.2, 5.4

Lochner, M., Natarajan, I., Zwart, J. T. L., Smirnov, O., Bassett, B. A., Oozeer, N. & Kunz, M., 2015, *ArXiv e-prints* 2

Majumdar, L., Das, A., Chakrabarti, S. K. & Chakrabarti, S., 2012, *Research in Astronomy and Astrophysics*, **12**, 1613 1

Makrymallis, A. & Viti, S., 2014, *ApJ*, **794**, 45 1.4, 3

Martin, J., Ringeval, C., Trotta, R. & Vennin, V., 2014, *JCAP*, **3**, 39 2

Mathis, J. S., Rumpl, W. & Nordsieck, K. H., 1977, *ApJ*, **217**, 425 1.2.4

McElroy, D., Walsh, C., Markwick, A. J., Cordiner, M. A., Smith, K. & Millar, T. J., 2013, *A&A*, **550**, A36 1.3.2, 3, 3.2.3

McKee, C. F. & Ostriker, J. P., 1977, *ApJ*, **218**, 148 1.1

Metropolis, N., Rosenbluth, A., Rosenbluth, M., Teller, A. & Teller, E., 1953, *Journal of Chemical Physics*, **21**, 1087 1.4, 2, 3

Molinari, S., Swinyard, B., Bally, J., Barlow, M., Bernard, J.-P., Martin, P., Moore, T., Noriega-Crespo, A., Plume, R., Testi, L., Zavagno, A., Abergel, A., Ali, B., Anderson, L., André, P., Baluteau, J.-P., Battersby, C., Beltrán, M. T., Benedettini, M., Billot, N., Blommaert, J., Bontemps, S., Boulanger, F., Brand, J., Brunt, C., Burton, M., Calzoletti, L., Carey, S., Caselli, P., Cesaroni, R., Cernicharo, J., Chakrabarti, S., Chrysostomou, A., Cohen, M., Compiegne, M., de Bernardis, P., de Gasperis, G., di Giorgio, A. M., Elia, D., Faustini, F., Flagey, N., Fukui, Y., Fuller, G. A., Ganga, K., Garcia-Lario, P., Glenn, J., Goldsmith, P. F., Griffin, M., Hoare, M., Huang, M., Ikhenaode, D., Joblin, C., Joncas, G., Juvela, M., Kirk, J. M., Lagache, G., Li, J. Z., Lim, T. L., Lord, S. D., Marengo, M., Marshall, D. J., Masi, S., Massi, F., Matsuura, M., Minier, V., Miville-Deschênes, M.-A., Montier, L. A., Morgan, L., Motte, F., Mottram, J. C., Müller, T. G., Natoli, P., Neves, J., Olmi, L., Paladini, R., Paradis, D., Parsons, H., Peretto, N., Pestalozzi, M., Pezzuto, S., Piacentini, F., Piazzo, L., Polychroni, D., Pomarès, M., Popescu, C. C., Reach, W. T., Ristorcelli, I., Robitaille, J.-F., Robitaille, T., Rodón, J. A., Roy, A., Royer, P., Russeil, D., Saraceno, P., Sauvage, M., Schilke, P., Schisano, E., Schneider, N., Schuller, F., Schulz, B., Sibthorpe, B., Smith, H. A., Smith, M. D., Spinoglio, L., Stamatellos, D., Strafella, F., Stringfellow, G. S., Sturm, E., Taylor, R., Thompson, M. A., Traficante, A., Tuffs, R. J., Umana, G., Valenziano, L., Vavrek, R., Veneziani, M., Viti, S., Waelkens, C., Ward-Thompson, D., White, G., Wilcock, L. A., Wyrowski, F., Yorke, H. W. & Zhang, Q., 2010, *A&A*, **518**, L100 1.1

Monje, R. R., Lis, D. C., Roueff, E., Gerin, M., De Luca, M., Neufeld, D. A., Godard, B. & Phillips, T. G., 2013, *ApJ*, **767**, 81 3.5

Montmerle, T., 2001, in T. Montmerle & P. André (eds.), *From Darkness to Light: Origin and Evolution of Young Stellar Clusters*, volume 243 of *Astronomical Society of the Pacific Conference Series*, p. 731 5.4

Morbidelli, A., Chambers, J., Lunine, J. I., Petit, J. M., Robert, F., Valsecchi, G. B. & Cyr, K. E., 2000, *Meteoritics and Planetary Science*, **35**, 1309 1

Mukherjee, P., Parkinson, D. & Liddle, A. R., 2006, *ApJL*, **638**, L51 2, 3.1

Neufeld, D. A., Goicoechea, J. R., Sonnentrucker, P., Black, J. H., Pearson, J., Yu, S., Phillips, T. G., Lis, D. C., de Luca, M., Herbst, E., Rimmer, P., Gerin, M., Bell, T. A., Boulanger, F., Cernicharo, J., Coutens, A., Dartois, E., Kazmierczak, M., Encrenaz, P., Falgarone, E., Geballe, T. R., Giesen, T., Godard, B., Goldsmith, P. F., Gry, C., Gupta, H., Hennebelle, P., Hily-Blant, P., Joblin, C., Kołos, R., Krełowski, J., Martín-Pintado, J., Menten, K. M., Monje, R., Mookerjea, B., Perault, M., Persson, C., Plume, R., Salez, M., Schlemmer, S., Schmidt, M., Stutzki, J., Teyssier, D., Vastel, C., Cros, A., Klein, K., Lorenzani, A., Philipp, S., Samoska, L. A., Shipman, R., Tielens, A. G. G. M., Szczerba, R. & Zmuidzinas, J., 2010*a*, *A&A*, **521**, L10 3, 3.2.2

Neufeld, D. A., Roueff, E., Snell, R. L., Lis, D., Benz, A. O., Bruderer, S., Black, J. H., De Luca, M., Gerin, M., Goldsmith, P. F., Gupta, H., Indriolo, N., Le Bourlot, J., Le Petit, F., Larsson, B., Melnick, G. J., Menten, K. M., Monje, R., Nagy, Z., Phillips, T. G., Sandqvist, A., Sonnentrucker, P., van der Tak, F. & Wolfire, M. G., 2012, *ApJ*, **748**, 37 3.5

Neufeld, D. A., Sonnentrucker, P., Phillips, T. G., Lis, D. C., de Luca, M., Goicoechea, J. R., Black, J. H., Gerin, M., Bell, T., Boulanger, F., Cernicharo, J., Coutens, A., Dartois, E., Kazmierczak, M., Encrenaz, P., Falgarone, E., Geballe, T. R., Giesen, T., Godard, B., Goldsmith, P. F., Gry, C., Gupta, H., Hennebelle, P., Herbst, E., Hily-Blant, P., Joblin, C., Kołos, R., Krełowski, J., Martín-Pintado, J., Menten, K. M., Monje, R., Mookerjea, B., Pearson, J., Perault, M., Persson, C., Plume, R., Salez, M., Schlemmer, S., Schmidt, M., Stutzki, J., Teyssier, D., Vastel, C., Yu, S., Cais, P., Caux, E., Liseau, R., Morris, P. & Planesas, P., 2010*b*, *A&A*, **518**, L108 3, 3.2.2, 3.4, 3.5

Nisini, B., Benedettini, M., Codella, C., Giannini, T., Liseau, R., Neufeld, D., Tafalla, M., van Dishoeck, E. F., Bachiller, R., Baudry, A., Benz, A. O., Bergin, E., Bjerkeli, P., Blake, G., Bontemps, S., Braine, J., Bruderer, S., Caselli, P., Cernicharo, J., Daniel, F., Encrenaz, P., di Giorgio, A. M., Dominik, C., Doty, S., Fich, M., Fuente, A., Goicoechea, J. R., de Graauw, T., Helmich, F., Herczeg, G., Herpin, F., Hogerheijde, M., Jacq, T., Johnstone, D., Jørgensen, J., Kaufman, M., Kristensen, L., Larsson, B., Lis, D., Marseille, M., McCoey, C., Melnick, G., Olberg, M., Parise, B., Pearson, J., Plume, R., Risacher, C., Santiago, J., Saraceno, P., Shipman, R., van Kempen, T. A., Visser, R., Viti, S., Wampfler, S., Wyrowski, F., van der Tak, F., Yıldız, U. A., Delforge, B.,

Desbat, J., Hatch, W. A., Péron, I., Schieder, R., Stern, J. A., Teyssier, D. & Whyborn, N., 2010, *A&A*, **518**, L120 5, 5.2

Noble, J. A., Congiu, E., Dulieu, F. & Fraser, H. J., 2012, *MNRAS*, **421**, 768 1.3.2

Occhiogrosso, A., Viti, S., Ward, M. D. & Price, S. D., 2012, *MNRAS*, **427**, 2450 1.3.2

Offner, S. S. R., Bisbas, T. G., Bell, T. A. & Viti, S., 2014, *MNRAS*, **440**, L81 1.3.3

Offner, S. S. R., Klein, R. I., McKee, C. F. & Krumholz, M. R., 2009, *ApJ*, **703**, 131 1

Ó Ruanaidh, J. & Fitzgerald, W., 1996, *Numerical Bayesian Methods Applied to Signal Processing.*, **Springer–Verlag: New York** 2

Pan, L. & Padoan, P., 2009, *ApJ*, **692**, 594 3.2.2

Persson, C. M., Black, J. H., Cernicharo, J., Goicoechea, J. R., Hassel, G. E., Herbst, E., Gerin, M., de Luca, M., Bell, T. A., Coutens, A., Falgarone, E., Goldsmith, P. F., Gupta, H., Kaźmierczak, M., Lis, D. C., Mookerjea, B., Neufeld, D. A., Pearson, J., Phillips, T. G., Sonnentrucker, P., Stutzki, J., Vastel, C., Yu, S., Boulanger, F., Dartois, E., Encrenaz, P., Geballe, T. R., Giesen, T., Godard, B., Gry, C., Hennebelle, P., Hily-Blant, P., Joblin, C., Kołos, R., Krełowski, J., Martín-Pintado, J., Menten, K., Monje, R., Perault, M., Plume, R., Salez, M., Schlemmer, S., Schmidt, M., Teyssier, D., Péron, I., Cais, P., Gaufre, P., Cros, A., Ravera, L., Morris, P., Lord, S. & Planesas, P., 2010, *A&A*, **521**, L45 3, 3.2.2, 3.4, 4.4

Persson, C. M., De Luca, M., Mookerjea, B., Olofsson, A. O. H., Black, J. H., Gerin, M., Herbst, E., Bell, T. A., Coutens, A., Godard, B., Goicoechea, J. R., Hassel, G. E., Hily-Blant, P., Menten, K. M., Müller, H. S. P., Pearson, J. C. & Yu, S., 2012, *A&A*, **543**, A145 (document), 3, 3.2.1, 3.1, 3.2.1, 3.2.2, 3.4

Persson, C. M., Hajigholi, M., Hassel, G. E., Olofsson, A. O. H., Black, J. H., Herbst, E., Müller, H. S. P., Cernicharo, J., Wirström, E. S., Olberg, M., Hjalmarson, Å., Lis, D. C., Cuppen, H. M., Gerin, M. & Menten, K. M., 2014, *A&A*, **567**, A130 3, 3.2, 3.2.1

Pety, J. & Falgarone, É., 2000, *A&A*, **356**, 279 3.2.2

Pilbratt, G. L., Riedinger, J. R., Passvogel, T., Crone, G., Doyle, D., Gageur, U., Heras, A. M., Jewell, C., Metcalfe, L., Ott, S. & Schmidt, M., 2010, *A&A*, **518**, L1 1

Pineau Des Forêts, G. & Roueff, E., 2000, in *Astronomy, physics and chemistry of H^+_3*, volume 358 of *Royal Society of London Philosophical Transactions Series A*, p. 2549 1.3.1, 3.4

Placek, B., Knuth, K. H. & Angerhausen, D., 2014, *ApJ*, **795**, 112 2

Prasad, S. S. & Tarafdar, S. P., 1983, *ApJ*, **267**, 603 1.2.1, 4.1.4

Price, R. J., Viti, S. & Williams, D. A., 2003, *MNRAS*, **343**, 1257 3

Qiu, K., Zhang, Q., Wu, J. & Chen, H.-R., 2009, *ApJ*, **696**, 66 5

Rawlings, J. M. C., 1988, *MNRAS*, **232**, 507 4

Rawlings, J. M. C., Drew, J. E. & Barlow, M. J., 1993, *MNRAS*, **265**, 968 (document), 4.1.1, 4.1

Rawlings, J. M. C., Hartquist, T. W., Menten, K. M. & Williams, D. A., 1992, *MNRAS*, **255**, 471 3

Rawlings, J. M. C., Redman, M. P. & Carolan, P. B., 2013*a*, *MNRAS*, **435**, 289 5

Rawlings, J. M. C., Redman, M. P., Keto, E. & Williams, D. A., 2004, *MNRAS*, **351**, 1054 5

Rawlings, J. M. C., Taylor, S. D. & Williams, D. A., 2000, *MNRAS*, **313**, 461 5, 5.2, 5.2, 5.2, 5.3, 5.4, 5.4, 5.5

Rawlings, J. M. C., Williams, D. A., Viti, S., Cecchi-Pestellini, C. & Duley, W. W., 2013*b*, *MNRAS*, **430**, 264 1.2.1

Richings, A. J., Schaye, J. & Oppenheimer, B. D., 2014*a*, *MNRAS*, **440**, 3349 1, 1.3.3, 1.3.3, 3.2.3

Richings, A. J., Schaye, J. & Oppenheimer, B. D., 2014*b*, *MNRAS*, **442**, 2780 1.3.3

Roberts, H., van der Tak, F. F. S., Fuller, G. A., Plume, R. & Bayet, E., 2011, *A&A*, **525**, A107 4.4

Roberts, J. F., Rawlings, J. M. C., Viti, S. & Williams, D. A., 2007, *MNRAS*, **382**, 733 1.3.2, 3

Röllig, M., Abel, N. P., Bell, T., Bensch, F., Black, J., Ferland, G. J., Jonkheid, B., Kamp, I., Kaufman, M. J., Le Bourlot, J., Le Petit, F., Meijerink, R., Morata, O., Ossenkopf, V., Roueff, E., Shaw, G., Spaans, M., Sternberg, A., Stutzki, J., Thi, W.-F., van Dishoeck, E. F., van Hoof, P. A. M., Viti, S. & Wolfire, M. G., 2007, *A&A*, **467**, 187 3, 3, 4, 4.2, 4.2, 4.3

Rollins, R. P. & Rawlings, J. M. C., 2012, *MNRAS*, **427**, 2328 1.3.2, 3, 4

Rollins, R. P., Rawlings, J. M. C., Williams, D. A. & Redman, M. P., 2014, *MNRAS*, **443**, 3033 5

Ruiz de Austri, R., Trotta, R. & Feroz, F., 2011, *SuperBayeS: Supersymmetry Parameters Extraction Routines for Bayesian Statistics*, Astrophysics Source Code Library 2

Rybicki, G. B. & Hummer, D. G., 1991, *A&A*, **245**, 171 1.3.3

Savage, B. D., Cardelli, J. A. & Sofia, U. J., 1992, *ApJ*, **401**, 706 4.2

Savage, B. D. & Mathis, J. S., 1979, *Annual Review of Astronomy and Astrophysics*, **17**, 73 1.2.4

Schilke, P., Neufeld, D. A., Müller, H. S. P., Comito, C., Bergin, E. A., Lis, D. C., Gerin, M., Black, J. H., Wolfire, M., Indriolo, N., Pearson, J. C., Menten, K. M., Winkel, B., Sánchez-Monge, Á., Möller, T., Godard, B. & Falgarone, E., 2014, *A&A*, **566**, A29 3, 3.5

Schmidt, M., 1959, *ApJ*, **129**, 243 1.2.2

Schöier, F. L., van der Tak, F. F. S., van Dishoeck, E. F. & Black, J. H., 2005, *A&A*, **432**, 369 5.4

Seifried, D., Banerjee, R., Pudritz, R. E. & Klessen, R. S., 2015, *MNRAS*, **446**, 2776 1.1

Shaw, J. R., Bridges, M. & Hobson, M. P., 2007, *MNRAS*, **378**, 1365 2.1.1, 2.1.1, 2.2.1, 2.4

Shu, F. H., 1977, *ApJ*, **214**, 488 1.1, 5

Skilling, J., 2004, in R. Fischer, R. Preuss & U. von Toussaint (eds.), *Bayesian Inference and maximum entropy methods in science and engineering*, volume 735 of *Amer. Inst. Phys. Conference Proc.*, p. 395 2.1, 2.1, 2.3, 3.1, 3.3

Smith, I. W. M., Herbst, E. & Chang, Q., 2004, *MNRAS*, **350**, 323 1.2.1

Sobolev, V. V., 1960, *Moving envelopes of stars* 1.3.3, 3.2.3

Spaans, M., Hogerheijde, M. R., Mundy, L. G. & van Dishoeck, E. F., 1995, *ApJL*, **455**, L167 5.4

Spitzer, L., 1978, *Physical processes in the interstellar medium* 1.3.2

Stäuber, P., Doty, S. D., van Dishoeck, E. F. & Benz, A. O., 2005, *A&A*, **440**, 949 5.4

Stecher, T. P. & Williams, D. A., 1967, *ApJL*, **149**, L29 4.1.2

Sternberg, A., 2004, in S. Pfalzner, C. Kramer, C. Staubmeier & A. Heithausen (eds.), *The Dense Interstellar Medium in Galaxies*, p. 423 4

Sternberg, A. & Dalgarno, A., 1995, *ApJS*, **99**, 565 (document), 1.4, 3.4

Sternberg, A., Dalgarno, A. & Lepp, S., 1987, *ApJ*, **320**, 676 1.3.2, 4.1.4

Sun, P. & Freund, R. M., 2004, *Operations Research*, **52**, 690 2.1.1, 1, 3.1

Surjanovic, S. & Bingham, D., 2015, *Virtual Library of Simulation Experiments: Test Functions and Datasets*, `http://www.sfu.ca/~ssurjano`, Retrieved July 15 2015 (document), 2.1, 2.2.1

Tan, J. C., Kong, S., Butler, M. J., Caselli, P. & Fontani, F., 2013, *ApJ*, **779**, 96 1.1

Taquet, V., Ceccarelli, C. & Kahane, C., 2012, *A&A*, **538**, A42 1.3.2

Tielens, A. G. G. M. & Hollenbach, D., 1985, *ApJ*, **291**, 722 1.3.3, 3.2.3, 4, 4, 4.1.3

Tielens, A. G. G. M., Meixner, M. M., van der Werf, P. P., Bregman, J., Tauber, J. A., Stutzki, J. & Rank, D., 1993, *Science*, **262**, 86 (document), 1.3

Truelove, J. K., Klein, R. I., McKee, C. F., Holliman, II, J. H., Howell, L. H., Greenough, J. A. & Woods, D. T., 1998, *ApJ*, **495**, 821 1.3.3

Umebayashi, T. & Nakano, T., 1980, *Publications of the Astronomical Society of Japan*, **32**, 405 1.3.2

van der Tak, F. F. S., Black, J. H., Schöier, F. L., Jansen, D. J. & van Dishoeck, E. F., 2007, *A&A*, **468**, 627 5.2

van Dishoeck, E. F., 1988, in T. J. Millar & D. A. Williams (eds.), *Rate Coefficients in Astrochemistry*, volume 146 of *Astrophysics and Space Science Library*, pp. 49–72 (document), 4, 4.1.1, 4.1.1, 4.1, 4.1.2, 4.2, 4.1.3

van Dishoeck, E. F. & Black, J. H., 1982, *ApJ*, **258**, 533 4.1.1

van Dishoeck, E. F. & Black, J. H., 1988, *ApJ*, **334**, 771 1.3.2, 4

van Kempen, T. A., Kristensen, L. E., Herczeg, G. J., Visser, R., van Dishoeck, E. F., Wampfler, S. F., Bruderer, S., Benz, A. O., Doty, S. D., Brinch, C., Hogerheijde, M. R., Jørgensen, J. K., Tafalla, M., Neufeld, D., Bachiller, R., Baudry, A., Benedettini, M., Bergin, E. A., Bjerkeli, P., Blake, G. A., Bontemps, S., Braine, J., Caselli, P., Cernicharo, J., Codella, C., Daniel, F., di Giorgio, A. M., Dominik, C., Encrenaz, P., Fich, M., Fuente, A., Giannini, T., Goicoechea, J. R., de Graauw, T., Helmich, F., Herpin, F., Jacq, T., Johnstone, D., Kaufman, M. J., Larsson, B., Lis, D., Liseau, R., Marseille, M., McCoey, C., Melnick, G., Nisini, B., Olberg, M., Parise, B., Pearson, J. C., Plume, R., Risacher, C., Santiago-García, J., Saraceno, P., Shipman, R., van der Tak, F., Wyrowski, F., Yıldız, U. A., Ciechanowicz, M., Dubbeldam, L., Glenz, S., Huisman, R., Lin, R. H., Morris, P., Murphy, J. A. & Trappe, N., 2010, *A&A*, **518**, L121 5

van Kempen, T. A., van Dishoeck, E. F., Güsten, R., Kristensen, L. E., Schilke, P., Hogerheijde, M. R., Boland, W., Nefs, B., Menten, K. M., Baryshev, A. & Wyrowski, F., 2009, *A&A*, **501**, 633 5.2, 5.4

Vigeland, S. J. & Vallisneri, M., 2014, *MNRAS*, **440**, 1446 2

Visser, R., Kristensen, L. E., Bruderer, S., van Dishoeck, E. F., Herczeg, G. J., Brinch, C., Doty, S. D., Harsono, D. & Wolfire, M. G., 2012, *A&A*, **537**, A55 5, 5.4

Visser, R., van Dishoeck, E. F. & Black, J. H., 2009, *A&A*, **503**, 323 1.3.2, 4, 4

Viti, S., Collings, M. P., Dever, J. W., McCoustra, M. R. S. & Williams, D. A., 2004, *MNRAS*, **354**, 1141 1.3.2

Viti, S., Jimenez-Serra, I., Yates, J. A., Codella, C., Vasta, M., Caselli, P., Lefloch, B. & Ceccarelli, C., 2011, *ApJL*, **740**, L3 1.3.2

Viti, S. & Williams, D. A., 1999, *MNRAS*, **305**, 755 1.3.2

Wakelam, V., Herbst, E., Loison, J.-C., Smith, I. W. M., Chandrasekaran, V., Pavone, B., Adams, N. G., Bacchus-Montabonel, M.-C., Bergeat, A., Béroff, K., Bierbaum, V. M., Chabot, M., Dalgarno, A., van Dishoeck, E. F., Faure, A., Geppert, W. D., Gerlich, D., Galli, D., Hébrard, E., Hersant, F., Hickson, K. M., Honvault, P., Klippenstein, S. J., Le Picard, S., Nyman, G., Pernot, P., Schlemmer, S., Selsis, F., Sims, I. R., Talbi, D., Tennyson, J., Troe, J., Wester, R. & Wiesenfeld, L., 2012, *ApJS*, **199**, 21 1.3.2

Wakelam, V., Loison, J.-C., Herbst, E., Pavone, B., Bergeat, A., Béroff, K., Chabot, M., Faure, A., Galli, D., Geppert, W. D., Gerlich, D., Gratier, P., Harada, N., Hickson, K. M., Honvault, P., Klippenstein, S. J., Le Picard, S. D., Nyman, G., Ruaud, M., Schlemmer, S., Sims, I. R., Talbi, D., Tennyson, J. & Wester, R., 2015, *ApJS*, **217**, 20 1.3.2

Waldmann, I. P., Tinetti, G., Rocchetto, M., Barton, E. J., Yurchenko, S. N. & Tennyson, J., 2015, *ApJ*, **802**, 107 3

Whittet, D. C. B., Gerakines, P. A., Hough, J. H. & Shenoy, S. S., 2001, *ApJ*, **547**, 872 1

Williams, D. A. & Viti, S., 2014, *Observational Molecular Astronomy* (document), 1.1

Wolniewicz, L., Simbotin, I. & Dalgarno, A., 1998, *ApJS*, **115**, 293 3

Woodall, J., Agúndez, M., Markwick-Kemper, A. J. & Millar, T. J., 2007, *A&A*, **466**, 1197 1.3.2, 4.1.1, 4.1.4, 4.2, 5.2

Woods, P. M., Kelly, G., Viti, S., Slater, B., Brown, W. A., Puletti, F., Burke, D. J. & Raza, Z., 2012, *ApJ*, **750**, 19 1.3.2, 3.5

Wootten, A. & Thompson, A. R., 2009, *IEEE Proceedings*, **97**, 1463 1

Yıldız, U. A., Kristensen, L. E., van Dishoeck, E. F., Belloche, A., van Kempen, T. A., Hogerheijde, M. R., Güsten, R. & van der Marel, N., 2012, *A&A*, **542**, A86 5.4

Yıldız, U. A., van Dishoeck, E. F., Kristensen, L. E., Visser, R., Jørgensen, J. K., Herczeg, G. J., van Kempen, T. A., Hogerheijde, M. R., Doty, S. D., Benz, A. O., Bruderer, S., Wampfler, S. F., Deul, E., Bachiller, R., Baudry, A., Benedettini, M., Bergin, E., Bjerkeli, P., Blake, G. A., Bontemps, S., Braine, J., Caselli, P., Cernicharo, J., Codella, C., Daniel, F., di Giorgio, A. M., Dominik, C., Encrenaz, P., Fich, M., Fuente, A., Giannini, T., Goicoechea, J. R., de Graauw, T., Helmich, F., Herpin, F., Jacq, T.,

Johnstone, D., Larsson, B., Lis, D., Liseau, R., Liu, F.-C., Marseille, M., McCoey, C., Melnick, G., Neufeld, D., Nisini, B., Olberg, M., Parise, B., Pearson, J. C., Plume, R., Risacher, C., Santiago-García, J., Saraceno, P., Shipman, R., Tafalla, M., Tielens, A. G. G. M., van der Tak, F., Wyrowski, F., Dieleman, P., Jellema, W., Ossenkopf, V., Schieder, R. & Stutzki, J., 2010, *A&A*, **521**, L40 5, 5.2

Yusef-Zadeh, F., Muno, M., Wardle, M. & Lis, D. C., 2007, *ApJ*, **656**, 847 1.3.2

. . . yesterday is all that does count. What else is there?

C. McCarthy