Advancements in Galactic Chemical Evolution

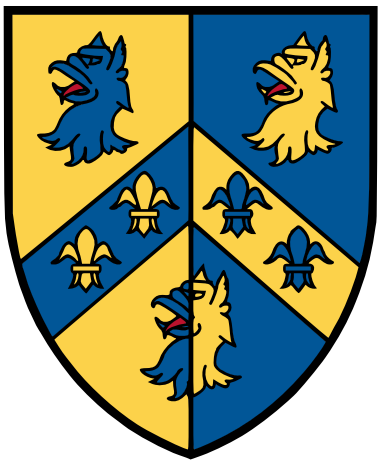

Jack Fraser-Govil

Trinity College

University of Oxford

A thesis submitted for the degree of

Doctor of Philosophy

Hilary 2022

Acknowledgements

The greatest of thanks must go to my supervisor, Dr. Ralph Schönrich. For the past five years, through both my Master's and my DPhil, I have had the privilege of working with one of the most incisive minds I have ever, or will ever, know. He has been an invaluable mentor, ally and my vociferous proponent through these tumultuous years. His support and dedication underpins every word of this thesis.

I would also like to give thanks to Dr. John Magorrian. Whilst he would probably argue that he has been a supervisor in name only, the support he has provided extends far beyond the academic. He has been a perpetual guiding hand through the bureaucracy of the system, and an infinite source of (bad) humour when I needed it most. I also thank the Royal Society for providing the funding for me to undertake my research, and who made this DPhil financially possible.

Over the years, I have had the pleasure of working with a great number of people who have enriched and improved my DPhil experience . To name but a few: James Binney, Rimpei Chiba, Dominic Dootson, Andy Everall, and Chengdong Li have all been invaluable colleagues, always there to lend a helping hand or provide incisive questioning. Amery Gration has been a true friend, always willing share a coffee, or lend me the wisdom of his considerable years. Douglas Boubert took a gamble on my abilities, which has resulted in a fruitful, lengthy, and rewarding collaboration. Jennifer Friske has been a sister-in-arms in recent months, it is a travesty it took us this long to actually work together.

I am forever in the debt of my wonderful family: my mother, Claire, my father David and my brother, Sam. From them, I learned everything that I am. Although I am now all grown up, I know that there will always be a place for me at home, with a warm cup of tea and a welcoming hug.

This thesis also holds dear the memories of my Nan and my Nanna, Hazel Quinn and Cindy Fraser, both of whom we lost over the course of my DPhil. Though they are no longer with us, memories of trifle, mug-squash, and grandmotherly love have kept me going strong through long, dark nights of the soul.

Finally, I thank Shubhangi, the love of my life, and my soon-to-be wife. She is my entire world, and this work has only been possible due to her infinite grace, patience and support for this endeavour.

Abstract

In this thesis we examine several improvements which can be made to the canonical understanding of Galactic Chemical Evolution (GCE) models in light of modern developments in stellar physics and new modelling techniques. Not only is GCE a vital component of our understanding of the history and composition of the universe in its own right, but since it incorporates physics on the scale of the cosmological to the subatomic, it may also be used as an independent probe and constraint on physics across many different fields.

We develop two brand new, state-of-the-art models for studying galactic chemical evolution under two differing paradigms: the Simple Analytical Chemical Evolution Model (SACEM) which allows for rapid and complete scans of the high-dimensional parameter space associated with GCE, and the updated version of the Radial Migration in Chemical Evolution Simulation (RAMICES II), which emulates the underlying physics chemical evolution as faithfully, and completely, as possible.

These models allow us to investigate the origin of r-process nucleosynthesis within the Milky Way, in particular, identifying the ability of 'collapsar' events to contribute meaningfully to the global r-process budget. Utilising the SACEM code, we are able to evaluate in excess of 10^{12} potential models of the galaxy, and evaluate them against intentionally weak constrains on the resulting chemical and temporal properties, and hence provide a strong upper bound on the ability of such events to contribute to the r-process abundance of the Milky Way at late times.

We study the impact of advances in stellar evolution modelling on GCE models, utilising new rotationally-resolved stellar evolution codes and nucleosynthetic yield grids. We show that the output from these new stellar models is fundamentally irreconcilable with observed data, and hence develop a number of tools which allow future stellar modelling to be rapidly diagnosed as incompatible, providing a new constraint on the internal evolution of stellar bodies.

Finally, we apply our new RAMICES II model to briefly study the phenomena of metallicity-offset, strongly peaked curves in the gas-phase abundance diagrams – so called 'caustics' – and how these imprint themselves in non-trivial ways on the stellar abundance plane, using them to inform future studies of both r-process nucleosynthesis and the assembly history of the galaxy.

Declaration of Authorship

The work in this thesis is based on research I carried out at the University of Oxford between Michaelmas Term 2018 and Hilary Term 2022.

Jack Fraser-Govil
22nd March 2022

Peer-Reviewed Publications

1. Fraser, J. & Schönrich, S., 2022a, Metallicity-suppressed collapsars cannot be the dominant r-process source in the Milky Way. **MNRAS 509**, 6008-6027

Chapters 3 and 5 are adapted from this work. I wrote the SACEM code and performed the simulations and analysis presented in this paper, and wrote the majority of the text. Dr. Ralph Schönrich provided the RAMICES I code and invaluable insight, guidance and direction in the writing of the paper.

Firstly, note that my name has changed since the publication of this work (from Fraser to Fraser-Govil), and secondly that in the light of the development of the RAMICES II code, some the simulations in Chapter 5 have been updated from the published version, in order to use this new code.

Works Under Development

1. Fraser-Govil, J., Friske, J. & Schönrich, S., 2022, The RAMICES II Code: Next Generation Galactic Chemical Evolution. **Unpublished Work**

2. Fraser-Govil, J. & Schönrich, S., 2022b, Constraining Rotational Stellar Yields with Galactic Chemical Evolution. **Under Review**

Chapters 4 and 7 (for Fraser-Govil, Friske & Schönrich 2022), as well as Chapter 6 (for Fraser-Govil & Schönrich 2022b) are adapted from the content of these works which are have either been submitted, or are due to be submitted for publication in the near future. In both of these works, I wrote the codes and performed the simulations and analysis contained within. My coauthors provided theoretical contributions, computational calibration and style and directional advice for these papers.

Contents

List of Figures

List of Tables

List of Abbreviations, Definitions & Symbols

$\alpha\beta\gamma$ . . . . . Humorous shorthand for the paper Alpher, Bethe & Gamow (1948), which first proposed the mechanism for primordial nucleosynthesis

A09 . . . . . Asplund et al. (2009). The literature standard for solar abundances

B^{2}FH . . . . . Burbidge et al. (1957). The work which provided a solid grounding for the mechanisms of stellar nucleosynthesis

CCSN . . . . . Core-Collapse Supernova, a general term for SNII and SNIb events, indicating that they are associated with the death of massive stars

CGM . . . . . Circum-Galactic Medium, the extended halo of gas which comprises the nearest part of the IGM to the Milky Way.

DTD . . . . . Delay-Time Distribution, the probability of a star or compact remnant undergoing a supernova event at a given time

ECSN . . . . . Electron-Capture Supernova, a subset of CCSN in which the collapse is triggered by electron-capture processes in a O-Ne-Mg core, instead of an iron core.

IGM . . . . . Inter-Galactic Medium, the diffuse gas which fills the space between galaxies, with a density of approximately 1 atoms cubic metre. Note that we use IGM interchangably with CGM throughout this work, as the portion of the IGM of interest to our work is that closest to our galaxy.

ISM . . . . . Inter-Stellar Medium, the diffuse gas which fills the space between stars, with a density of approximately 10^6 atoms per cubic metre.

IMF . . . . . Initial Mass Function, a function, $\zeta(M)$, which gives the fraction of stars born in the mass range $[M, M + \mathrm{d}M]$

GCE . . . . . Galactic Chemical Evolution, the study of the evolution of the chemical abundance of galaxies

LGRB . . . . . Long Gamma Ray Burst, a $\mathcal{O}(100\mathrm{s})$ burst of gamma rays, thought to originate from exotic events such as collapsars.

LC18 . . . . . Limongi & Chieffi (2018). Published rotational yield grids which are the focus of Chapter 6

Metallicity . . . . . Denoted Z, the metallicity is the fraction of matter which is neither hydrogen nor helium

NSM . . . . . Neutron Star Merger

$M_\odot$ . . . . . A solar mass, equal to 1.988435×10^{30}kg.

M22 . . . . . Magg et al. (2022). Updated values to the solar abundances, challenging A09

RA13 . . . . . Ramírez-Agudelo et al. (2013). A work studying the rotational properties of stars

SB09 . . . . . Schönrich & Binney (2009a), a foundational work for this thesis, introduced the original RAMICES code.

SFR . . . . . Star Formation Rate, the mass of new stars which formed at a given time.

Si19 . . . . . Siegel et al. (2019), a work arguing for the dominance of Collapsar models in r-process synthesis

SGRB . . . . . Short Gamma Ray Burst, a $\mathcal{O}(1\mathrm{s})$ burst of gamma rays, thought to originate from NSM.

Sn08 . . . . . Sneden et al. (2008), a work demonstrating various properties of r-process rich stars

SNIa, SN1a . . . . Supernova 1a, an observational class of supernovae events, theorised to arise from the thermonuclear disruption of a white dwarf.

Tinsley Diagram . Plots of the form [X/Fe] vs [Fe/H], following the pioneering work of Tinsley (1979)

ZAMS . . . . . Zero-Age Main Sequence, the point at which a proto-star enters the main sequence and becomes a 'proper' star. Stellar Ages are given relative to this point.

In the beginning the Universe was created. This has made a lot of people very angry and been widely regarded as a bad move.

— Douglas Adams

1 Introduction

Contents

In the beginning, there was heat, matter, and energy.

Our prevailing cosmology shows us that, at the very beginning of time, the universe was nothing more than a hot, dense soup of primordial matter. This fluid was undergoing rapid expansion, a 'Big Bang', as matter, energy and gravity fought to reshape the fabric of space and time.

At this earliest stage, the universe was almost unrecognisable from the form we know today: the energy densities were so extreme that stable nuclei - even the component hadrons themselves - were ripped apart before they could form, and the electrically charged melange was so dense that photons could not propagate. The universe was a homogeneous, uniform fluid which stretched across the cosmos. One must only look around to see that there have been more than a few changes since that time.

We explain these changes through the fact that, as it expanded, the universe cooled such that the 'soup' began to thicken and condense. After only a handful of microseconds, the change in temperature was enough for the subatomic particles to become 'confined', and stable protons and neutrons first appeared. These protons and neutrons rapidly exchanged forms via neutrino and electron captures, until the temperature decreased further to disfavour such interactions, freezing the proton-neutron ratio to about 6:1.

Over the next few minutes, the cosmological temperature caused high energy interactions between the nucleons, forcing them together in fusion events. Since neutrons are unstable unless bound into a nucleus, all of the neutrons either found themselves fused into nuclei, or decayed into protons. From the 6:1 proton to neutron ratio, Alpher, Bethe & Gamow (1948) (referred to as $\alpha\beta\gamma$) showed that the resulting universe must be approximately 75% Hydrogen (free protons) by mass, 25% Helium-4 (2 protons & 2 neutrons), with trace amounts of deuterium, lithium and other light elements.

As the universe cooled further, the fusion events would cease, electrons could recombine with the nuclei to form neutral atoms, and the universe would be filled with this primordial gas mixture of hydrogen, helium and minuscule amounts of light elements.

Despite this prediction, I am writing this thesis as a carbon-based lifeform, breathing in an atmosphere of nitrogen and oxygen, and typing on silicon-based computer technology. The building I sit in has a framework of iron-alloys, surrounded by calcium-based concrete, whilst the ring I proposed to my fiancée with is made of gold and platinum.

We know, however, that these materials are not simple aggregates of hydrogen and helium – they are fundamentally different and distinct on a nuclear level. There is a veritable cornucopia of elemental abundance and diversity in the world around us, which simply cannot be explained by the mechanisms of primordial nucleosynthesis proposed by the $\alpha\beta\gamma$ paper.

How, then, did the universe come to be populated by elements other than the primordial few?

1.1 Into the Hearts of Stars

The narrative of the primordial synthesis above ended with the formation of the neutral atoms of hydrogen, helium and some trace pollutants[1], approximately 300,000 years after the Big Bang.

Up until this point, the ordinary matter and the radiation had been tightly coupled: the electrically charged soup had been pushed and pulled around by the photons, preventing the matter from clumping together or collapsing under the force of gravity. The dark matter, however, had no such limitation as it is (as far as we know), totally decoupled from all forces aside from gravity. Whilst our primordial soup had been pushed around, the dark matter had been undergoing *hierarchical structure formation*: small overdensities collapsed in on themselves to cause regions of high density, which in turn led to further collapse. By the time

[1]For historical reasons, astronomers have chosen to term these pollutants 'metals', and the measure of their abundance is called the 'metallicity'. This is in spite of the fact that the most abundant 'metals' in the universe (Oxygen, Carbon, Nitrogen) are not metals in the chemical sense. We appreciate this is nonsensical, and yet we persist.

of the so-called 'Cosmic Dark Ages', these clumps had become 'halos', enormous clumps of dark matter, which created gravitational wells, into which the primordial gases flowed.

Over the next 100 million years, the gas continued to stream down into the wells, and began to cool and become more dense. Much like the formation of the halos, small pockets of gas eventually became dense enough to undergo runaway Jeans collapse and become the first stars in the universe. With the ignition of the first stars came not only light (and hence a re-ionization of a significant portion of the gas), but also the creation of the first new nuclei since those first few minutes after the Big Bang.

That stars source their energy by the nuclear 'burning'[2] of hydrogen into helium had been first suggested in the 1920s by Arthur Eddington, whilst Gamow famously elaborated on this mechanism with his quantum mechanical treatment of the tunnelling processes necessary for fusion at reasonable energies. Hans Bethe used these arguments to predict two important fusion chains: the proton-proton chain (which transmutes hydrogen to helium) and the carbon-nitrogen-oxygen (CNO) cycle.

However, these arguments were all concerned with the energy production of these mechanisms - the elements themselves were merely a means to an end. How the elements got there, and what happened to them after they donated their energy to star, was left unsaid. It was not until Hoyle (1946) was it suggested that these processes were in fact the *source* of the elements within our universe. Hoyle (1954) elaborated on this further, however the crown jewel which cemented the theory of stellar nucleosynthesis as the dominant paradigm was the paper of Burbidge, Burbidge, Fowler & Hoyle (1957). The B^2FH paper, as it is known, comprehensively analysed the theoretical arguments of Hoyle, along with the nuclear physics of Fowler and the observational data of the Burbidges, and presented the following picture as the source of the chemical abundances in the universe.

As a star shines, it converts one element into another in a nuclear furnace in the core of the star, starting with hydrogen, the most abundant element. As the star ages and depletes their hydrogen fuel, they can (if they have enough mass) eventually begin to feed on the helium – both primordial and newly synthesised – fusing it into carbon. Massive stars are able to continue up this chain, cycling through carbon, nitrogen and oxygen, and producing the 'alpha' elements: neon, magnesium, silicon, sulphur, argon and calcium.

Each subsequent element can only be produced if the star is massive enough - else the star cannot ignite the next fusion stage, and cools down to become a white dwarf. Across

[2]Some artistic liberty is taken with the term "burning" in this context. Instead of a combustion reaction, 'hydrogen burning' is a nuclear fusion process whereby 4 protons are fused into a helium nucleus.

the stellar lifetime, and in the final stages of white dwarf formation, material from the stellar surface is expelled through the stellar wind, bringing with it the newly synthesised materials into the universe.

Even more dramatically, if the nuclear fusion chain continues unabated – which is only possible for stars greater than $\sim 8M_{\odot}$ – the star is able to build up a core of iron-56. This leads to the core becoming unstable, leading to photodisintegration and total collapse. As the core collapses, shocks and neutrino physics (which are still not adequately understood in many respects) force the outer envelope of the star outwards in a dramatic event, a **Core-Collapse Supernova** (CCSN), in which additional fusion events occur. This directly scatters the newly synthesised material out into the surrounding galaxy where, like the material ejected through the stellar wind, it is mixed in with the Interstellar Medium (ISM), where it can be subsequently incorporated into the next generation of stars and (importantly for our existence) any planets surrounding them.

Stars are therefore the source of a great deal of the elemental diversity we see around us today, and the exact mechanisms by which different elements are produced, destroyed, recycled and eventually ejected back into the galaxy are subject to continued debate and study. The internal structure of stars remains a complex topic, but we do know that stars can alter their structure drastically in response to minuscule changes to their elemental structure, as the presence of metal pollutants changes how the light and matter within the star interact. The large scale elemental abundances that we observe around us is incredibly sensitive to the microscopic physics occurring deep within the hearts of stars.

Nowhere is this better exemplified than in the prediction of the so-called 'Hoyle Resonance' of carbon-12. Hoyle's argument that such a state must exist hinged solely on the fact that, without this resonance the triple-alpha process (the mechanism by which helium-burning occurs) would stall at beryllium-8, a highly unstable isotope, which would immediately fragment back into helium. Were this the case, the entire mechanism of stellar nucleosynthesis would fail, and our universe would only be capable of transmuting hydrogen into helium. Since this is observably not the case, Hoyle argued, there must be an atomic conspiracy – a resonance at 7.65MeV – which ensured that enough beryllium would fuse into carbon, and allow the chain to continue. This resonance line was experimentally observed soon after: a microscopic discovery motivated by cosmological observations.

1.2 Creation From Beyond the Grave

The story of nucleogenesis, however, does not end with the death of the stars. Although the coherent picture presented by B^2FH robustly addressed the vast majority of observations, there remained a number of unexplained phenomena which would remain puzzling for a number of decades.

The standard Minkowski-Baade classification of supernovae is based upon a phenomenological approach: Type II supernova are characterised by a distinct hydrogen signal in their spectrum, whilst Type I supernovae show no hydrogen. Both types are then further subdivided: Type 1a show silicon lines, Type 1b shows neutral helium but no silicon, and Type 1c shows neither silicon nor helium.

This phenomenological categorisation, however, does not necessarily indicate a common origin for the different types. In fact, mechanisms for Types II and Type Ib/c could all be understood to fall under the broad umbrella of core-collapse supernovae in stars of different types: Type II being the most common CCSN, whilst Type Ib/c being CCSN in massive stars which have started to shed their outermost layers.

Type Ia, however, eluded a comprehensive understanding for a significant amount of time (and debate still surrounds their exact nature). It was recognised as early as 1960 (Hoyle & Fowler, 1960) that they required burning of degenerate matter, but the remarkable homogeneity of the spectra of Type 1a Supernovae (SNIa) provided a significant challenge: within a small calibration, all SNIa look almost identical to all other SNIa. This is in contrast to other CCSN events, where there is significant differences between individual events. Whilst this homogeneity proved a fascinating tool as a standardisable candle, a consensus on the mechanisms which could produce SNIa remained elusive.

The current favoured model for SNIa formation first emerged in the 1970s (Wheeler & Hansen, 1971; Whelan & Iben, 1973), though it remained controversial for much of the next decade (Woosley & Weaver, 1986). Under the current consensus model, SNIa form when a low-mass binary star pair evolves enough that one member (the progenitor) exits the Main Sequence and dies, leaving behind a carbon-oxygen white dwarf[3]. The CO white dwarf is, in essence, the remaining core of the star which was too light to initiate the next phase of fusion, and so is in a dormant nuclear state. However, as the companion continues to evolve, it can overflow its Roche lobe and donate material onto the CO white dwarf.

[3]The requirement that the progenitor form a CO white dwarf places strict mass constraints on the progenitor: too light, and the white dwarf will be helium, which does not detonate; too massive and the star instead undergoes a CCSN or an Electron-Capture Supernova (ECSN), leaving behind a neutron star

As the white dwarf accretes matter, its mass increases. If this process were to continue unabated, the white dwarf would eventually reach the Chandresekhar limit ($\approx 1.4 M_{\odot}$), and collapse into a neutron star. However, as the WD has been accreting matter, the internal temperature has been rising as the gravitational forces compress the new matter. In a normal star, the pressure would correspondingly increase, venting the excess energy. The degenerate nature of WDs, however, prevents this, so the temperature rise continues unabated until the temperature rises enough to power a convection phase ('simmering'), and eventually a deflagration front passes through the white dwarf, causing the ignition of the remaining nuclear fuel.

The result is a complete nuclear disruption of the white dwarf, resulting in rapid burning of the material all the way up to the nuclear statistical equilibrium at ^{56}Ni. The result (after subsequent nuclear decay) is the formation of $\sim 1 M_{\odot}$ of ejecta primarily composed of the iron peak elements: iron, cobalt, manganese, chromium.

Iron produced in this fashion is thought to make up approximately half of the total galactic iron budget. This might seem surprising given that White Dwarfs were long thought to simply be the dormant 'corpses' of stars whose nucleosynthesis days had long since passed.

It is perhaps even more surprising, therefore, that neutron stars and black holes – the remnants of CCSN events – are also thought to contribute to the nucleosynthesis. Despite being more alien (and not composed of elemental material), these objects can undergo incredibly energetic mergers, such as that detected in the famous multi-messenger observation of GW170817 (Abbott et al., 2017). These interactions produce neutron rich material suitable for r-process nucleosynthesis, which is the subject of Chapter 5.

The existence of both SNIa and Neutron Star Mergers (NSM) as significant contributors to the elemental abundance ratio tells us that, although stars are the supreme source of post-primordial elements, the timescales over which they contribute their materials to the galaxy can vary hugely, and are not simply confined to the standard stellar lifetimes we would naïvely expect.

1.3 A Dance in Spaces & Time

We infer the chemical abundance of the galaxy through spectroscopic surveys, which peer into the upper atmospheres of stars and determine the chemical composition. As the outermost layers of the star are (modulo any lifetime nucleosynthesis and mixing) composed of the same gas which initially formed the star, we can then compute the abundances of the star forming gas when this star was formed.

In an ideal world, astrometry would provide us with an accurate location of the star, whilst spectroscopy could then be employed to measure the chemical content of a star, and then determine its age. We would then have a map of the chemical abundance of the cold, star-forming gas at various points in space and time, which would then constrain and correct our theoretical models.

It turns out, however, that the most intuitive coordinates, those of everyday space and time, are often far more fraught than we would ideally like.

Firstly: stars are not stationary objects. Although the mean velocity of stars in the Milky Way is almost entirely azimuthal, individual stars themselves often exhibit significant deviation from circularity. Over the course of their lifetime, stars can be scattered by various axisymmetric features and resonances such as the spiral structure of the Milky Way, or its central bar. As a result, a star at the solar radius may have come from the very inner galaxy, or the outermost edges, and hence encode a vastly different chemical history than the rest of our surroundings.

Secondly, it is often incredibly difficult to accurately determine the age of a star. The Milky Way has a remarkably short chemical equilibrium timescale, settling into an equilibrium ~ 1Gyr after a perturbation. Given that the current best estimates on stellar ages often have errors significantly in excess of this time, any time-dependence in the enrichment would be washed out in the noise of the stellar ages. Efforts to use temporal coordinates as a primary driver of historical studies are, at the current time, doomed to almost certain failure.

The field of galactic archaeology, therefore follows the pioneering work of Tinsley (1979) where, instead of absolute abundances vs. time, the object of interest is the relative, logarithmic abundances of the elements, [X/Y] such that:

$$\begin{aligned}\left[\frac{\mathrm{X}}{\mathrm{Y}}\right] &\equiv \log_{10}\left(\frac{n_{\mathrm{X}}}{n_{\mathrm{Y}}}\right) - \log_{10}\left(\frac{n_{\mathrm{X}}^{\odot}}{n_{\mathrm{Y}}^{\odot}}\right) \\ &= \log_{10}\left(\frac{m_{\mathrm{X}}}{m_{\mathrm{Y}}}\right) - \log_{10}\left(\frac{m_{\mathrm{X}}^{\odot}}{m_{\mathrm{Y}}^{\odot}}\right)\end{aligned} \tag{1.1}$$

Here n_{X} and m_{X} are the number and mass densities of element X respectively, and $\odot$ represents the solar value. By comparing the abundance of elements with different source profiles, we can therefore gain an insight into the processes which affect the galactic chemistry on both a galactic and stellar timescale: the total metal content correlates strongly with the star formation rate (SFR) in the galaxy, whilst the $\left[\frac{\alpha}{\mathrm{Fe}}\right]$ ratio indicates the balance of CCSN and SNIa events.

Whilst the spatial and temporal properties of the stars are in themselves insufficient for a rigorous study of the galactic chemical history, they can provide an invaluable tool for further constraining the relative abundance data. Since the relative abundances are necessarily a convolution of many different physical processes, it is necessary to ensure that a given model replicates the observations in multiple different chemical planes, as well as reproducing the known temporal and spatial properties of the galaxy.

1.4 Revolutions & Revelations: This Thesis

The primary goal of this thesis is to collect and collate the cutting edge advances made in Galactic Chemical Evolution (GCE) and its associated fields over the past decade into a cohesive whole, which will allow us to shine a spotlight onto some long-unanswered questions.

We begin, in Chapter 2, by introducing the basic observational and theoretical framework which will underpin the rest of this thesis. After discussing the observational data which our models must reproduce to be considered valid, we shall outline both the general features which all of our chemical evolution models possess, before moving onto the specific properties of the SAGEM (Chapter 3) and RAMICES II (Chapter 4) models which form the two prongs of our subsequent studies.

In Chapter 5 we use these models to study the half-century old question of the origin of r-process material within the galaxy, particularly focusing on the conflict between the traditional (and near-consensus) NSM model, in which inspiralling Neutron Stars eject vast quantities of neutron rich material sufficient to ignite the r-process, and the collapsar model, in which jet-envelope interactions in a collapsing, rapidly-rotating star provide a potential alternative site. We investigate the effect that the theoretical properties of collapsars would have on the resulting chemical histories by exhaustively exploring the vast parameter space associated with our chemical evolution models.

Chapter 6 adapts the SAGEM formalism to scrutinise the new stellar evolution models of Limongi & Chieffi (2018) which incorporate, for the first time, stellar rotation as a major driver of stellar nucleosynthetic variability. We use this as a testbed to develop a simple tool which enables future works to rapidly assess their models and identify incompatibilities with the current understanding of GCE.

In Chapter 7, we utilise RAMICES II to briefly study the phenomena of 'gas caustics', in which 'hoop' and 'loop'-like structures in the gaseous plane imprint themselves in drastic and often unpredictable fashion on the stellar abundance plane.

We conclude with Chapter 8, in which we summarise our findings, and highlight potential future uses for the formalisms and tools developed over the course of this research.

Everything starts somewhere, although many physicists disagree.

— Terry Pratchett

2 Background & Theory

Contents

In this chapter we shall endeavour to give a brief overview of the observational evidence and theoretical underpinning of GCE , which we shall then use to motivate our further studies in the subsequent chapters.

2.1 Observational Data

2.1.1 Global Distributions

Figure 2.1 reproduces Figure 1 from B^2FH, showing the (smoothed) general trends of the solar neighbourhood abundance distribution. B^2FH identified 5 major features of this curve, and offered the following explanations for their origins:

1. **Exponential decrease in abundance with increasing atomic mass**

   Explanation: Most nucleosynthesis processes are sequential in transmuting one element to the next, with processes capable of reaching higher atomic weights being rarer than those capable of reaching lower atomic weights.

2. **Under abundance of D, Li, Be B**

   Explanation: These elements are consumed in many stellar processes, leading to inefficient production and even net destruction

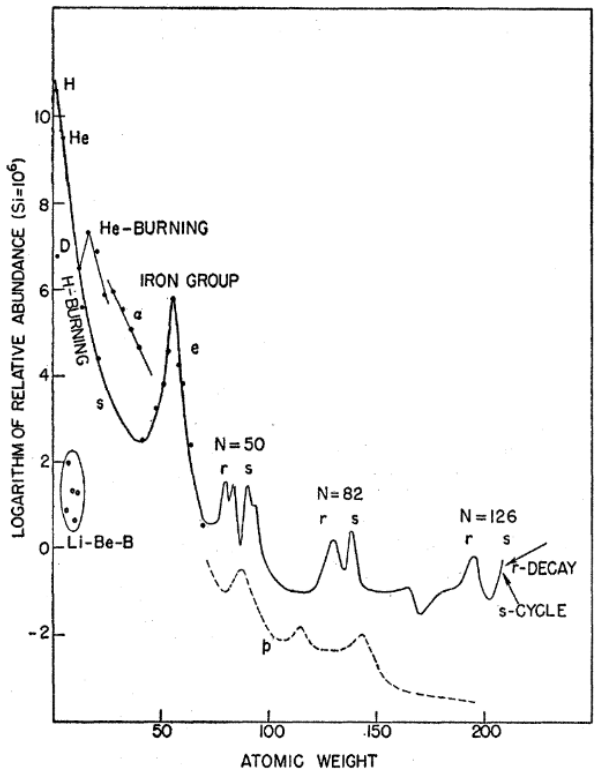

Figure 2.1: General features of the solar neighbourhood abundance distribution, reproduced from Figure 1 of B^2FH, and derived from the data of Suess & Urey (1956). B^2FH identified a number of key features in this plot and were able to explain them using a variety of nuclear and stellar processes.

3. **High abundance of O, Ne, Ca, Mg, Ti etc.**

   Explanation: These "α" elements are produced through a highly efficient helium burning process.

4. **Strong peak around Fe**

   Explanation: Explosive Nucleosynthesis (supernovae) which reaches nuclear statistical equilibrium per Hoyle (1954), producing matter peaked around ^{56}Ni

5. **Double-peaks around the $N = 50, 82$ and 126 magic numbers**

   Explanation: The rapid (r) and slow (s) neutron capture processes (see Chapter 5)

2.1.2 Tinsley Diagrams

Whilst these broad explanations for the general trends were enough to solidify stellar nucleogenesis as the origin of the chemical elements in the universe, there remain a number of other observational trends which require a more nuanced understanding of the history of the galactic enrichment. In particular, following Tinsley (1979), the most common tool used in GCE are relative enrichment plots, usually in the form of [X/Fe]-[Fe/H] diagrams, with $\left[\frac{\mathrm{X}}{\mathrm{Y}}\right]$ defined as in Eq. (1.1).

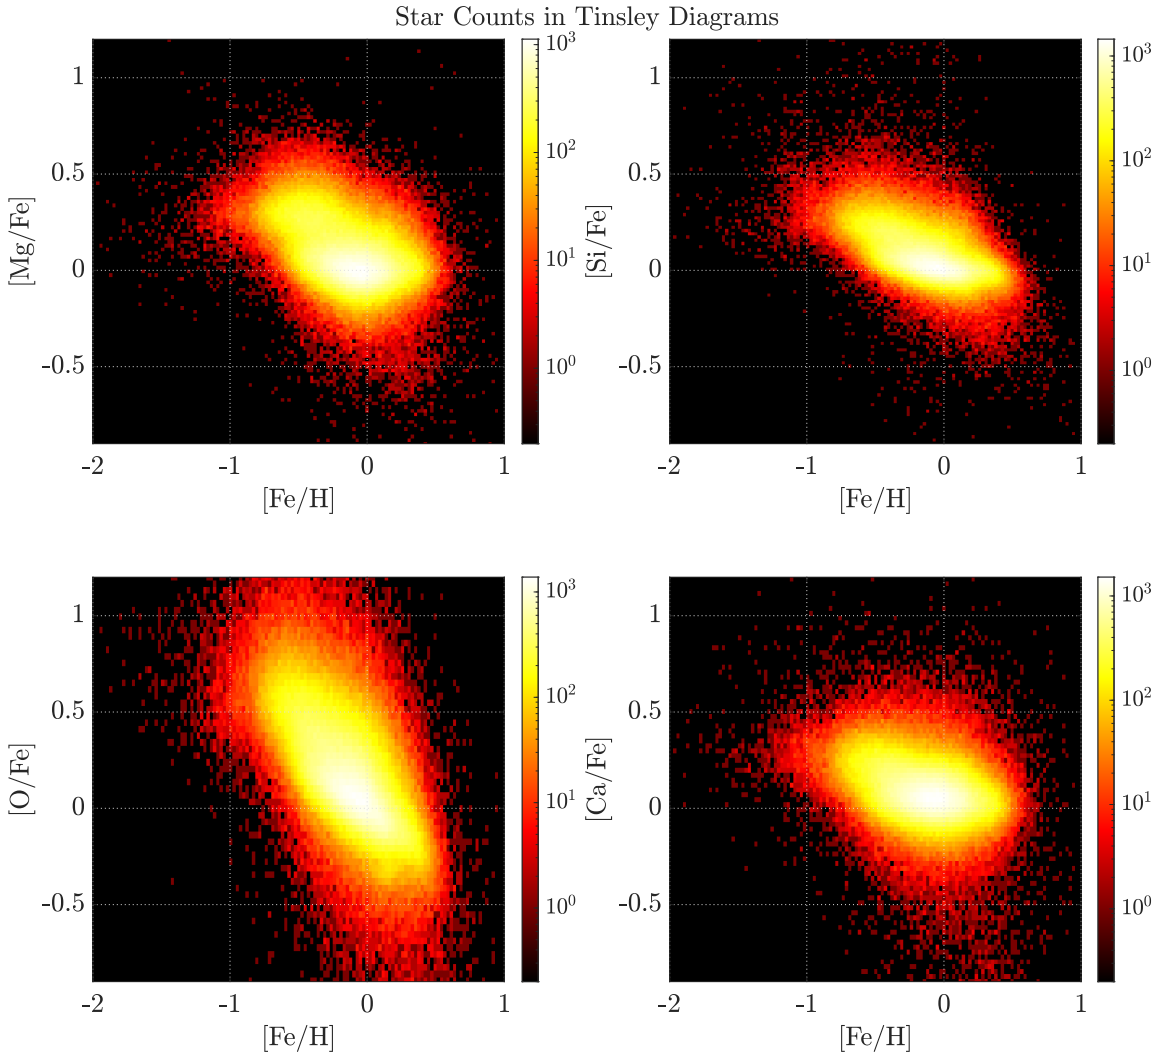

Figure 2.2: Tinsley Diagrams for several α elements, showing the differing behaviour across different elements for or the 58500 stars from the GALAH DR3 (Buder et al., 2021) database with Fe, Mg, O, Si, Ca and Al measurements with both upper and lower error bounds, and which pass the quality cuts ($\sigma_X < 0.1$dex) and chemical-halo cuts (see Hayes et al. 2018), for (*clockwise from top left*) [Mg/Fe], [Si/Fe], [Ca/Fe] and [O/Fe] vs [Fe/H].

Figure 2.2 show a number of such 'Tinsley Diagrams' for four α elements: [Mg/Fe], [Si/Fe], [Ca/Fe] and [O/Fe]. We see indications of additional structures, and also see that although these are all α elements and hence, through the picture presented in B^{2}FH, produced in the same nucleosynthetic process, these structures differ between the individual elements.

The most notable structure is the infamous α-bimodality, which is most visible in [Mg/Fe]. There is a clear indication of two different populations: one metal poor[1] but α-rich – which means having a higher α/Fe ratio than the sun, not necessarily more abundant in α in absolute terms – and the other more metal rich, but α-poor. Figure 2.3 shows a projected

[1] [Fe/H] is often used as a proxy for net metallicity.

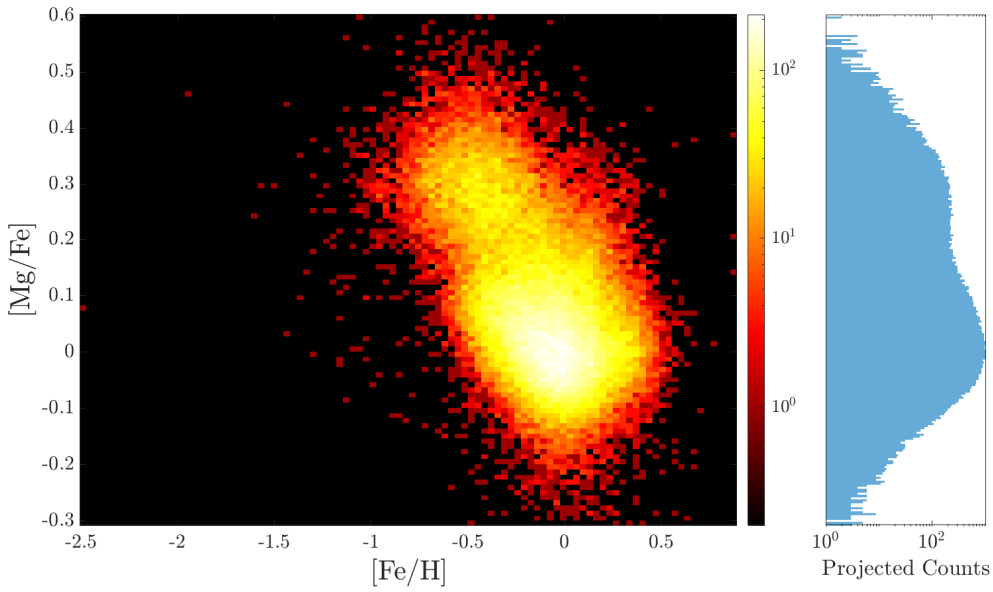

Figure 2.3: (*left*) The first panel from Figure 2.2 with the further restriction that the observed error, $\sigma_{[\mathrm{Mg/Fe}]} < 0.2$dex, with (*right*) a projected star count in [Mg/Fe], demonstrating the clear population bimodality.

histogram of stars in $\left[\frac{\mathrm{Mg}}{\mathrm{Fe}}\right]$, where the two populations are evident.

As we might expect given that the universe started out from primordial material with little-to-no metals, we also see in Figure 2.4 a clear temporal evolution in [Fe/H]: as the universe ages, the [Fe/H] abundance increases. We therefore also see that the two α populations are separated in age: the high-α population is generally much older, whilst the low-α population is younger.

In kinematic space, these high-α and low-α stars strongly overlap with the kinematic thick and thin discs respectively. Given the significant overlap in kinematic space, it is therefore common to delineate these components by means of a cut in [Mg/Fe]-[Fe/H] space which isolates the older of the two bimodal populations. When we refer to the 'thick disc', it should therefore be understood that we are referring to this chemically identified population, which is older and α rich, but metal poor, and equally the 'thin disc' refers to the younger, α-poor but metal rich population.

2.1.3 Halo Pollution

We note that the sample of stars presented in Figs. 2.2 and Figs. 2.3 are explicitly those stars which passed a cut designed to remove members of the stellar halo. The stellar halo is an approximately spherical, spatially extended and kinematically distinct population which

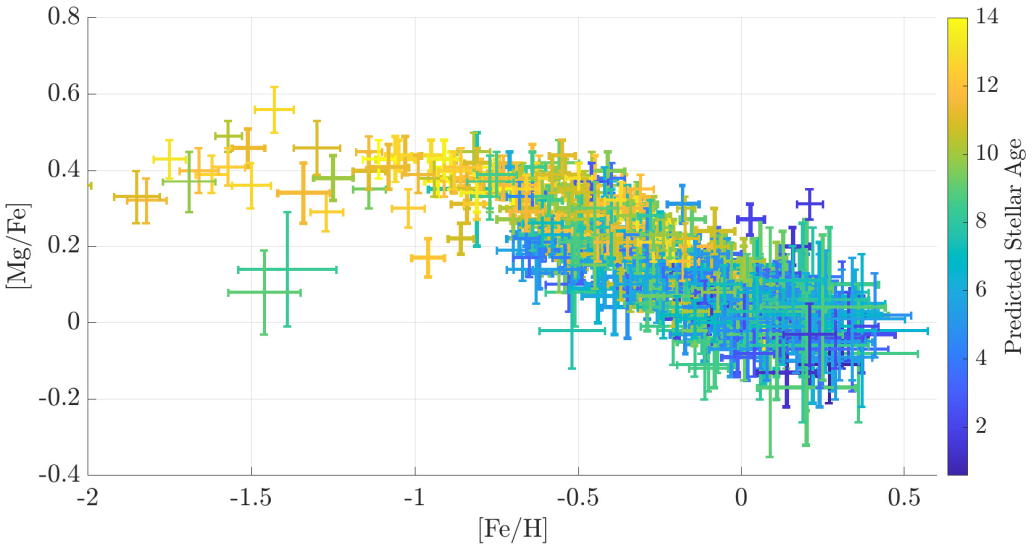

Figure 2.4: A Tinsley Diagram in [Mg/Fe] demonstrating the variation in ages of stars. Data is 593 G and F dwarf stars drawn from Bensby et al. (2014) – note that this is a different sample than Figure 2.2 by necessity of accurate age determinations.

surrounds the galaxy out to large distances, with many of its stars composed of those accreted during galaxy mergers (Binney & Tremaine, 2008; McCarthy et al., 2012).

These stars, therefore, represent a superposition of chemical histories of multiple galaxies of different ages and masses. Whilst they undoubtedly contain useful information regarding the evolution and assembly history of the Milky Way, when attempting to study the in-situ chemical evolution, they only serve to obscure any signal which might be present.

We therefore follow the recommendation of Hayes et al. (2018), and use a cut in [Al/Fe] to determine members of the stellar halo in order to remove them from our samples.

2.1.4 Radial Structure in the Galaxy

As we shall see in § 4.7 , it is rather easier to generate an estimate of the properties of the cold, star forming gas than it is to predict the expected properties of a given spectroscopic catalogue, as the latter requires synthesising the particular nuances of the chosen survey, and convolving those with the entire star forming history of the model galaxy.

It is, however, far harder to actually observe the chemical enrichment of the star forming gas than it is to sample the spectra of stars. It is therefore convenient to find a population of stars which have most recently 'sampled' the star forming gas, and use those to infer the cold gas metallicity. The classical Cepheids are a well studied population, and their formation mechanism requires them to be high mass, and short lived: ideal candidates for such a probe.

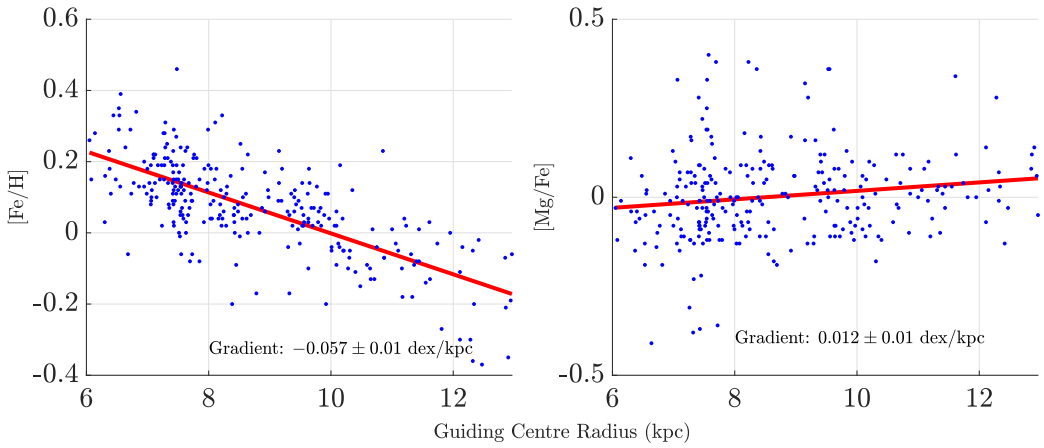

Figure 2.5: Radial abundance data, as given by Luck & Lambert (2011), for classical Cepheid stars as a function of the guiding centre radius, R_g. The large masses (and hence short lifetimes) of Cepheids mean that they sample the recent abundance of the star forming gas. We see the classical metallicity gradient of the galaxy with $\mathrm{d}\left[\frac{\mathrm{Fe}}{\mathrm{H}}\right]/\mathrm{d}R \approx -0.06$ dex kpc^{-1}, and an inverted gradient in $\mathrm{d}\left[\frac{\mathrm{Mg}}{\mathrm{Fe}}\right]/\mathrm{d}R$

Figure 2.5 shows the radial distance of a sample of such stars, demonstrating the radial distribution of chemicals: more iron rich in the centre of the Milky Way, but more α-poor[2].

2.2 The Galactic Canon

Having introduced the basic observational data which we are attempting to replicate, we first take a moment to consider the accepted descriptions of the origin of many of the structures and substructures that we have introduced.

2.2.1 The Origin of α-Bimodality

The canonical explanation for the origin of these two distributions lies in the time dependence of two key processes which produce α elements and Fe: CCSN and SNIa.

CCSN arise from the death of stars with masses $\gtrsim 10M_\odot$, and give rise to ejected material which is rich in both α elements and iron peak elements (Chieffi & Limongi, 2004). The precise amount of this material is still subject to study, and varies depending on the initial properties of the star, but is generally super-solar in $\left[\frac{\alpha}{\mathrm{Fe}}\right]$ (Limongi & Chieffi, 2018).

Though their precise mechanism is still under debate, SNIa are generally believed to arise from either a merger between two Carbon-Oxygen White Dwarfs (COWD) in a tight binary system, or mass-accretion between a COWD and a stellar companion. Whichever

[2]Due to the difficulty in disentangling 3D NLTE effects in Cepheid stars, it is convenient to use Mg as a probe for the whole set of α elements

mechanism dominates, the subsequent explosive nucleosynthesis is expected to completely disrupt the progenitor, and lead to ejecta composition close to that predicted by nuclear stellar equilibrium: highly enriched in the iron peak elements relative to the α elements (see Table 6.5, derived from Iwamoto et al. (1999)).

The lifetimes of stars sufficiently massive enough to go CCSN are much smaller than the evolution timescale of the galaxy, and so CCSN are often seen as co-temporal with star formation events. SNIa, however, require that a lower mass star $\left(3 < \frac{M}{M_{\odot}} < 8.5\right)$ dies, and then undergoes a period of inspiralling and/or accretion from their binary companion. This imposes a temporal delay between the onset of star formation, and the time at which the first SNIa can occur - this value is often quoted to be $150 - 600$Myr.

CCSN are therefore occurring almost immediately after the onset of star formation in the galaxy, and hence the chemical composition of the Inter-Stellar Medium (ISM) is being polluted almost entirely by material of the same composition: the early time $\left[\frac{\alpha}{\mathrm{Fe}}\right]$ value is therefore representative of the $\left[\frac{\alpha}{\mathrm{Fe}}\right]$ abundance of CCSN, averaged over, for example, the different masses of stars producing the material. We explore this in detail in Chapter 6.

This CCSN-dominated picture is interrupted as soon as SNIa 'turn-on' after their prescribed temporal delay. The SNIa ejecta have a massively sub-solar $\left[\frac{\alpha}{\mathrm{Fe}}\right]$, so the SNIa onset acts to pull the $\left[\frac{\alpha}{\mathrm{Fe}}\right]$ abundances down. This downward 'knee' is seen in many elements, not just the α elements, since it arises from a sudden increase in iron production: the only elements which we would expect not to show a similar pattern are those which are co-produced with iron in such events (Co, Mn etc.).

We note that this explanation merely serves to explain why there is a temporal evolution from high-α to low-α, not why there is a seemingly discontinuous bimodality in the two populations. This 'discontinuity' is often used to argue for a distinct break between two phases of star formation within the galaxy - an early time phase which formed the thick disc, and a later time phase which formed the thin disc, with a distinct gap between the two (Chiappini et al., 1997).

We, however, follow the argument that this discontinuity is present due to a combination of an exponentially declining SFR (contributing to a rapid chemical enrichment period before SNIa turned on), and short gas depletion timescale, which causes the system to move rapidly between equilibria after a perturbation is applied. Prior to SNIa turn on, the galaxy approached a chemical equilibrium at high $\left[\frac{\alpha}{\mathrm{Fe}}\right]$, which causes a buildup of star formation at the thick disc abundance values. After SNIa turn on, the SNIa ejecta disrupts the equilibrium and the system rapidly shifts to the new, low- $\left[\frac{\alpha}{\mathrm{Fe}}\right]$ equilibrium, where

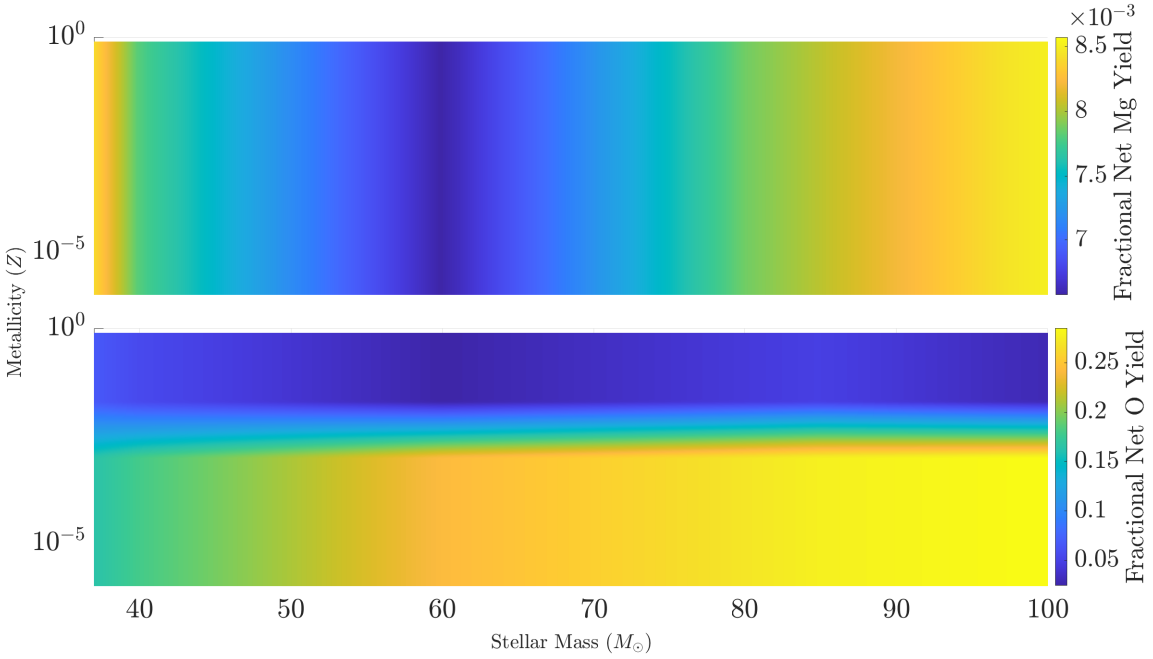

Figure 2.6: The (net) Magnesium (*top*) and Oxygen (*bottom*) yields expected for a star of initial mass $M > 35M_\odot$ and metallicity Z, drawn from the heterogeneous yield sets discussed in §4.6. Whilst the magnesium shows no metallicity dependence at these masses, the oxygen yield shows a strong dependence, with the net production being highly suppressed at high metallicities.

another buildup occurs. This produces two distinct populations with a smooth transition region in between. Schönrich & Binney (2009a) (henceforth SB09) showed that this arises in all cases with an exponentially declining SFR.

2.2.2 Differences in Elements

As the α elements are expected to be produced in the same process, and hence co-temporally with each other, the differences between α elements, as seen in Figure 2.2, must therefore be explained via metallicity dependencies in their yields, and on their relative presence in SNIa ejecta.

Figure 2.6 shows the mass and metallicity dependence of the net Mg and O yield for stars of $M > 35M_\odot$, as drawn from Maeder (1992), Marigo (2001), Chieffi & Limongi (2004) and Limongi & Chieffi (2008). We see that although the Mg yield shows no variation in metallicity at these masses, the Oxygen yield shows a drastic drop in Oxygen synthesis, due the strong stellar wind-metallicity coupling, and the subsequent atmospheric stripping and exposition of the outer core at high metallicities (see Maeder (1992) for more details).

Although these stars are comparatively rare, their short lifetime means that they rapidly and heavily impact the enrichment of the ISM. Oxygen, therefore, undergoes a strong drop in synthesis rates at higher metallicities. Since the galaxy exhibits a metallicity gradient

(higher Fe at lower radii), this amplifies the [O/Fe] gradient such that [O/Fe] is lower in the inner, metal rich galaxy than it is in the outer, metal poor galaxy.

This is therefore the origin of the large (~ 0.6dex) drop in [O/Fe] between thick and thin discs, as well as the inclined thin disc structure observed in Figure 2.2, though we note that this inclined structure is amplified in part by the selection function of the survey (see §4.7).

We also note, from Table 6.5, that although SNIa produce material with $\left[\frac{\alpha}{\mathrm{Fe}}\right] \ll 0$, the individual α elements have differing [X/Fe] ratios in the ejecta from SNIa events - with $\left[\frac{\mathrm{O}}{\mathrm{Fe}}\right] \approx \left[\frac{\mathrm{Mg}}{\mathrm{Fe}}\right] \approx -1.4$, $\left[\frac{\mathrm{Ca}}{\mathrm{Fe}}\right] \approx -0.3$ and $\left[\frac{\mathrm{Si}}{\mathrm{Fe}}\right] \approx -0.4$. Therefore O and Mg will see a steeper drop from their thick values (with O magnified further by the decreased synthesis effect), whilst Ca and Si will exhibit smaller drops, and hence their bimodal populations will strongly overlap.

Other Elements

This, of course, is not the full story, but provides a bird's-eye view of the accepted explanations for the largest scale structures within the body of observational evidence.

Some elements strongly deviate from these basic patterns – Fluorine, for example, shows an almost perfectly inverted Tinsley diagram to [Mg/Fe], with [F/Fe] *increasing* (seemingly) at the SNIa onset Guerço et al. (2019). This difference, however, lies in the highly debated production mechanisms for Fluorine, with strong metallicity dependencies and long-lived stars likely to play a major role.

2.3 Basic GCE Theory

Having presented a description of the system and observables we are attempting to model, we now present an outline of the underlying physics of our GCE models, and discuss the common ingredients and physics which motivates our both the rapid, parameter space scanning model SACEM and the full simulation model RAMICES II.

2.3.1 Fundamental Assumptions

Throughout all of our work, we make a number of fundamental assumptions which simplify our models considerably.

Continuum Approximation

We explicitly limit ourselves to the the case of a large 'continuum' galaxy where, for example, the metal content of a single star or supernova is negligible compared to the metal content of the remainder of the galaxy. We are then able to concern ourselves primarily with spatially

resolved rates of birth and death and probabilistic motion of material, rather than tracking the location of each star in the galaxy.

Well-Mixed Approximation

The RAMICES II model is spatially resolved, meaning that we have different 'bins' representing the gas and stars at different radii, and both RAMICES II and SACEM differentiate between hot and cold gas.

Within a given spatial and thermal bin, we assume that the chemical abundances are uniform and instantaneously mixed into the bulk of the gas. Chemical imhogeneity is modelled through differences and subsequent interactions between bins

Atmospheric Abundance Representation

We assume that the spectroscopic abundances measured from stellar atmospheres represent the chemical abundance of the cold gas which formed the star at the time and position of the star's formation, neglecting effects such as lifetime nucleosynthesis, mass-fractionation (heavier elements sinking into the centre of the star) and dredge up events. In practice, we must therefore be careful in certain instances with comparing our observations against spectroscopic catalogues for some elements in particular. For example, the CNO cycle can drastically alter the balance of individual CNO elements from their birth ratios, with nitrogen being particularly susceptible, leading to us omitting it from our yield network entirely.

Pre-Ordained Stellar Lifetimes

We assume that all stars evolve as isolated entities without outside interference. The stellar properties at all subsequent times are therefore uniquely defined by their initial mass, abundances and stellar rotation.

Strongly Coupled Star Formation

We assume throughout that the SFR is strongly coupled to the gas content of the reservoir in which the stars are forming, and that this is the same reservoir which the nucleosynthetic events are polluting. This is in contrast to some other models in the literature in which the SFR is externally imposed onto the gas reservoir (Siegel et al. 2019, henceforth Si19), or the star forming efficiency fine tuned by hand as a function of time to match observed data (Chiappini et al., 1997).

2.3.2 From Gas to Stars, & Back Again

As the central engines of nucleosynthesis, the rate at which stars are being born (depleting the gas reservoir) and dying (returning metals to the ISM) is the most important ingredient in a GCE model.

Star Formation

The SFR, $\rho_{\rm sfr}(t)$, tracks the rate (in mass-units) at which cold gas is being compressed into stars at each moment in time, such that the mass of stars in the reservoir M_* and the star forming gas, $M_{\rm c}$ (see §2.3.3) are related by the star formation law,

$$\left(\frac{{\rm d}M_*}{{\rm d}t}\right)_{\rm sfr} = -\left(\frac{{\rm d}M_{\rm c}}{{\rm d}t}\right)_{\rm sfr} = \rho_{\rm sfr}(t) \tag{2.1}$$

Following the work of Schmidt (1959), it has been common to express the SFR in terms of the surface density of the stars, Σ_* and the total gas surface density $\Sigma_{\rm gas} = \Sigma_{\rm c} + \Sigma_{\rm h}$ such that

$$\dot{\Sigma}_* = \kappa \Sigma_{\rm gas}^n. \tag{2.2}$$

Kennicutt (1998) found that $n = 1.4 \pm 0.15$, and also indicated that there should be a break in the power law at a density $\Sigma_{\rm break}$ such that

$$\dot{\Sigma}_* = \begin{cases} \kappa \Sigma_{\rm gas}^{n_1} & \Sigma_{\rm gas} > \Sigma_{\rm break}, \\ \kappa \Sigma_{\rm break}^{n_1 - n_2} \Sigma_{\rm gas}^{n_2} & \text{else.} \end{cases} \tag{2.3}$$

The recommended value of n_2 is ~ 4, but this value (and the value of $\Sigma_{\rm break}$) is still subject to significant uncertainty, and GCE models typically calibrate the values to match observed stellar surface densities.

Each spatial bin ('zone') is assumed to have a constant surface area A_i and we assume the density properties are uniform across each zone. The SFR itself, $\rho_{{\rm sfr}\,i}(t)$ in each spatial zone is therefore simply computed from

$$\rho_{{\rm sfr}\,i}(t) = A_i \dot{\Sigma}_{*,i}. \tag{2.4}$$

We note that, under the well-mixed assumption (§2.3.1), star formation depletes the cold gas reservoir uniformly, thereby removing synthesised metals from the ISM. This is important in the case where the gas reservoir is then 'topped up' by another source, such as radial flows of gas between radially-segregated reservoirs or by the inflow of metal-poor gas from the Circum-Galactic Medium (CGM, see §2.3.5), as this results in a decrease in the metal abundance as well as the absolute mass – we would require a continuous source of

new metals (i.e. supernovae) to counter this process. This in turn suggests that the galaxy will eventually reach an equilibrium between the synthesis processes and the star formation and accretion processes, resulting in a finite maximum metallicity.

Population Makeup

Whilst $\rho_{\rm sfr}(t)$ gives the total *mass* of stars which are forming, it gives us no information about the distribution of the properties of the individual stars. To describe these properties, we introduce the Initial Mass Function (IMF), ζ, which determines the properties of a population of stars produced through a single burst of star formation. We assume that the IMF is a constant function of metallicity, such that $\zeta(M)\mathrm{d}M$ gives the fraction of stars born with masses in the range $[M, M + \mathrm{d}M]$.

We follow the Chabrier (2003) form for the IMF, which obeys the Salpeter (1955) form of the IMF for $M > M_\odot$, but with corrections for lower mass stars:

$$\zeta(M) = \begin{cases} aM^{-1} \exp\left(-\frac{\log_{10}\left(\frac{M}{0.08}\right)^2}{2\times 0.69^2}\right) & \text{if } M < M_\odot, \\ bM^{-\alpha} & \text{else.} \end{cases} \tag{2.5}$$

The values a and b are completely constrained by the requirement that the function is continuous and normalised across the domain $[0, \infty]$. The value of α is determined by observations, and typically constrained to be 2.3 ± 0.3.

Over an infinitesimal timestep $\mathrm{d}t$, a mass of new stars $M_{\rm new} = \rho_{\rm sfr}(t)\mathrm{d}t$ is formed. We discretise the masses of stars onto a grid of N bins with edges $M_0 < M_1 \ldots < M_N$. The number of stars formed with masses in the $i^{\rm th}$ bin is given by:

$$\mathrm{d}N_i = \frac{\rho_{\rm sfr}(t)\mathrm{d}t}{\int_0^\infty M\zeta(M)\mathrm{d}M} \int_{M_i}^{M_{i+1}} \zeta(M)\mathrm{d}M \tag{2.6}$$

Over a finite timestep Δt, the true number of stars formed ΔN should be an integer. However, with reference to the continuum approximation (§2.3.1), we allow the non-physical case of non-integer star counts.

Stellar Death

Once a star has been born, it persists within the galaxy, holding an imprint of the chemical abundance from which it was formed in its stellar atmosphere, until the point at which the star undergoes the process of stellar death. At this point, those materials are then variously ejected back into the ISM, into the CGM or folded into remnant reservoirs such as white dwarfs, neutron stars and black holes. The stellar lifetimes are therefore of vital importance in determining when synthesised metals are made available for stellar and planetary formation.

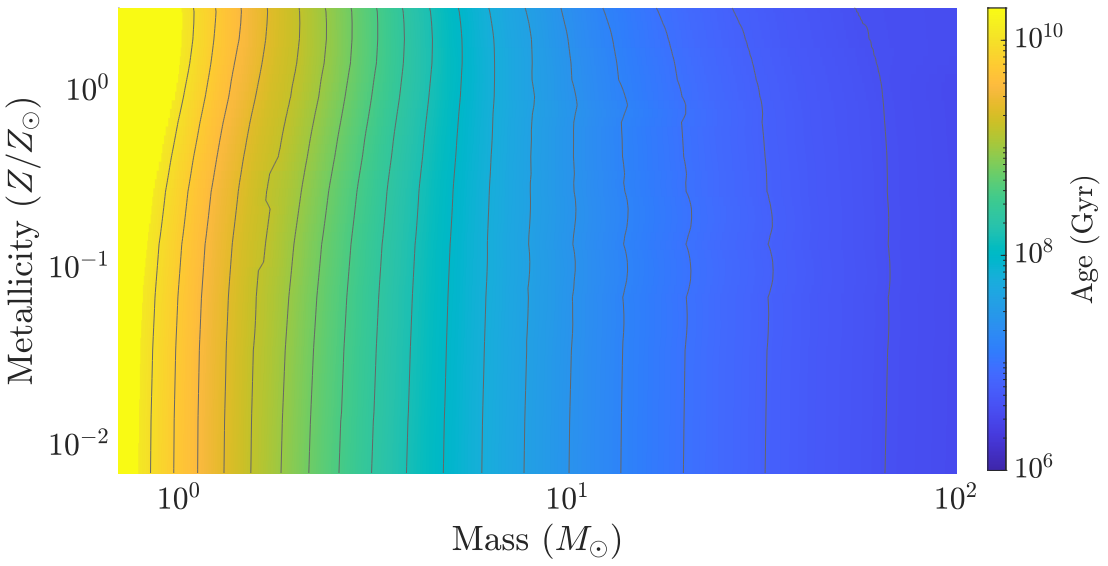

Figure 2.7: The extracted stellar lifetimes from the PARSEC code as a function of initial mass and metallicity. Overplotted are contours of constant age with spacing $\Delta \log(\text{Age/Gyr}) = 0.2$. For simplicity, stars with lifetimes greater than 14Gyrs are not computed - we simply assume these stars to be 'immortal' and set their age to be greater than the simulation runtime.

Following §2.3.1, we make the simplifying assumption that each star evolves as an isolated body, and hence the evolution of the star is perfectly constrained by the Zero-Age Main Sequence (ZAMS) properties, which consist entirely of the mass (M), metallicity (Z) and rotation speed (v_{zams}).

The impact of rotation is the subject of Chapter 6 of this work, and remains a field of study in its infancy. We therefore neglect v_{zams} for the time being, and instead determine the lifetime of stars based on (M, Z).

Utilising the stellar evolution tracks of the PARSEC code (Bressan et al., 2012; Tang et al., 2014; Chen et al., 2015), extended to include the TP-AGB phase (Marigo et al., 2017; Pastorelli et al., 2020), we are able to compute the expected time between the star entering the main sequence and the beginning of the remnant phase as a function of the initial properties (M, Z). The results of these computations are shown in Figure. 2.7.

Remnant Behaviour

After the stellar lifetime is over, a portion of the mass is returned to the ISM or CGM as either 'untouched' gas, or with additions due to stellar nucleosynthesis. However, it is known that even when undergoing CCSN events, there is not a complete disruption: there is a remnant left behind.

Table 2.1: The fate of stars, and their remnants, given the initial mass of the star. We note that, for simplicity, we have omitted metallicity dependence on the nature of the stellar remnants (though it can affect the remnant mass).

Progenitor Mass ($M_\odot$)	**Remnant Type**	**Remnant Fate**
< 0.5	Helium White Dwarf	Immortal
$0.5 \leq M < 8.5$	CO White Dwarf	Most immortal, some SNIa
$8.5 \leq M < 40$	Neutron Star	Most immortal, some NSM
> 40	Black hole	Immortal

This remnant can take one of several forms, which are enumerated in Table 2.1. As noted in the final column of this table, many of these remnants are not sterile, dormant objects – they can still undergo high energy processes which contribute the the chemical budget of the galaxy. In particular, CO white dwarfs are the progenitors of SNIa events, and Neutron stars can undergo mergers which we have observed to produce r-process elements (Chornock et al., 2017; Rosswog et al., 2018). We therefore need to keep track of these remnants as important sources of chemical synthesis.

There are two important caveats to this, however - firstly, not all CO white dwarfs (COWD) are viable SNIa progenitors - under the current understanding, COWDs can only undergo a SNIa event if the COWD is sufficiently massive ($\gtrsim 3.2 M_\odot$), and the progenitor existed in a tight binary system with another star which can donate material onto the COWD. Equally, Neutron Star Merger (NSM) events can only occur in binary systems where both progenitors have undergone CCSN into neutron stars. The exact fraction of stellar progenitors which undergo these processes is highly unconstrained at the moment, and subject to a number of degeneracies with other parameters, and so the associated fractions will have to be calibrated by hand.

Secondly, there is a hard limit on the amount of time which can pass between the onset of star formation and the time at which the subsequent remnant processes can occur – this accounts for both the original stellar lifetime, and the subsequent inspiralling and accretion phases. Even after this lower time limit has occurred, there is a large variability in the time needed for a remnant event to be triggered, a function of, for example, the initial separation of the progenitor stars in the binary system. After the minimum time scale has passed, there is therefore then a period of stochastic decay where the remnants gradually undergo these processes. This is the origin of the Delay-Time Distribution (DTD) in Eq. (3.12).

2.3.3 Full of Hot (and Cold) Air

In the absence of a true treatment of the thermodynamics of the gas in each radial zone of the galaxy, we simplify the problem by assuming that there are two kinds of gas: a *cold* reservoir with mass $M_{\rm c}$ and temperature $T_{\rm c} \sim 10$K, and hot gas reservoir ($M_{\rm h}$) with temperature $T_{\rm h} \sim 10^7$K. The cold gas reservoir is primarily in-disc and able to form stars, whilst the hot gas models both gas outside the disc, and recently heated gas from supernova events.

Within a given gas reservoir, we assume that both the cold gas and the hot gas are individually well mixed, each with a different, uniform elemental abundance. The two reservoirs can interact through point heating events, and through a continual thermal exchange, which dominantly cools hot gas to cold, and hence is termed 'passive cooling'.

Point heating

The point heating events are associated with stellar phenomena: primarily supernovae events, and stellar feedback which occurs during bursts of star formation. Whenever a synthesis event (CCSN, SNIa etc.) returns an amount of gas $M_{\rm gas}$ to the reservoir, a fraction f_j of that gas is added to the hot reservoir, whilst the remainder is added to the cold reservoir. The fraction f_j varies depending on the event (i.e. $f_{\rm ccsn}, f_{\rm sn1a}$), as the differing properties of these events and their ejecta results in a different thermal profile.

It is common to parameterise stellar feedback through the 'mass loading factor', η. Whenever a mass M_* of cold gas is removed from the reservoir to form a new population of stars, a mass ηM_* of cold gas is heated into the hot phase

With $\rho_{\rm sfr}$ as defined in Eq. (2.4), the point-heating equations are therefore:

$$\left(\frac{{\rm d}M_{\rm c}}{{\rm d}t}\right)_{\text{point-heat}} = -\eta\rho_{\rm sfr}(t) + \sum_j (1 - f_{h,j})Y_j(t), \tag{2.7}$$

$$\left(\frac{{\rm d}M_{\rm h}}{{\rm d}t}\right)_{\text{point-heat}} = +\eta\rho_{\rm sfr}(t) + \sum_j f_{h,j}Y_j(t), \tag{2.8}$$

where $Y_j(t)$ is the amount of matter being returned to the galaxy by the point-heating event j, and the sum runs over all such point-heating processes.

Passive Cooling

The cooling of hot gas, however, is a much more complex topic, and still subject to significant uncertainty. The cooling happens primarily at the interface between the hot and the cold gas layers, where turbulent mixing creates fractal surfaces over which this cooling occurs (El-Badry et al., 2019; Fielding et al., 2020). This means that correctly modelling the

physics requires microscopic resolution across a galactic scale: clearly a very complex task, which remains under heated discussion.

To that end, we deliberately leave the cooling law exposed and subject to further investigation. The nominal functional form which we choose is based on a simple Newtonian Cooling law, expressed in terms of $\xi \equiv M_{\rm h}/(M_{\rm h} + M_{\rm c})$ the fraction of the gas which is hot:

$$\dot{\xi} = \lambda \left(\xi_\infty^n - \xi^n\right). \tag{2.9}$$

Here λ is the cooling rate, which we take to be ~ 1 Gyr^{-1} as a nominal value. ξ_∞ is the equilibrium hot fraction which the galaxy would otherwise support in the absence of point heating. The exponent n is 1 in the nominal Newtonian case, however Fielding et al. (2020), for example, finds that $n \approx -3/8$, however in this case the cooling was driven by interactions solely with the cold in-disc gas, whereas our model allows for interactions with the CGM as well. Given that $n < 0$ would drive a divergence in the cooling rate as $M_h \to 0$, we choose $n \approx 1$ as our nominal model.

Since cooling conserves the gas mass $M_{\rm c} + M_{\rm h}$, the passive cooling equations are:

$$\left(\frac{{\rm d}M_{\rm h}}{{\rm d}t}\right)_{\rm cool} = \lambda(M_{\rm c} + M_{\rm h}) \left(\xi_\infty^n - \left(\frac{M_{\rm h}}{M_{\rm c} + M_{\rm h}}\right)^n\right) \tag{2.10}$$

$$\left(\frac{{\rm d}M_{\rm c}}{{\rm d}t}\right)_{\rm cool} = -\left(\frac{{\rm d}M_{\rm h}}{{\rm d}t}\right)_{\rm cool} \tag{2.11}$$

We note that in the $n = 1$ case where $\xi_\infty \ll \xi \ll 1$, this reduces to a simple exponential cooling law:

$$\left(\frac{{\rm d}M_{\rm h}}{{\rm d}t}\right)_{\rm cool} \approx -\lambda M_{\rm h} \tag{2.12}$$

This is the cooling law used in SB09 for the original RAMICES. We also see that, although we termed this the 'cooling law', it is possible for this law to heat gas up in the case where $\xi < \xi_\infty$. In practice, this is only likely to occur in a multi-zoned model at the very outer regions of the galaxy where there is very little heating and large amounts of cold gas flowing in from the CGM.

2.3.4 Return of the Metals

With the basic framework of how stars are born and die within a galaxy, we must now encode how this reflects back into the chemical abundance of the gaseous components.

Given the variety of ways in which metals can be ejected into the ISM, it is convenient to talk about them in general terms - we will therefore use the language of 'processes', which should be understood to encompass both CCSN, ejecta from low mass stars (covered by the catchall term 'AGB death'), SNIa, NSM, ECSN and so on.

If we suppose that source process, j, is associated with a star of mass M and element content[3] $\mathbf{z}$ occurs at a rate $R_j(t, M, \mathbf{z})$ at the time t, then the total rate at which an element X is being returned from the stellar reservoir is:

$$y_{X,j}(t) = \int_0^\infty \mathrm{d}M \int_0^1 \mathrm{d}\mathbf{z}\ R_j(t, M, \mathbf{z}) Y_{j,X}(M, \mathbf{z}) \tag{2.13}$$

The development of suitable functions R_j in an analytical model is discussed in the context of SACEM in §3.3. In RAMICES II, R_j is indirectly computed for CCSN and AGB-death events through the tracking of stellar population lifetimes, and §4.3.3 considers how R_j is computed for SNIa and NSM events.

In this section, we focus on the gross yield function, $Y_{j,X}(M, \mathbf{z})$, the total ejected mass of material X which a star of initial mass M and abundance $\mathbf{z}$ returns to the ISM when it undergoes process j.

We first note that we have made an explicit effort to denote that $Y(M, \mathbf{z})$ is a function of the entire elemental content of the star forming gas, instead of just Z. This is important as, especially for low mass stars, the majority of the ejected material has undergone no nucleosynthesis and is simply the original metal content incorporated into the star at birth. The ejecta must therefore have a precise memory of the elemental distribution which formed it.

However, when stellar modellers produce nucleosynthetic yield grids, they do not produce a model for every possible combination of $\mathbf{z}$, it is instead assumed that stars follow a roughly solar-like elemental distribution which is then scaled linearly by the total metallicity Z. This is less than ideal for our purposes, as it is clear that a star which formed from gas which was 1% magnesium by mass will behave differently to one which was 1% plutonium, but the alternative is an evidently unfeasibly high dimensional yield grid.

We therefore dampen the impact of this by considering the *net* yields such a model produces, instead of the gross yield. The yield function therefore takes the form

$$Y_{j,X}(M, \mathbf{z}) = M \times \left(N_X(M, Z) + z_X\right), \tag{2.14}$$

where $N_X(M, Z)$ is the net change in the fractional abundance of element X predicted by the stellar models, and z_X is the fraction of the star forming gas which was element X (i.e. the X^{th} component of $\mathbf{z}$). We note that whilst $N < 0$ is perfectly possible (and in fact, will be the norm when $X =$ Hydrogen), we should be careful to ensure that $N + z_X > 0$, as we cannot destroy more of an element than originally existed in a star.

[3] In the most extreme case, $\mathbf{z}$ is a vector of arbitrary length detailing the precise elemental and isotopic composition of the star.

For SNIa and NSM events, we assume that $z_X = 0$, as, by definition, the progenitors of these events have been heavily altered from their birth abundances through the action of both degeneracy pressure and nuclear fusion such that no amount of the initial material is released. The yields of SNIa events are therefore gross, and taken from the W70 model of Iwamoto et al. (1999), the literature standard model of SNIa explosive nucleosynthesis. The NSM yields are highly unconstrained, and highly degenerate with $f_{\rm NSM}$, the fraction of neutron stars which are viable NSM candidates. As NSM are only tracked to account for Eu (our r-process tracer), we assume that $N_{\rm Eu,\ NSM} = 1$, i.e. the entire neutron star mass is converted to r-process material, and calibrate $f_{\rm NSM}$ by hand.

2.3.5 Onfalls, Outflows

The final key ingredient that we require is that the gas reservoirs be continuously supplied with a fresh supply of material. This onfall from the CGM is required in order for the galaxy to sustain observed SFRs: without this onflow, the gas reservoirs would have been rapidly depleted in the early galaxy, and the majority of the material locked up in the stellar phase (Chiosi, 1980). This cannot be the case for the Milky Way, as we can observe a sustained SFR throughout galactic history, with the current value observed to be in the region 1-2$M_\odot {\rm yr}^{-1}$ (Mor et al., 2019; Chomiuk & Povich, 2011; Robitaille & Whitney, 2010; Aumer & Binney, 2009).

A common prescription involves two periods of onflow, a rapid, early-time phase (thought to form the thick disc), and a slow-burn inflow (responsible for the thin-disc):

$$\dot{M}_{\rm onfall}(t) = \frac{M_1}{\tau_1} e^{-t/\tau_1} + \frac{M_2}{\tau_2} e^{-t/\tau_2}. \tag{2.15}$$

Let us, then, not make our division as we did before, with a view to all...for that has already brought upon us the proverbial penalty, of having made too much haste, resulting in too little speed

— Plato, Statesman

3

SACEM: An Analytical Model

Contents

When constructing a model for predicting the chemical evolution of the Milky Way (or indeed, any other galaxy), one must first decide the level of accuracy which one wishes to replicate. As we have seen in the previous sections, there are vast amounts of physics which must be accounted for in a 'perfect' model - but this necessarily requires simulation-heavy models and precludes the use of basic analytical tools.

We also note that, given the large degree of unconstrained parameters and observables required for a GCE model, a sufficiently flexible model is easily able to reproduce a small set of observational constraints. The ability of a model to replicate the data is not necessarily indicative of the model's physicality, as many divergent models are equally able to replicate the same portions of the data.

In order to draw conclusions about the physicality of a model, one must therefore look for where data is *not* reproduced, as in the case of a highly flexible model, this suggests the region is strongly forbidden, even under weak constraints. The model introduced in this chapter is therefore, by design, only weakly constrained and highly flexible, to allow for rapid searches of the parameter space, and hence the location of such 'excluded regions' of parameter space.

The resulting Simple Analytical Chemical Evolution Model (SACEM) is a single-zoned, thermally-phased analytical model of the chemical evolution of the galaxy. The model

was originally developed in order to search for sources of r-Process nucleosynthesis, and hence is the centrepiece of Chapter 5.

3.1 Galactic Mass & Star Formation

SACEM follows the general prescription for heating, cooling and star formation laid out in §2.3 in determining the rate at which cold gas is transformed into stars.

Starting with three mass reservoirs (M_c, the cold gas, M_h, the hot gas, and M_* the mass locked up in stars), we make a simplifying approximation to Eq. (2.4), such that the SFR is linear in the cold gas mass:

$$\rho_{\rm sfr} \propto M_c \tag{3.1}$$

This alternative form of the SFR is validated by observational work (Gao & Solomon, 2004; Krumholz et al., 2009; Heiderman et al., 2010; Kennicutt & Evans, 2012), and hence allows us determine the SFR in a linearised fashion, along with the feedback prescriptions and thermal fractionation presented in Eqs.(2.7)-(2.8), whilst the passive cooling law uses the simplified form of Eq. (2.12).

Instead of tracking stellar lifetimes via isochrones, we use a simple 'exponential death' model for stars returning their material back to the ISM[1], such that:

$$\dot{M}_{*,die} = \mu M_* \tag{3.2}$$

We initially assume the galaxy is composed only of cold gas ($M_c(t=0) = M_0$, $M_h(0) = M_*(0) = 0$). Subsequent infall from the CGM is parameterised by the exponential form of Eq. (2.15), where the parameters $\{M_i\}$ and $\{\beta_i = 1/\tau_i\}$ set the infalling mass and timescales respectively. Together, this produces the following coupled system of differential equations:

$$\begin{aligned}\dot{M}_c = &\sum_i M_i\beta_i \exp\left(-\beta_i t\right) + (1-f_h)\mu M_* \\ &+ \lambda_{\rm cool} M_h - (1+\delta)\nu_{\rm sfr} M_c\end{aligned} \tag{3.3}$$

$$\dot{M}_h = f_h \mu M_* + \delta \nu_{\rm sfr} M_c - \lambda_{\rm cool} M_h \tag{3.4}$$

$$M_{\rm t} \equiv M_c + M_h + M_* \tag{3.5}$$

$$\dot{M}_{\rm t} = \sum_i M_i \beta_i \exp\left(-\beta_i t\right) \iff M_{\rm t} = \bar{M} - \sum_i M_i e^{-\beta_i t} \tag{3.6}$$

This can be analytically solved for M_c, and hence the SFR $\rho_{\rm SFR}(t) = \nu_{\rm SFR} M_c(t)$. Because of the linearised assumptions we have made, the solution is expressible in terms of a sum of exponential terms.

[1]Note that this is solely for the purposes of the SFR and does not alter the chemical evolution. A more complex prescription could be used, but the final result would be equally replicated by altering $\nu_{\rm sfr}$ or μ, and therefore would add nothing to the model except complexity

3.2 Elemental Synthesis & Return of Metals

If a nucleosynthesis pathway j produces an amount $y_{j,X}(t)$ of element X at time t (Eq. 2.13), then the amount of X due to j present in the cold gas is given by M_{xcj}, equally the corresponding amount for the hot reservoir is given by M_{xhj}. They are linked via the thermal fractionation of the yield $y_{j,X}$ and the cooling and feedback mechanisms previously discussed:

$$\dot{M}_{xcj} = (1 - f_{h,j})y_{j,X}(t) + \lambda_j M_{xhj} - (1 + \delta)\frac{M_{xcj}}{M_c}\rho_{\rm sfr}(t) \tag{3.7}$$

$$\dot{M}_{xhj} = f_{h,j}y_{j,X}(t) - \lambda_j M_{xhj} + \delta\frac{M_{xcj}}{M_c}\rho_{\rm sfr}(t) \tag{3.8}$$

This reduces to:

$$\dot{M}_{xcj} = (1 - f_{h,j})y_{j,X}(t) + \lambda_j M_{xhj} - (1 + \delta)\nu_{\rm sfr}F_{\rm mod}M_{xcj} \tag{3.9}$$

$$\dot{M}_{xhj} = f_{h,j}y_{j,X}(t) - \lambda_j M_{xhj} + \delta\nu_{\rm sfr}F_{\rm mod}M_{xcj} \tag{3.10}$$

Here $F_{\rm mod}$ has been introduced as a 'lockup modification factor', such that the lockup rate is proportional to $F_{\rm mod}\rho_{\rm sfr}$, instead of just $\rho_{\rm sfr}$. This modification is introduced to allow for the fact that y is the rate at which new material is synthesised. Since stars are formed from polluted gas, as long as they do not destroy the material, they can release metals which they did not synthesise. If $F_{\rm mod} < 1$, therefore, we reduce the rate at which material is being locked up by mimicking the recycling of previously synthesised material. This is a (crude) approximation to the full form of Eq. (2.14).

3.3 Yield Functions

To compute the rate at which elements are synthesised via process j, we must compute the value of $y_{X,j}$ from Eq. (2.13), which is in turn a function of $R_j(t)$, the rate at which the event is occurring. We therefore invoke a DTD, $\Psi_j(t)$, giving the probability of a star undergoing stellar death a time t after it was formed[2]. The mass-rate of events j (i.e. the stellar mass loss rate through channel j) occurring at a time t is therefore given by:

$$R_j(t) = \int_0^t \rho_{\rm SFR}(t - t')\Psi_j(t'){\rm d}t' \tag{3.11}$$

Swapping the integration variable $t' \rightarrow \tau = t - t'$, it follows that the yield from event j is given by:

$$y_{X,j}(t) = \int_0^\infty \zeta(M){\rm d}M \int_0^t \rho_{\rm SFR}(\tau)\Psi(t - \tau, M, Z_{\rm cg}(\tau)) \times Y_{X,j}(M, Z_{\rm cg}(\tau)){\rm d}\tau \tag{3.12}$$

[2] $\Psi_j(t) = \Psi_j(M, Z, t)$ is explicitly a function of the progenitor mass and metallicity in this formulation, but this is omitted from the notation for convenience

Where $\zeta(M)$ is the IMF, $Y_{X,j}(M,Z)$ is the gross yield of X from a star of mass M and initial metallicity Z dying through process j, and $Z_{\rm cg}(t)$ is the cold-gas metallicity at a time t. With Eq. (3.12) in hand, we are able to derive three equations for the cases of CCSN events and delayed/inspiral events such as SNIa and NSM as follows:

3.3.1 Core Collapse Supernovae

Eq. (3.12) simplifies for the case of CCSN from high-mass progenitors. Such CCSN occur at the end of a lifetime T, such that the DTD becomes a Dirac delta function: $\Psi(t) = \delta(t-T)$. In addition, for stars massive enough to go CCSN, this lifetime is short compared to the timescale over which galactic properties change, such that we may approximate $T \approx 0$:

$$y_{X,\rm CCSN}(t) = \rho_{\rm SFR}(t) \int_0^\infty \zeta(M) Y(M, Z(t)) {\rm d}M \tag{3.13}$$

In the case where the yields are independent of the metallicity (which, by comparison with yield grids such as Chieffi & Limongi (2004); Limongi & Chieffi (2018), we see holds to a good approximation for both Fe and Mg), we may compute the integral over M to find the characteristic yield function of X via j, $\bar{Y}_{X,j}$, giving the synthesis rate via CCSN as:

$$\begin{aligned} y_{X,\rm CCSN}(t) &= \bar{Y}_{X,j} R_{\rm CCSN}(t) \\ &= \bar{Y}_{X,j} \times \rho_{\rm SFR}(t) \end{aligned} \tag{3.14}$$

3.3.2 Delayed Yields

Finally, we consider the yields of SNIa and NSM events. These events do not occur uniquely at the end of a stellar lifetime. Instead, after the stellar lifetime has passed, there exists a period of probabilistic decay, whilst the system continues to evolve until finally the progenitors inspiral (for double degenerate SNIa and NSM events), or accrete enough matter from their companion (for single degenerate SNIa). In addition, there exists a non-trivial time delay before the first events can start occurring.

Whilst the common DTD for SNIa is typically given as $\propto t^{-1}$, in order to continue our ability to easily analytically integrate them, we follow the work of SB09in using an exponential DTD for SNIa events - for a discussion of the validity of this approach, see Weinberg et al. (2017). With $\Theta(x)$ the usual Heaviside step function, we find that the DTD becomes

$$\Psi_j(t) \propto \Theta(t-\tau_j) \exp\left(-\nu_j t\right). \tag{3.15}$$

Assuming metallicity independence, we can express this as

$$\begin{aligned} y_{X,\mathrm{delay}}(t,\nu,\tau) &= \bar{Y}_{X,j}\mathcal{I}\left[t,\rho_{\mathrm{SFR}},\nu,\tau\right] \\ &= \bar{Y}_{X,j}\Theta(t-\tau)\int_{\tau}^{t}\rho_{\mathrm{SFR}}(t-x)\exp\left(-\nu x\right)\mathrm{d}x. \end{aligned} \tag{3.16}$$

The literature standard for SNIa yields is the (metallicity-independent) Iwamoto et al. (1999) W70 model. More recent efforts such as have attempted to account for progenitor metallicity, Travaglio et al. (2005); Gronow et al. (2021), and find that Fe production is altered by less than 6 per cent from the W70 model between $0.1\mathcal{Z}_{\odot}$ and $Z_{\odot}$, such that we consider the metallicity-independent model a good approximation.

3.4 Calibrating Yields

The yield functions described in §3.3 contain an undetermined multiplicative factor, $Y_{X,j}$, the effective yield of element X through process j. This factor encapsulates both the absolute yield from the star, the weighting of the IMF, as well as effects such as galactic ejection and remnant lockup.

A true GCE model (such as the RAMICES II code developed in Chapter 4) would try to derive these prefactors from first principles. For the purposes of SACEM, we instead fix these values by requiring that the curves of interest to us replicate some chosen calibration values. The chosen calibration points, and their values in the nominal model, are determined by the purpose of the model. In Chapter 5, we fix the end-point abundances of [Eu/Fe], [Mg/Fe] and [Fe/H], and the contribution fraction of NSM and Collapsar events to the final Eu budget, in order to study the early-time abundances of [Eu/Fe].

We note again that this strategy allows additional degrees of freedom into our model, but that since we are investigating which regions of parameter space are excluded, this in fact strengthens any conclusions we might draw: additional constraints on $\bar{Y}_{Xj}$ would serve to reduce the viable parameter space, not expand it.

3.5 SACEM: A Recipe

We are now able to set out a procedure to derive a chemical history as a function of the parameters of a given SACEM model. A full list of parameters is given in 5.2.

1. Solve differential equations Eq. (3.3)-(3.6) to generate a SFR:

$$\rho_{\mathrm{SFR}}(t) = \rho(t|M_0, M_1, M_2, \beta_1, \beta_2, \nu_{\mathrm{sfr}}, \mu, \delta) \tag{3.17}$$

2. Use Eq. (3.14) and Eq. (3.16) to generate event rates for each process:

$$R_{\rm ccsn}(t, \rho_{\rm sfr}) = \rho_{\rm sfr}(t) \tag{3.18}$$

$$R_{\rm k}(t, \rho_{\rm sfr}, \nu_{\rm k}, \tau_{\rm k}) = \mathcal{I}\left[\rho_{\rm sfr}, \nu_{\rm k}, \tau_{\rm k}\right] \quad k \in \{\rm snIa, nsm\} \tag{3.19}$$

3. Solve Eq. (3.9)-(3.10) to find the (uncalibrated) cold-gas mass in the disc produced by each process as a function of the relevant yield, thermal and lockup parameters.

$$M_k = M\left[t, R_k, f_{h,k}, \lambda_k, \delta, \nu_{\rm SFR}, F_{\rm mod}\right] \tag{3.20}$$

4. For each element under consideration, construct models for the mass of this element within the cold gas as a function of unknown prefactors:

$$M_{\rm X}(t) = \sum_k \alpha_{X,k} M_k(t) \tag{3.21}$$

5. Determine the values of $\alpha_{X,k}$ by analytically solving for the observational constraints.

This method generates analytical functions which can be used to plot Tinsley diagrams for the chosen elements, and forms the core of SACEM.

3.6 The Action of Delays & Thermal Phasing

Central to understanding our chemical evolution models are the differing DTDs of yields, and delays of injection to and freeze-out from the hot gas phase.

Events such as CCSN occur almost immediately as the lifetime of high-mass stars is negligible (≈10Myr) compared to the time over which the galactic properties change. SNIa and NSM, however, rely on the death of a previous population of stars for their formation mechanism, giving a minimum timescale of about 200Myr for SNIa and 50Myr for NSM.

In addition to the absolute rate of synthesis, Schönrich & Weinberg (2019) showed that the rate at which polluted gas becomes available for star formation also has a large impact on GCE models of r-process synthesis, because we expect different timescales for CCSN, NSM and SNIa gas-availability. For example: the majority of CCSN occur close to regions of high star formation (and hence close to the feedback 'chimneys' of Norman & Ikeuchi 1989), so we expect a large contribution of the gas to be placed in the hot reservoir, or temporarily ejected from the galaxy, only to 'fountain' back in, as per (Shapiro & Field, 1976). NSM, however, have a non-trivial time delay, and also experience natal kicks, so can be expected to be far removed both spatially and temporally from feedback chimneys, and their synthesis of

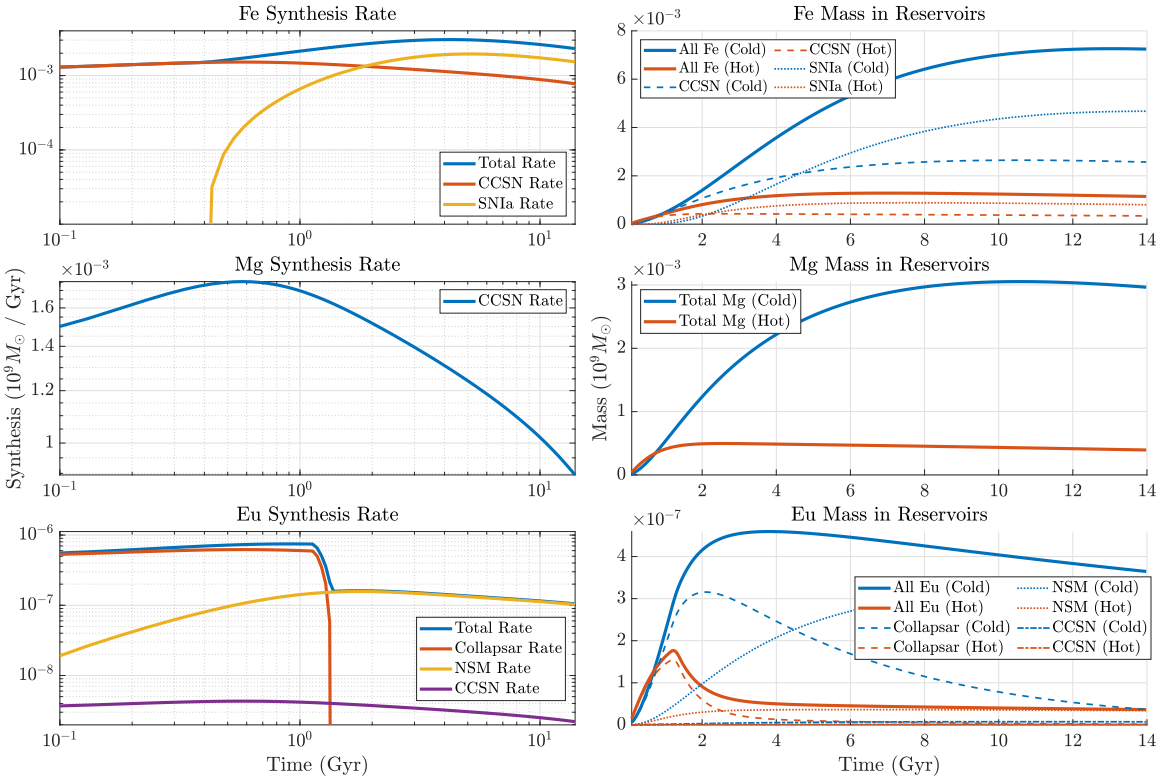

Figure 3.1: *Left:* The rate of synthesis of the three species within SACEM and *right:* the resulting mass of each element stored in the hot and cold gas phases for a calibrated SACEM instantiation with a collapsar cutoff time $\tau_{\rm coll} = 1.35$Gyr. The Eu contributions (χ, σ, κ) were selected to amplify the signal of the differing behaviours, rather than generate a viable chemical history.

high-mass, neutron rich material would lead to strong line cooling of the ejecta, making the material available for star formation much more rapidly. Microphysics such as dust-formation may also have an impact on the availability of gas for star formation.

It was for this reason that §2.3.3 introduced the 'hot gas injection fractions', in which a portion f_j of the material synthesised by process j into the hot gas reservoir, with the remainder going into the hot-gas reservoir [3]. Following Schönrich & Weinberg (2019), we adopt $(f_{\rm CCSN}, f_{\rm Coll}, f_{\rm NSM}, f_{\rm SNIa}) = (0.75, 0.75, 0.4, 0.99)$ and $\lambda_j = \lambda_{\rm CCSN} = 1$Gyr as our nominal model.

Figure 3.1 shows how these differing temporal and thermal processes lead to differing amounts of material found within the various reservoirs over time. The left hand panels show the (calibrated) rates at which the three tracked metals are synthesised within a given SACEM initialisation, whilst the right hand panels show the total mass of material in the hot and cold gas reservoirs at each time. The version of SACEM used here is the

[3] SACEM does not include a total-ejection fraction as this takes the form of a multiplicative prefactor $y_{\text{non-eject}} = (1 - f_{\rm eject}) y_{\rm X}$, as $y_{\rm X}$ is calibrated, we can omit the factor and calibrate the value of $f_{\rm eject}$ simultaneously with the absolute yield.

one described in Chapter 5, and hence tracks Fe, Mg and Eu synthesis, with an additional 'collapsar' channel for Eu. The collapsar channel operates identically to CCSN, with the exception that it 'turns off' at a prescribed time.

If I have seen further it is by standing on the shoulders of Giants.

— Isaac Newton

4

RAMICES II: A Simulational Model

Contents

Much of the work presented in this thesis is expanding on the discoveries made in Schönrich & Binney (2009a) (SB09) and the follow up works Bilitewski & Schönrich (2012), Schönrich & McMillan (2017), Schönrich & Weinberg (2019). These works were centred around a simulation developed for SB09: the Radial Migration Chemical Evolution Simulation (RAMICES).

RAMICES is a spatially resolved, thermally-phased simulational model of GCE. The chemodynamics portion of the code is capable of highly powerful synthesis of the full action-angle distribution of the stellar component and accurately synthesises a stellar catalogue of candidates meeting the selection criteria of the Geneva-Copenhagen Survey (GCS, Nordström et al. 2004) or the Apache Point Observatory Galactic Evolution Experiment (APOGEE, Ahumada et al. 2020). The code, however, suffers from a number of difficulties in utilisation, as a result of the iterative, organic growth of the solutions, and foundational design choices which make subsequent modifications to the internal physics difficult without disrupting the remainder of the model

In order to advance the field further, we therefore produced a complete redesign of the software, RAMICES II, incorporating new physics and processes, and exposing some of the

inner workings to the user, such that the space of possible models can be explored more fully[1].

4.1 Aims for this model

Whilst RAMICES II is heavily inspired by its predecessor, and follows many of the same prescriptions, it includes several improvements, both from the point of view of increasing the physics which governs the evolution, and from the way in which the code can be modified and altered by the user to study different physical scenarios.

The primary aim for RAMICES II is that it be a publicly available codebase which can be used and modified by anybody wishing to make a useful study of GCE. Whilst this does not particularly impact the physics or processes which will be incorporated into the code, it does require a significant dedication of resources to the design of the software and the user interface. In particular, RAMICES II is highly modular in its design, allowing for new processes to be easily incorporated into the existing framework.

This design intent also manifests itself in the open-ended way some of the models are expressed – for example, the cooling law Eq. (2.10) is parameterised in a very open-ended way, specifically for this purpose.

We also aim to make it easier to study for example, how the choices regarding the CGM infall impact the resulting chemical evolution, and study how additional processes, such as ECSN impact galactic chemistry.

In its current iteration, we make no attempt for RAMICES II to accurately reproduce the action-based dynamics in the way that RAMICES does. Whilst we hope to incorporate such an optional feature at a later date, the lessons learned from RAMICES indicate that this was a highly-resource expensive component of the code ($> 90\%$ of runtime), and had a negligible impact on the distributions in chemical space.

4.2 Basic Outline & Disc Structure

We model the galaxy as a series of N_r concentric annuli, each of width ΔR_i. An annulus is labelled by the central radius, R_i, and hence extends between $R_i \pm \Delta R_i/2$. The model explicitly supports considerations of non-uniform ring width, and hence allows more detailed studies of denser regions such as the galactic centre, though we do not use this feature in our nominal model.

[1]Although not the focus of the work here, we also note that we were strongly motivated by advances in both available hardware and improvements in the C++ design standard (ISO, 2020), allowing for significant conceptual and performance improvements.

Each annulus contains a gas reservoir with both hot and cold components, and a stellar reservoir tracking both living stars, and the remnants of previous generations. The gas reservoir differentiates gas by three sources: accreted/primordial, stellar and remnant material. Material which is ejected from stars in CCSN or AGB events is added into the stellar pool, whilst material in the remnant category originated in SNIa or NSM events. Material which originated in either stellar or remnant events and is subsequently ejected into the CGM (§4.4.4) and then accreted back into the galaxy through onflows is 'converted' to the accretion pool. In this fashion we are able to easily determine the dominant factor for the origin of chemicals in each annulus - particularly useful for identifying the annuli where the onflow prescription is the dominant driver of the abundance patterns.

A given RAMICES II model begins with a pristine galaxy of mass M_0, and evolves it over the course of a number of timesteps of width Δt until the prescribed final time $t_{\rm end}$.

4.3 Stars & their Corpses

4.3.1 Stellar Populations

Following §2.3.2, we take the SFR to be the Kennicutt-Schmidt formulation, with prefactor $\kappa_{\rm schmidt}$ and cutoff density $\Sigma_{\rm cut}$ chosen such that the model recovers the stellar surface density in the solar neighbourhood. Figure 4.1 shows the evolution of our calibrated model, compared to the surface densities of the various galactic components reported by McKee et al. (2015), whilst Figure 4.2 shows the recovered SFR, compared to literature values of the star forming history of the Milky Way. We find good agreement with these observations when using the parameters of Table 4.2.

We allow stars to form with masses in the range $[0.1, 100]M_\odot$, though we make the simplifying choice that a number of stars can be considered 'immortal', i.e. they have lifetimes considerably in excess of the simulation duration at all possible metallicities, which is true for stars with $M < 0.7M_\odot$. Tracking these stars in individual mass bins would serve no purpose for the majority of the simulation - we simply treat these low-mass stars as a low mass monolith. The distributions are reconstructed during the synthesis of observations, but otherwise we treat these 'immortal' stars as inert mass particles.

Whilst stars are allowed to form down to $0.1M_\odot$ (and the IMF normalised accordingly), we therefore only track stars in the region $[0.7, 100]M_\odot$. We split this mass range into N_m bins, which allows us to track the number of stars in each bin, using the Chabrier (2003) IMF and Eq. (2.6). As noted previously, we allow stars to form in non-integer numbers, such

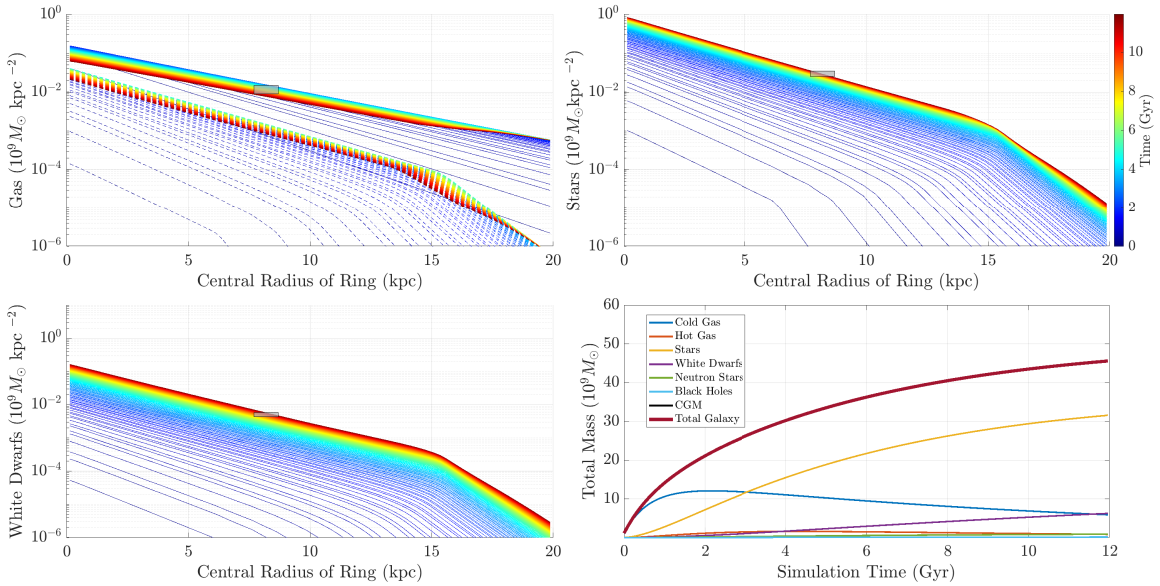

Figure 4.1: The evolution of the surface density of gas (*top left, cold gas solid lines, hot gas dashed lines*), stars (*top right*) and white dwarfs (*bottom left*), as well as the total mass budget (*bottom right*). The gray shaded boxes depict the measured surface densities of each component, as reported by McKee et al. (2015).

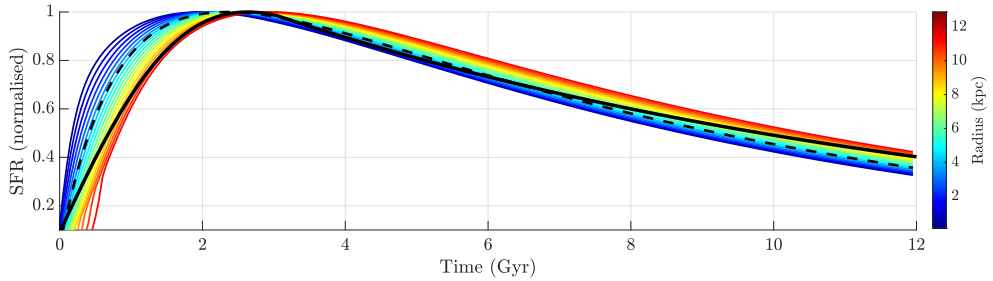

Figure 4.2: The evolution of the SFR, normalised to the maximum value at each annulus. The dashed line shows the (normalised) galactic SFR, whilst the solid black line shows the SFR of Just & Jahreiß (2007) for $t < 3$Gyr and Aumer & Binney (2009) for $t > 3$Gyr, find close agreement with our values for the solar annulus.

that each star represents a *probability* of a star, rather than attempting to directly simulate a galaxy. The only time that this probability is resolved into an integer number of stars is during the synthesis of the stellar catalogue, discussed in §4.7.

The N_m mass-bins are uniform in $\log(M)$, allowing us a higher mass resolution at lower masses. This important particularly when considering the lifetimes of stars relative to the temporal resolution of the simulation: stars greater than $\sim 25 M_\odot$ will live only a single timestep, and hence there is little to be gained from tracking these stars at high resolution, whilst stars with masses $\sim M_\odot$ have lifetimes which differ by many dozens of timesteps.

4.3.2 Modified Feedback

Due to the presence of Σ_h in the star formation law Eq. (2.3.2), we must be careful that our feedback prescription (intended to act as a self-limiting factor on star formation), does not lead to runaway growth, or unphysical star formation efficiencies, $\eta = \rho_{\rm sfr}/M_c$.

In the prescription laid out in §2.3.3, feedback heats gas from the cold gas reservoir into the hot reservoir. Since this leaves $\Sigma_c + \Sigma_h$ unchanged, but reduces Σ_c, it therefore drives up the star formation efficiency drastically. We therefore modify this previous prescription, and allow a fraction $e_{\rm feedback}$ of the feedback-heated gas to be ejected into the CGM, instead of the hot gas phase.

4.3.3 Stellar Remnants

Stellar death rarely completely disrupts the star, as shown in Table 2.1. Most remnants are dormant, in the sense that they only need to be tracked for the purpose of ensuring mass conservation. This is true for Black Holes and White Dwarfs from progenitors less than $\approx 3.5 M_\odot$.

However, neutron stars and COWDs in the appropriate mass range are important contributors to galactic chemistry, and must therefore be tracked in detail in order for us to model SNIa and NSM events.

Truly modelling SNIa and NSM events would require a vastly more sophisticated simulation than we have here, capable of resolving binary star systems and the associated mass transfer, gravitational wave emissions and so on. We therefore make the simplifying assumption that, for both NSM and SNIa a constant fraction of remnants in the appropriate mass range (which is all neutron stars, and COWDs from progenitors $> 3.2 M_\odot$) are candidates for SNIa and NSM progenitors, $c_{\rm nsm}$ and $c_{\rm snIa}$.

These candidates are then placed into a buffer, where a set length of time must pass, hence modelling the binary evolution and inspiralling which must occur. This time delay is chosen to be $\tau_{\rm nsm} \approx 0.05$Gyr and $\tau_{\rm snia} \approx 0.5$Gyr. Once this time has passed, the candidates become 'active', and can undergo their synthesis events.

Each active entry in the buffer then probabilistically decays given a frequency $\nu_{\rm snia}$ or $\nu_{\rm nsm}$, such that a fraction $(1 - \exp(-\nu \Delta t))$ of candidates undergo synthesis events at each timestep.

We note that, unlike RAMICES, we incorporate our remnants into the scattering framework (§4.5). This is particularly important given the long temporal delay between remnant formation and the return of metals. The ejecta of these events should be significantly more dispersed than those of the CCSN which formed them.

4.4 Onflows & Inflows

In §2.3.5, we showed that GCE models require a constant inflow of gas in order to support sufficient star formation to late times. The total amount of this onfall is commonly parameterised through the two-infall model of Eq. (2.15). In addition, there is observational evidence of gas within the galaxy flowing inwards in a 'galactic wind', which acts to sweep synthesised materials into the centre of the galaxy.

4.4.1 Exponential Enforcement

The requirement that we match the simulated SFR to the observable star forming history of the galaxy allows us to fine tune the amount of mass in the galaxy relatively precisely, and hence determine the parameters in Eq. (2.15). A more difficult question, however, is *where* in the galaxy the gas ends up. Intuitively, we expect the majority of the inflowing gas to enter the galaxy at the outer radii (this follows from the inferred alignment of the galaxy angular momentum and that of the gaseous filament from which it formed), however there is also evidence of in-disc interactions with the circum-galactic medium. The radial distribution of onfalling matter is therefore non-trivial and unconstrained.

We again follow SB09 in making the simplifying assumption that the cold gaseous disc had an approximate exponential surface density profile at all times, with only the total mass, $M_{\rm cg}$ and exponential scale length, $R_{\rm gas}$ changing with time:

$$\Sigma_{\rm cg}(R,t) \equiv \frac{M_{\rm cg}(t)}{2\pi R_{\rm gas}(t)^2} \exp\left(-\frac{R}{R_{\rm gas}(t)}\right) \tag{4.1}$$

Inside-Out Formation

We adopt the inside-out formation of Schönrich & McMillan (2017), such that the exponential scale length of the gas evolves according to

$$R_{\rm gas}(t) = R_0 + \mathcal{N}\left(\text{atan}\left(\frac{t-t_0}{t_g}\right) - \text{atan}\left(\frac{t_0}{t_g}\right)\right), \tag{4.2}$$

where $\mathcal{N} = \mathcal{N}(R_0, R_{\rm final}, t_{\rm final}, t_0, t_g)$ is a normalisation constant such that $R_{\rm gas}(t_{\rm final}) = R_{\rm final}$. Eq. (4.2) was chosen to be a smooth, infinitely differentiable function to prevent artefacts arising from discontinuities. The choice of parameters is shown in Table 4.1, and was chosen to provide a fast linear growth whilst $t \sim t_0$, before transitioning to a slow growth of the scale length, in line with observational estimates.

Table 4.1: Inside-Out Growth parameters, derived from Schönrich & McMillan (2017).

Parameter	Interpretation	Value
R_0	Initial gas scale length	1.5 kpc
t_{final}	End of growth period	12Gyr
R_{final}	Gas scale length at $t = t_{\text{final}}$	3.5 kpc
t_0	Delay time before linear transition	1Gyr
t_g	Growth timescale	2Gyr

4.4.2 Inflow versus Onflow

At each timestep, Eq. (2.15) prescribes an amount ΔM of new, cold gas to be added to the disc. Each annulus then uses the exponential profile Eq. (4.1) to infer what its 'new mass', $M_{i,\text{ new}}$ should be.

There are, however, two ways in which this requirement can be met: either the galactic wind can transfer matter from a higher radius (inflow), or the material can be directly accreted from the CGM (onflow). The outermost annulus has no 'higher radius', and hence its entire budget is met through accretion from the CGM.

Bilitewski & Schönrich (2012) demonstrated that angular momentum considerations could be used to constrain the ratio $\Delta M_{\text{inflow}}/\Delta M_{\text{onflow}}$. We adapt their argument for the case of non-uniform spatial grids and find that for the annulus at radius R_i

$$B_i = \frac{\Delta M_{\text{inflow}}}{\Delta M_{\text{onflow}}} = \frac{2}{R_i\left(\Delta R_i + R_{i+1}\right)}\left(\frac{a}{4R_{\max}}\left(R_{i+}^4 - R_{i-}^4\right) + \frac{(b-1)}{3}\left(R_{i+}^3 - R_{i-}^3\right)\right) \tag{4.3}$$

Where $R_{\max}$ is the radial extent of the disc and $R_{i\pm} = R_i \pm \frac{\Delta R_i}{2}$ are the inner and outer radius of the annulus. The parameters a and b control the circular velocity profile of the CGM, in terms of the galactic circular velocity V_c, such that:

$$v_{\text{cgm}}(R) = V_c\left(b + a\frac{R}{R_{\max}}\right) \tag{4.4}$$

The prescription for determining how each annulus accretes matter, assuming the set $\{M_{i,\text{new}}\}$ is known, is as follows:

1. Start from the innermost annulus
2. Compute the additional mass required to increase the mass to $M_{i,\text{new}}$, accounting for any mass lost in previous iterations of this loop, to find ΔM
3. Compute Eq. (4.3) (unless at final annulus), and hence find the inflowing matter $\Delta M_{\text{inflow}} = \frac{B_i}{1+B_i}\Delta M$

4. Take $\Delta M_{\rm inflow}$ from the annulus $i+1$, decreasing the mass of that annulus (see point 2).

5. Take the remainder of the mass, $\Delta M_{\rm onflow} = \Delta M - \Delta M_{\rm inflow}$ from the CGM

6. Iterate outwards through the galaxy

We note that whilst the parameters a and b controls the mean windspeed of the inflowing material[2], this prescription has a finite maximum speed at which a tracer propagates through the galaxy: one ring-width per timestep. In the case of our nominal model ($\Delta t = 0.02$Gyr, $\Delta R = 0.2$kpc), this corresponds to a windspeed of 10km/s, which is significantly above the expected windspeed for the majority of galactic history, and hence is of little concern.

4.4.3 Edge Cases

The early galaxy, presents two main edge-cases for the prescriptions laid out above. We first note that in the early, rapid accretion phase the wind speeds can begin to reach the limiting speed, which manifests through the required inflowing mass being greater than the gas available in the next annulus, $\Delta M_{\rm inflow} > M_{i+1,\rm gas}$, in such a case, we take the mass available, and make up the remainder from the CGM. This is equivalent to increasing the angular momentum of the onfalling gas, and can be prevented by using a finer temporal grid.

Care must also be taken when computing $M_{i,\rm new}$ during phases of rapid inside-out formation. In the naïve case, it is entirely possible for the redistribution of mass implied by the altering of $R_{\rm gas}$ to lead to a mass overbudget in the inner annuli: i.e. $M_{i,\rm new} < M_i$ - the annulus is trying to lose matter to higher radii by means of an outward flowing wind. Considering the assembly history of the galaxy, we consider such outward-oriented winds to be unphysical. Should such a scenario arise, we instead keep the gas mass in the annulus constant and increase the onfall in the outer galaxy, sacrificing the exponential profile of the gas disc in favour of smooth accretion function. We note that the prescription for re-balancing the matter in RAMICES II fixes an error present within the RAMICES prescription which caused infinitely divergent accretion when this condition was met.

4.4.4 Circum-Galactic Medium

Just as with the radial distribution of the onflowing matter, the exact chemical abundance of this onfalling material is unconstrained. We expect this gas to be low metallicity as it is predominantly extra-galactic gas which has not been enriched as much as the in-disc

[2]This is evident from the extreme case $(a, b) = (0, 1)$, in which $v_{\rm cgm} = V_c$, and hence $B_i = 0$, such that the wind speed is zero

gas. However, quasar absorption lines (Pettini et al., 2003) show the presence of heavy elements in the CGM even in the early universe, which indicates that there are feedback mechanisms for enriched gas to enter into the onfalling material. This is also evidenced by the 'Galactic Fountain' of Shapiro & Field (1976), in which supernova ejecta from the inner galaxy 'fountains' back onto the outer galaxy, enriching it.

We therefore model these complex processes by considering the matter ejected from the galaxy in supernova events. That a significant portion of the material synthesised in supernovae must be totally ejected from the galaxy follows from simple energetic considerations. We approximate that if the supernova return rate is $y_{X,j}(t)$ (Eq. 2.13), the amount of material being ejected into the CGM is:

$$r_{\text{eject}}(t) = \ell y_{X,j}(t) \tag{4.5}$$

We assume that this removes mass before the point heating (§2.3.3) is applied, and is unbiased in terms of which metals are ejected from the galaxy. Therefore the mass of X which is injected into the each phase of the annulus is:

$$i_{\text{cold, X},j} = (1 - f_j)(1 - \ell) y_{X,j}(t) \tag{4.6}$$

$$i_{\text{hot, X},j} = f_j(1 - \ell) y_{X,j}(t) \tag{4.7}$$

It is common in models for the 'ejected' mass to be truly ejected: the material is simply thrown away. Within RAMICES II, we instead use this material to physically motivate the abundance patterns of our onfalling gas, we hold another gas reservoir around our galaxy to mimic the CGM. 'Ejected' gas is incorporated into the CGM and allowed to mix and cool. Whenever the onflow prescription requires new gas be added, it is taken from this reservoir.

The CGM is initially assumed to be formed from cold gas of mass M_{CGM} in primordial elemental ratios. The ability for the galaxy to pollute the CGM, and hence the metal content of accreted matter, is controlled by M_{CGM}. The case $M_{\text{CGM}} \to \infty$ mimics the accretion of pure primordial across galactic history, whilst $M_{\text{CGM}} \approx M_{\text{galaxy}}$ ensures that, at late times, the entire CGM is composed of material ejected during supernova events and is hence very metal rich.

4.5 Radial Migration

In addition to the flow of gas through the disc, the presence of structures such as spiral arms and the central bar are expected to cause stars and gas to exchange angular momentum and hence change the radius of their orbit.

We follow SB09 in suggesting that, at each infinitesimal timestep, a test particle in ring i has a probability p of either moving to a neighbouring ring $j = i \pm 1$, or remaining in its current ring. This probability is independent of the previous history of the test particle, and hence this mechanism is a Continuous Time Markov Chain. We therefore introduce an infinitesimal scattering matrix S such that a given test particle distribution $\mathbf{m} = (m_{\text{ring }1}, m_{\text{ring }2} \ldots)$ is distributed after the scatter by:

$$\mathbf{m}(t + \mathrm{d}t) = S\mathbf{m}(t) \tag{4.8}$$

SB09 used a simple dimensional argument to suggest that the probability of a test particle flowing into a region should be proportional to the mass, M_j, of the absorbing region. Modifying their argument to account for non-uniform spatial grids and robustness against changes in the spatial resolution, we find that the probability should be proportional to $\frac{M_j}{\Delta R_i \Delta R_j}$, with $\Delta R_{i/j}$ the width of the donating/absorbing annulus. We also note that, although the test particle changes angular momentum, the global angular momentum must be preserved throughout the scattering, as well as the specific angular momentum at each annulus. This is achieved if the mass flowing into a region is equal to that flowing out: the mass vector $\mathbf{M} = (M_1, M_2, \ldots)$ must therefore be an eigenvector of the infinitesimal scattering operation, $\mathbf{M} = S\mathbf{M}$. Together, this implies that the infinitesimal scattering matrix must take the form:

$$S_{ij} = p(j \to i) = \begin{cases} 1 - p(j \to i+1) - p(j \to i-1) & i = j \\ \frac{\mu \mathrm{d}t}{\Delta R_i \Delta R_j} \frac{M_i}{M_{\max}} & j = i \pm 1 \\ 0 & \text{else} \end{cases} \tag{4.9}$$

Here we have introduced $M_{\max}$ (the maximum mass within the set of rings) as a normalisation parameter and μ is a Markov-mixing parameter with units $\mathrm{kpc}^2\mathrm{Gyr}^{-1}$. The choice of $M_{\max}$ is somewhat arbitrary - equally valid choices would be M_{tot}, the total mass in the galaxy, or even omitting this parameter entirely, and letting μ gain mass units. Although these do imply differing scattering histories, numerical experimentation shows the effect to be negligible, and so we choose to use $M_{\max}$ as it implies that the maximum transfer probability at every timestep is $\mathcal{O}(\frac{\mu \Delta t}{\Delta R^2})$, and hence we are easily able to estimate the required expansion order for §4.5.1. We note that $M_0 = M_{N+1} = 0$ (i.e. there is no mass flowing into the centre or over the very edge of the galaxy during a scattering operation). By construction, S is a stochastic matrix which is a function only of the global mass distribution (which is constant during the scattering), it therefore follows that, in the limit of $\mathrm{d}t \to 0$:

$$\dot{\mathbf{m}}(t) = \lim_{\mathrm{d}t \to 0} \left(\frac{S - I}{\mathrm{d}t} \right) \mathbf{m} \equiv K(\mathbf{M})\mathbf{m}(t) \tag{4.10}$$

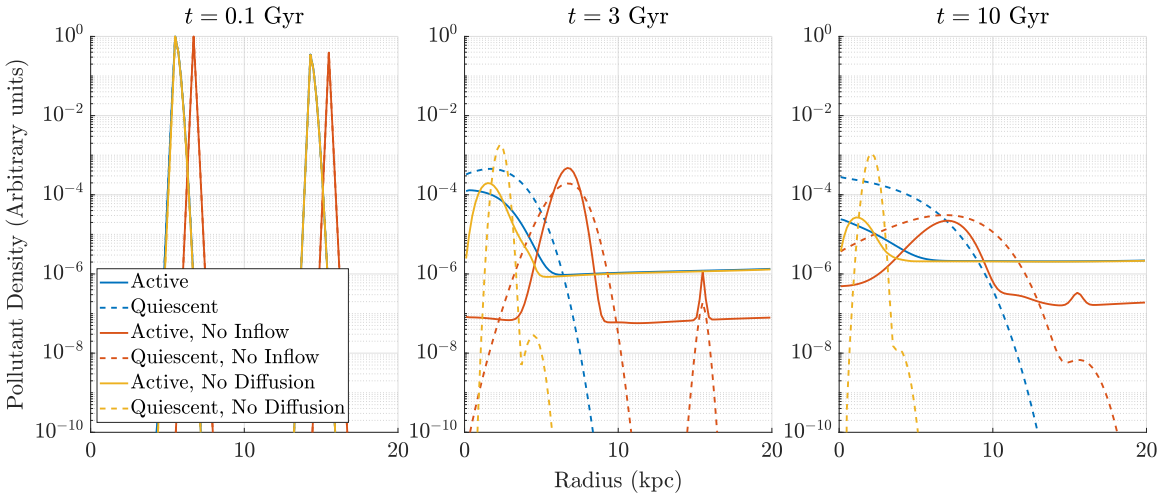

Figure 4.3: The evolution of the linear abundance density $\frac{m_{\text{pollutant}}}{m_{\text{hydrogen}}}$ after 'spiking' the model with a tracer pollutant at early times for a variety of different model parameters. Quiescent models have no star formation, whilst active models form stars and feed enriched material into the CGM for subsequent fallback. In 'No Inflow' models, all material moving into the galaxy is directly supplied by the CGM (§4.4.2). In 'No Diffusion' models, the Markov mixing strength is set to 0Gyr^{-1}, otherwise it is set to $0.5\text{kpc}^2\text{Gyr}^{-1}$.

Where

$$K_{ij} = \begin{cases} -\frac{\mu}{\Delta R_i M_{\max}} \left(\frac{M_{i+1}}{\Delta R_{i+1}} + \frac{M_{i-1}}{\Delta R_{i-1}} \right) & i = j, \\ \frac{\mu}{\Delta R_i \Delta R_j M_{\max}} M_i & j = i \pm 1, \\ 0 & \text{else.} \end{cases} \tag{4.11}$$

Therefore, over the duration of one finite timestep Δt, the scattering occurs according to:

$$\mathbf{m}(t + \Delta t) = \exp\Big(\Delta t \times K(\mathbf{M})\Big)\mathbf{m}(t) \tag{4.12}$$

In Figure 4.3 we see the effects of this dispersion relationship combined with various other mixing and transport processes, including the galactic fountain induced by stellar feedback and supernova ejection. We see that although the effects of the CGM interactions dominate the mixing processes, especially at larger radii, the inclusion of the mixing is by no means trivial, and it is therefore vital to model this effect.

4.5.1 Truncation

Especially at small Δt, it would be wasteful to fully compute the matrix exponential of Eq. (4.12), as the majority of the terms will be incredibly small. We therefore suggest simply expanding the exponential in polynomial terms of K, up to a specified order N_{expand}. We note that the properties of stochastic matrices guarantees that the conservation properties

of S are also obeyed by $\exp(aK)$, as well as all finite power law expansions thereof:

$$\mathbf{m}(t+\Delta t) = \left(\sum_{n=0}^{N_{\text{expand}}} \frac{(\Delta t K)^n}{n!} \right) \mathbf{m}(t) \tag{4.13}$$

The properties of K means that such a truncation will prohibit ring j from exchanging material with rings further away than $j \pm N_{\text{expand}}$ in a single timestep. The scattering method used in SB09 effectively fixed $N_{\text{expand}} = 2$ as this limitation on the diffusion speed was considered desirable.

The value N_{expand} can be chosen by the user at runtime in order to increase or decrease the accuracy (and, since this lies in the performance critical section of the code, correspondingly increase or decrease the runtime) of the scattering process, though we note from numerical experimentation that $N_{\text{expand}} \approx 10 \pm 5$ provides sufficient accuracy for all but the most extreme examples.

4.5.2 Stellar Pseudo-Scattering

Aside from the very last timestep, at which point we wish to synthesise observations (§4.7), there is no need to actually scatter the stellar component of the galaxy at every timestep. The only relevant quantity for the intermediary computations is the position of the stars when they die, such that their ejecta is successfully mixed in.

At each timestep, we therefore collate the ejected material produced in each ring, and then scatter this material according to the cumulative values of Eq. (4.12) which have accumulated since the time of the stellar birth. If the galactic mass distribution evolves according to $\mathbf{M} = \mathbf{M}(t)$, and a mass E of material is being returned to the ISM by a stellar population born in the j^{th} ring N_t timesteps ago, then the ejected material is assigned to the neighbouring rings according to:

$$\boldsymbol{\mathcal{E}} = \left(\prod_{i=0}^{N_t} \exp\left(\Delta t \times K(\mathbf{M}(t - (N_t - i)\Delta t) \right) \right) E\hat{e}_j \tag{4.14}$$

Figure 4.4 shows the result of this computation for stars born at 0Gyr and 4 Gyr. We see that stars born early on in galactic lifetime and at inner radii are able to diffuse almost completely, whilst stars in the outer galaxy disperse far less strongly, though we see that the outer-galaxy dispersion is more rapid for stars born at late times.

We further note that keeping the stars stationary during their lifetime does not affect intermediary calculations which rely on the stellar mass distribution (i.e. the computation of the scattering matrix itself), as the nature of the stochastic matrix (i.e. the angular

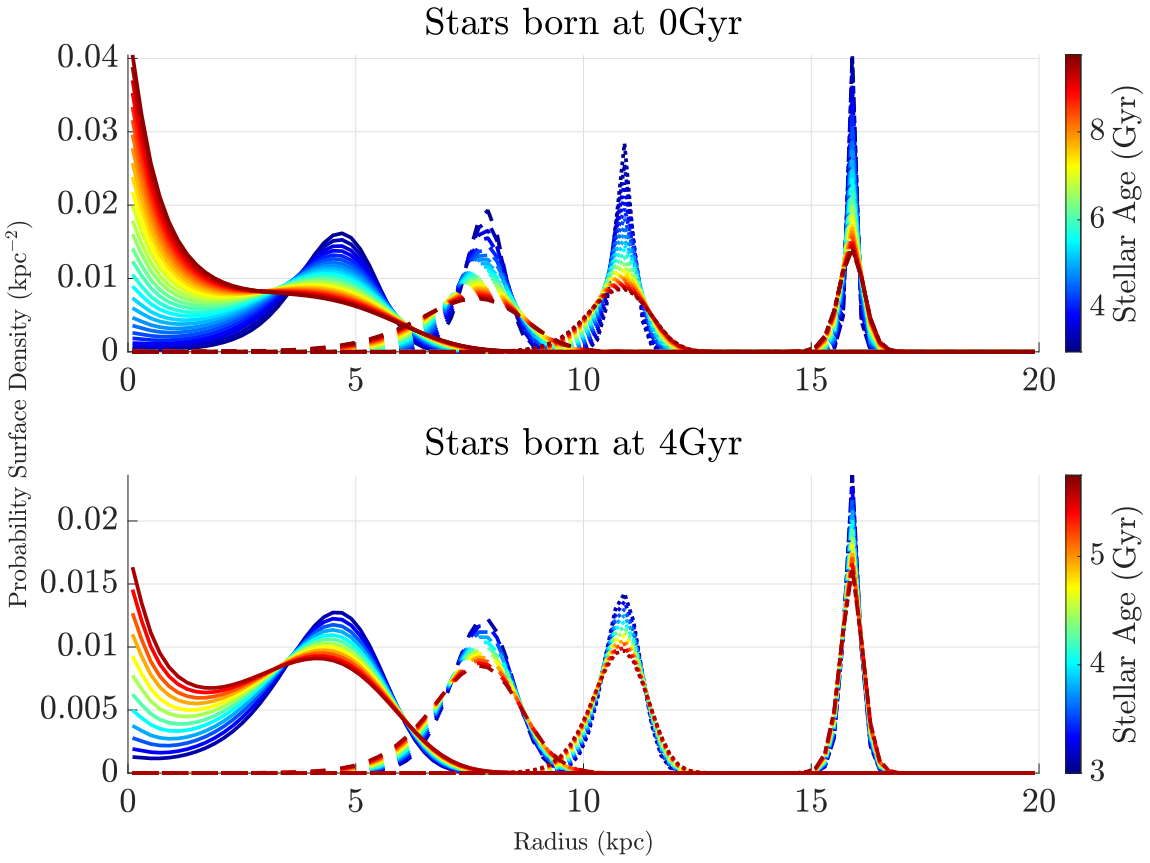

Figure 4.4: The probablity per unit area of finding a star at a given radius, for a variety of starting radii, and evolved over 5Gyr lifetime computed from Eq. (4.14), with starting times 0Gyr (*top panel*) and 4Gyr (*lower panel*). The diffusion strength is set to 0.5kpc^{2}Gyr^{-1} and the starting radii were 5, 8, 11 and 16kpc.

momentum conservation) ensures that the total stellar mass in each ring remains constant, even if the exact populations of those stars is altered.

4.6 Yields & Abundances

4.6.1 Metallicity Scale

There has been considerable work over the past few decades in attempting to discern the current surface elemental abundance of the sun. The traditional value has been taken to be $Z = 0.019 \pm 0.001$ (i.e. Grevesse & Sauval 1998). More recent work, however, has found that $Z \approx 0.014$ (Grevesse et al., 2007; Asplund et al., 2009; Caffau et al., 2011; Asplund et al., 2021), in particular as a result of Non-Local Thermodynamic Equilibrium (NLTE) assumptions and 3D modelling of stellar atmospheres, with the atmospheric abundance values of Asplund et al. (2009) (henceforth A09) becoming the literature standard.

However, there has remained considerable concern that these modern metallicity values are inconsistent with helioseismic constraints (Villante & Serenelli, 2020), and since the discrepancy arises primarily in stellar modelling, it produces a discontinuity between stellar-inferred ISM abundances (as in Figure 2.5), and direct measurements of HII regions which is very hard to reconcile.

New models (Bergemann et al., 2017; Magg et al., 2022), however, have pushed back against the downward revision of the metallicity scale, arguing that the reduced solar metallicity is an artefact of the particular calibration tools chosen, particularly for a blended nickel/oxygen line. In this work we use the abundance catalogue of Magg et al. (2022) (henceforth M22) to compute the solar-normalised abundances for our models. We note that, under this catalogue, the solar metallicity has a value $\mathcal{Z}_\odot = 0.017$.

We find that the differences between using the M22 abundances versus A09 is simply a vertical translation of, for example 0.1 dex in [O/Fe], as seen in Figure 4.5. This translation is not consistent across different elements: Whilst [O/Fe] was reduced by 0.1dex moving from A09 to M22, [Mg/Fe] was increased by 0.05dex, meaning that recalibrating i.e. the SNIa fraction cannot help us simultaneously fit the abundances to solar. We find that the A09 abundance scale is extremely difficult to simultaneously replicate solar values for [O/Fe] and [Mg/Fe], and similarly with other elements. The M22 abundance catalogue, however, shows good, consistent and simultaneous fitting of all the elements.

Whilst the differences induced by these different catalogues is has only a minor impact on our models, we emphasise that this work is still of vital importance for GCE studies, as the underlying models are then applied to infer the abundances of stellar catalogues: improvements in NLTE and 3D modelling alters not only the solar normalisation, but the abundance of the catalogues we are attempting to replicate. This is important because, unlike the solar abundance, this does not merely translate to a uniform vertical translation since more advanced models control for i.e. metallicity dependent offsets such as those found in Bergemann et al. (2017).

4.6.2 Processes

Unlike SACEM, in RAMICES II we are able to completely compute Eq. (2.14), and hence find the gross yield of a synthesis event $Y_X(M, Z)$ given the progenitor was born with some abundance distribution $\mathbf{z}$. We must therefore be able to determine the net yield, N_X for each nucleosynthetic process to be included.

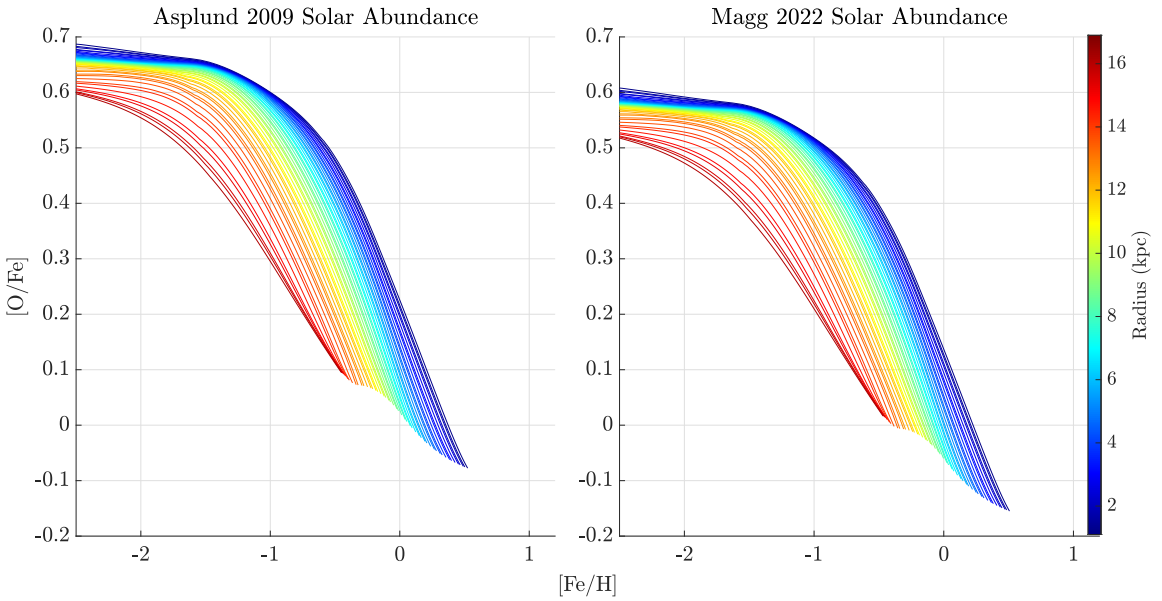

Figure 4.5: Tinsley Diagrams for [O/Fe] for RAMICES II models using *(left panel)* the A09 solar abundance catalogue and *(right panel)* the M22 catalogue.

We model each 'process' as a separate entity (this is in contrast to the treatment in RAMICES, where there is a continuum between, i.e. AGB and CCSN events), each with their own yield grid in M and Z. This allows RAMICES II to discriminate easily between material from different sources, which can prove useful in determining the dominant source of a given element.

RAMICES II models five different nucleosynthetic events, in addition to the usual CCSN and SNIa events, we include the yields (primarily stellar wind) from low mass stars (given the umbrella term 'AGB'), as well as NSM and, for the implementation in Chapter 5, collapsar events.

CCSN

All stars with masses greater than $M_{\rm ccsn} \sim 10 M_{\odot}$ undergo CCSN at the end of their stellar lifetime, which are of maximum 0.1Gyr. We derive the yield grids for CCSN events from a heterogeneous dataset derived from Maeder (1992); Chieffi & Limongi (2004); Limongi & Chieffi (2008).

We note that the grouping stellar lifetimes from continuous ages into timesteps has an undesired confluence with the temporal resolution. Changing, for example, $\Delta t = 0.01 \rightarrow 0.02$Gyr drastically alters the mass distribution of stars which die 'immediately' (i.e. in the same timestep that they were born), and hence also alters the early time abundance plane. In particular, this can result in large and unphysical variations in the [X/Fe] ratios one recovers during the first few timesteps, as demonstrated in Figure 4.6

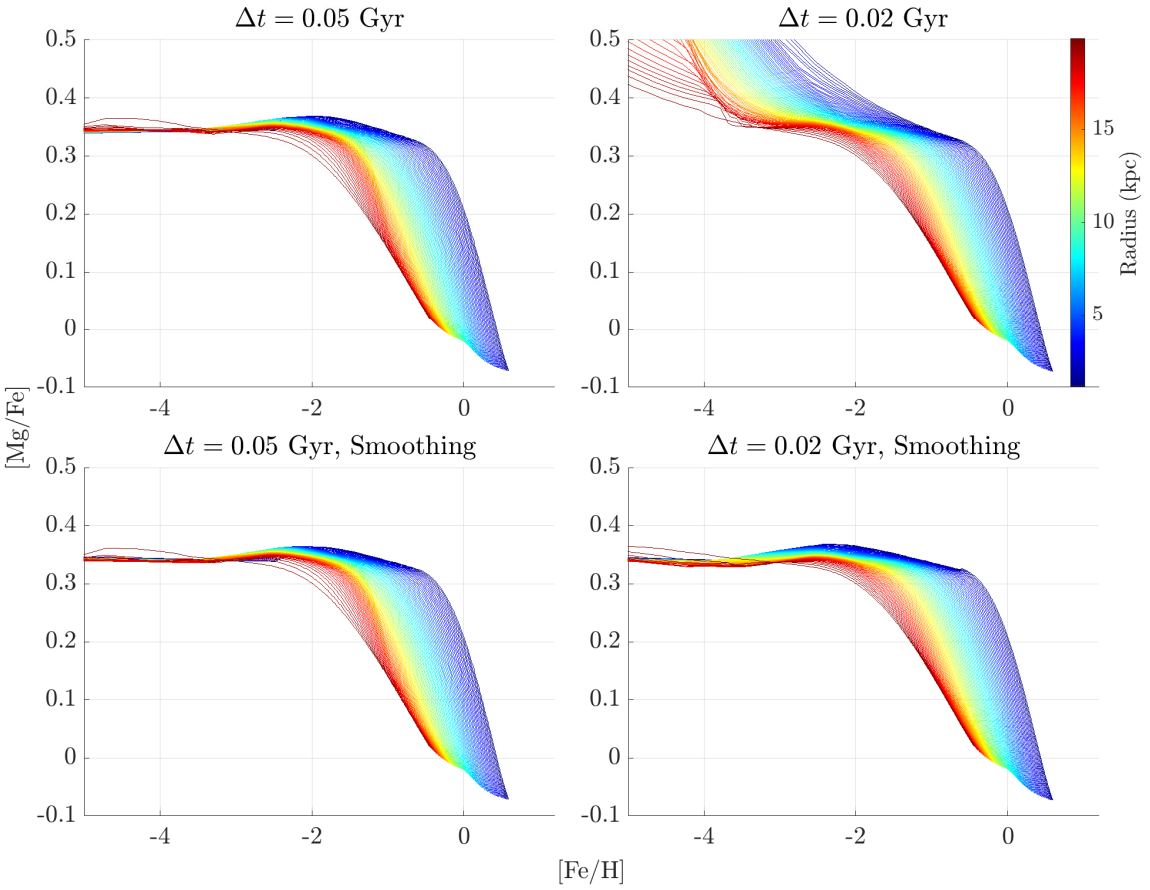

Figure 4.6: *(Top Panels)* The effect of altering the timestep on the early-time [Mg/Fe] abundances. This effect rises because a) [Mg/Fe] ejecta of massive stars varies strongly with mass, b) the timestep controls the lowest mass star able to detonate in the first few timesteps (and hence the [Mg/Fe] value of the mixed ejecta) c) the early galaxy is very easy to pollute. *(Lower Panels)* As we consider this effect unphysical, we add 'lifetime smoothing', which allows all CCSN-capable stars to explode in the first few timesteps, thereby recovering the continuum expected value.

This arises due to the strong variations in [X/Fe] of ejecta with mass, and the fact that our finite timestep effectively integrates over the IMF of stars allowed to detonate at each timestep. With small timesteps, our initial ejecta is therefore dominated by stars with lifetimes less than Δt, which have extremely high [X/Fe] for smaller Δt. Since the CGM is very easy to pollute at early times, this results a strong increase in [X/Fe], which then slowly descends to the thick-disc value. At larger Δt, we would 'integrate' over the entire population of CCSN stars, and recover a stable thick disc value at all early times. We note that the effects on the late-time chemical evolution are negligible, which is as to be expected, this is an artefact of the first few timesteps and does not persist until late times.

Whilst we have no doubt that lifetime-induced effects such as this are indeed important for studying the galactic chemistry at these early times, we note from §2.3.1 that we are

explicitly limiting ourselves to the continuum limit. The behaviour observed in Figure 4.6, however, arises in the domain where stochastic effects are important, and hence beyond the domain of our models to accurately model at the current time. To recover our continuum expected value, we impose the additional constraint that for $Z < 10^{-2.5}$, all CCSN-capable stars should be treated as having lifetimes of $\sim \Delta t$Gyr, and explode immediately. This effectively states that we have an uncertainty on our stellar lifetimes of the order 0.05Gyr, and in the case of such uncertainty, we round downwards generously. The effect of this rounding is shown in the lower panels of Figure 4.6.

ECSN

Electron-Capture Supernova (ECSN) are a hypothesised 'bridge' between the fate of low mass stars, and high-mass CCSN (Miyaji et al., 1980). In the transition region at $M{\approx}8.5M_{\odot}$, the star does not reach iron synthesis, with the nucleosynthesis instead stalling with a degenerate O+Ne+Mg core, which becomes unstable to collapse due to electron capture events negating the degeneracy pressure. The nature of the remnant left behind is hotly contested, but is generally believed to be a neutron star.

We use the Jones et al. (2019) yields for the ECSN ejecta, noting that the abundances are not dissimilar to a SNIa event, and hence has the effect of reducing the [Mg/Fe] ratio in the thick disc.

Per Jones et al. (2019), ECSN are thought to be rare events, with a rate $\sim 10^{-3}$ of that of usual CCSN. We model this rarity by approximating that any star in the region $8.5 - 10M_{\odot}$ has a probability $p_{\rm ecsn}$ of going supernova, with the remainder exploding as normal CCSN events. The parameter $p_{\rm ecsn}$ is then tuned such that ECSN occur at the desired rate relative to CCSN - we find that $p_{\rm ecsn} \sim 0.01$ reproduces the rate well.

We find no evidence that ECSN are important for the global GCE trends with $p_{\rm ecsn}$ this small. There is a minor contribution to the early time yields (essentially acting as an early time SNIa channel). However, we include ECSN in the final product for completeness, and for the purpose of future study.

AGB Death

Stars with masses lower than $M_{\rm ecsn} \sim 8.5M_{\odot}$ are too low mass to undergo either CCSN or ECSN events. Their ultimate fate is to end up as white dwarfs, ejecting their upper atmosphere in the form of stellar winds.

We give this process the umbrella term 'AGB Death', or simply 'AGB processes'. We use Maeder (1992); Marigo (2001) for H,He, and CNO yields these stars, under the assumption that very little other nucleosynthesis takes place in these stars. We know that this is probably false, and that AGB stars are an important site for s-process nucleosynthesis, and hence we intend to replace these yields with those of, for example, Karakas & Lugaro (2016), however, whilst verifying the model against RAMICES, we felt it best to retain the same yield grid.

The way in which lower-mass stars contribute ejecta to the ISM is significantly different to CCSN events, insofar as the stellar wind is a sustained injection mechanism, compared to the point-injection of a CCSN event. This would seemingly indicate that we must track the ejection of matter over the course of the stellar lifetime - however, we note that by studying the mass-loss rate within the isochrone data (§2.3.2), we can see that significant mass-loss is confined to the final few stages of the stellar lifetime (hence the umbrella term 'AGB'), which, to within the temporal precision of our simulation, is well approximated by being a point-injection event.

SNIa & NSM

SNIa and NSM are treated as very similar events, using the same 'temporal buffer' system discussed in §4.3.3, to determine which progenitors are 'active' at any given time, given the importance of temporal offsets for these processes.

We use the standard W70 Iwamoto et al. (1999) yields for SNIa events, again assuming metallicity independence. Due to the unconstrained yields of NSM events, and the degenerate yield/rate relationship, we make the simplifying assumption that 100% of the mass of a NSM is converted into r-process material, with the fraction of NS which are considered NSM progenitors then calibrating a desired [Eu/Fe] abundance.

Collapsars

As discussed in detail in Chapter 5, collapsars are thought to arise from a small fraction of stars with mass $M > M_{\rm coll}$. We note that, given the continuum approximation and the distribution-based nature of stars within the RAMICES II model, we do not stochastically choose a fraction ξ of stars to go collapsar vs. those which go supernova, instead injecting the yields in a fractional fashion (i.e. if 1 star of $M > M_{\rm coll}$ dies, we inject the yields for ξ collapsars, and $(1 - \xi)$ CCSN events).

Since there is a strong degeneracy between the ejected Eu mass, and the frequency with which these events occur, ξ, we make the same approximation as we did with NSM:

we assume that the entire star is converted into Eu ($N \sim 1$), and then choose ξ to reproduce the required Eu values.

As discussed in §5.2.1, collapsars cease to form at a metallicity $Z_{\rm c}$, which we implement by making ξ a function of metallicity, which is constant before linearly decreasing to 0, over a width ΔZ, starting at $Z = Z_{\rm c} - \Delta Z$.

4.7 Synthesising Stellar Observations

The galactic history of a model galaxy is well-encapsulated by the cold-gas abundances. However, we note that this is *not* the quantity which is measured during chemical abundance surveys, which is the photospheric abundances of those stars which meet sufficient criteria to enter into spectroscopic pipelines.

Whilst the cold gas abundances allows one to get a broad, global picture of the environment in which stars have been forming, to properly ensure that our models are reproducing the correct dynamical and chemical properties of the observed catalogues, we must translate our models into a mock abundance catalogue.

4.7.1 A Note of Caution

We first issue a note of caution that it is not feasible to truly replicate the entire set of selection effects induced in large scale surveys. Instead, we merely wish to replicate the effects to a 'close enough' degree of precision that our mock sample and the real sample are mimicking similar behaviours.

We therefore make a number of simplifying approximations in our construction of the spatial and photometric cuts which determine the number of simulated stars included in the final catalogue.

4.7.2 Survey Properties & Selection Function

In order to synthesise a mock catalogue from the stellar populations held within the code, we must therefore apply a number of cuts to simulate the selection function of the survey. We deliberately leave the nature of these cuts open to future changes - for the present work, we simulate the selection function of the GALAH+ DR3 survey, as detailed in Buder et al. (2021).

For simplicity, we simulate only the 'GALAH Main' portion of the catalogue, which constitutes the bulk of the data, though note that extending this to the other components is simple. Although the GALAH survey makes some complex tiling decisions as to the

locations of their individual fields of view, our per-ring homogeneity assumptions renders these particular nuances moot, such that we may assume the criteria for entering the catalogue is:

- Visual Magnitude[3] in range $12 < V < 14$
- Off-Disc (but non-halo), such that Galactic Latitude $10° \leq |b| \leq 30°$
- Visible to Southern-Hemisphere[4], such that $\delta < 10°$

Although the GALAH survey does not discriminate or favour giant stars against dwarfs, flags are applied to candidates based on their inferred properties, with giants being those with $T_{\text{eff}} < 5500\text{K}$ and $M_K < 2\text{mag}$, the remainder being flagged as dwarfs.

We can use the same PARSEC isochrones utilised for the stellar lifetimes (§2.3.2) to determine the photometric properties of each stellar population within the ring. The fraction of stars of absolute visual magnitude M_V and age τ in ring i which meet the photometric and angular cuts of GALAH is then computed from:

$$f_{\text{GALAH}}(i, M_V, \tau) = \int_{-\infty}^{\infty} \mathrm{d}z \int_{R_i - w_i/2}^{R_i + w_i/2} \rho \mathrm{d}\rho \int_{-\pi}^{\pi} \mathrm{d}\varphi S\left[d(\rho, \varphi, z), M_V\right] n_i(\rho, \varphi, z, \tau) \tag{4.15}$$

Where $n_i(\rho, \varphi, z, \tau)$ is the spatial distribution function of stars of age τ, and

$$S\left[d(\rho, \varphi, z), M_V\right] = \begin{cases} 1 & \text{if } \begin{cases} 10^{\frac{2-M_V}{5}} \leq \sqrt{d_\odot^2 + \rho^2 + z^2 - 2d_\odot^2 \rho \cos(\varphi)} \leq 10^{\frac{4-M_V}{5}} \\ 10° \leq |b| \leq 30° \\ \delta_{\text{eq}} < 10° \end{cases} \\ 0 & \text{else} \end{cases} \tag{4.16}$$

Here $d_\odot$ is the solar orbital radius, and the angular coordinate φ is chosen such that $\varphi_\odot = 0$, and d has implicit units of kpc. δ_{eq} is the equatorial declination angle and b_{gal} is the galactic latitude.

We find that the position of the 'Southern Sky cut', $\delta < 10°$, when projected into Galactic (ℓ, b) coordinates are extremely well approximated by the following Fourier series:

$$\tan\left(\left(b_{(\delta<10)}(\ell)\right)\right) = 0.715 - 1.24\cos(\ell°) + 1.915\sin(\ell°) - 0.114\cos(2\ell°) - 0.2553\sin(2\ell°) \tag{4.17}$$

[3]GALAH actually uses a synthetic estimate for the visual magnitude $V \approx V_{JK} = K + 2(J - K = 0.14) + 0.382\exp\left(\frac{1}{2}(J - K - 0.2)\right)$, which is justified by comparison with Figure 1 in Sharma et al. (2018)

[4]Although the tiles are uniformly distributed in the Southern Hemisphere, in galactic coordinates this translates to an asymmetric focus around the galactic centre, with stars at $\ell = 0°$ only visible at heights $b > 30°$.

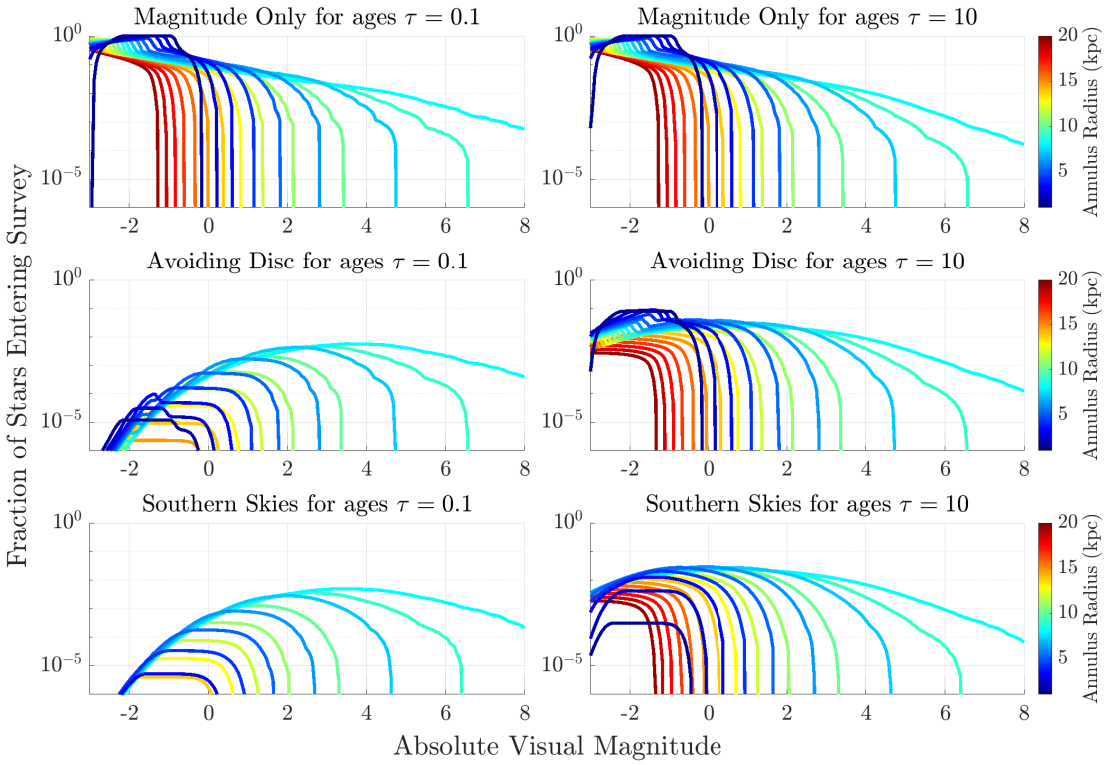

Figure 4.7: The fraction of stars within a given annulus entering into the survey (derived from Eq. (4.15)) for stars of age $\tau = 0.1$Gyr (left) and $\tau = 10$Gyr (right). The top row shows the result of only including the magnitude selection, the middle row the magnitude and the cut $|b| > 10°$, and the final row the complete set of constraints, including $\delta < 10$. We note that this is explicitly only the fraction of *surviving* stars - it is extremely unlikely that a star which reaches $M_V = -2$ has a lifetime in excess of 10Gyrs, but we have included this for completeness.

4.7.3 Simulating Vertical Structure

The in-disc homogeneity assumptions of RAMICES II imply that the angular and azimuthal components of the spatial distribution function are trivial:

$$n_i(\rho, \varphi, z, \tau) = \frac{1}{2\pi R_i w_i} \xi(z, \tau) \tag{4.18}$$

It therefore remains only to determine the vertical distribution function, $\xi(z, \tau)$ for stars of age τ. Since the stellar discs are observed to have an exponential profile, with the thick disc having scale length $\bar{z} \sim 1$kpc and the thin disc having $\bar{z} \sim 0.3$kpc, we parameterise the vertical structure as:

$$\xi(z, \tau) = \frac{1}{2\bar{z}(\tau)} \exp\left(-\frac{|z|}{\bar{z}(\tau)}\right) \tag{4.19}$$

The temporal evolution of $\bar{z}(\tau)$ then determines the vertical extent of the stellar components at each age. Binney et al. (2000) finds a lower floor of $\sigma_z \sim 4$km/s in the velocity dispersion

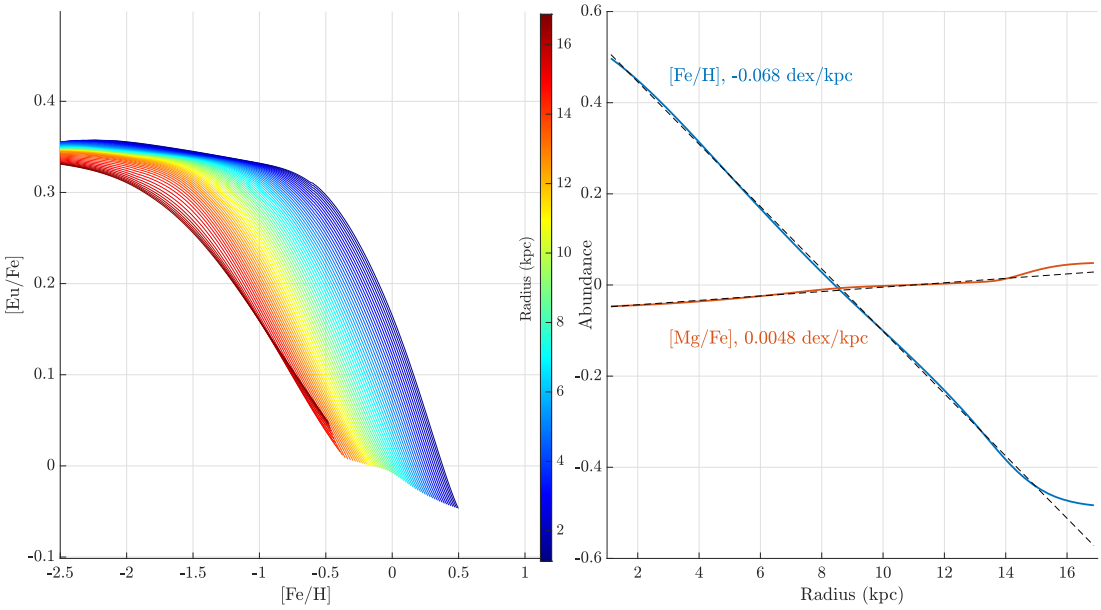

Figure 4.8: The [Mg/Fe] evolution for the nominal model, which is calibrated to replicate the gas properties of RAMICES. *(Left Panel)* The associated Tinsley Diagram, and *(Right Panel)* The final-time gradients of [Fe/H] and [Mg/Fe], showing values which are consistent with Figure 2.5.

at very low ages, followed by a relationship $\sigma_z \propto \tau^{1/3}$. This lower bound in σ_z equally implies a lower bound, $\bar{z}_{\rm zams}$ in the vertical scale length at $\tau = 0$. Similarly, since $\bar{z} \propto \sigma_z^2$ at low altitudes, we predict:

$$\bar{z}(\tau) = \bar{z}_{\rm zams} + k\tau^{2/3} \tag{4.20}$$

The values of $\bar{z}_{\rm zams}$ and k are then determined by requiring that the chemical thin and thick discs have appropriate composite scale heights in the solar neighbourhood.

We note again the warning of §4.7.1, in that the assumption of iso-temporal populations having an exponential profile is not self-consistently motivated. The choice can be seen as an analytical approximation to a model in which stars form isothermal populations distributed according to $n(z) \propto \exp(-\Delta\Phi(z)/\sigma_z(\tau))$, where $\Phi(z)$ is the potential of the galaxy. The results here are close to the self-consistent value, but have the added advantage of being trivially analytically integrable, which allows us to simplify the integral Eq. (4.15), as demonstrated in Appendix A.1

4.8 The Nominal Model

In order to inherit as many of the calibrations performed on RAMICES over the years (Schönrich & Binney, 2009a; Bilitewski & Schönrich, 2012; Schönrich & McMillan, 2017; Schönrich &

Table 4.2: The parameter values used in the RAMICES II nominal model (see also Table 4.1). For brevity, we use the shorthand $m_9 = 10^9 M_\odot$ to indicate our mass units.

Parameter	Physical Meaning	Value
$N_{\rm rings}$	Number of annuli	100
ΔR	The (uniform) width of each annulus	0.2kpc
Δt	Duration of a simulation timestep	0.02Gyr
$t_{\rm final}$	Final age of the galaxy at the end of the simulation	12Gyr
$\kappa_{\rm schmidt}$	Prefactor in the SFR, Eq. (2.4)	$2.34 \times m_9^{-0.4}\mathrm{Gyr}^{-1}$
$\eta_{\rm feedback}$	'Mass loading' factor for stellar feedback, Eq. (2.8)	0.3
$e_{\rm feedback}$	Feedback-ejection parameter §4.3.2	0.9
ℓ	Ejection fraction, Eq. (4.5)	0.25
M_0	Initial gaseous mass of the galaxy	$0.1m_9$
M_1	Mass of the first onfall, Eq. (2.15)	$4m_9$
M_2	Mass of the second onfall, Eq. (2.15)	$100m_9$
b_1	Timescale of the first onfall, Eq. (2.15)	0.8Gyr
b_2	Timescale of the second onfall, Eq. (2.15)	9Gyr
$M_{\rm cgm}$	Mass of the CGM, §4.4.4	$200m_9$
$\tau_{\rm cool}$	Cooling timescale	$1\ \mathrm{Gyr}^{-1}$
$q_{\rm snia}$	Fraction of progenitors able to undergo SNIa	0.075
$\tau_{\rm snia}$	Delay time before SNIa turn-on	0.4Gyr
$\nu_{\rm snia}$	Exponential decay time in the DTD for SNIa events	1.5Gyr
$q_{\rm nsm}$	Fraction of progenitors able to undergo NSM	1×10^{-5}
$\tau_{\rm nsm}$	Delay time before NSM turn-on	0.05Gyr
$\nu_{\rm nsm}$	Exponential decay time in the DTD for NSM events	4Gyr
$f_{h,\rm ccsn}$	Hot gas injection fraction for CCSN	0.7
$f_{h,\rm agb}$	Hot gas injection fraction for AGB	0.7
$f_{h,\rm snia}$	Hot gas injection fraction for SNIa	0.99
$f_{h,\rm nsm}$	Hot gas injection fraction for NSM	0.4
$\mu_{\rm mix}$	Mixing strength parameter, Eq. (4.11)	$0.5\ \mathrm{kpc}^2\mathrm{Gyr}^{-1}$
$a_{\rm Inflow}$	The 'Bilitweski a factor', §4.4.2.	0.4
$b_{\rm Inflow}$	The 'Bilitweski b factor', §4.4.2.	0.3
$c_{\rm snIa}$	The fraction of COWD which are SNIa progenitors	0.07
$c_{\rm nsm}$	The fraction of NS which are NSM progenitors	0.00015
$c_{\rm ecsn}$	The fraction of stars in the range 8.5-10$M_\odot$ which are ECSN progenitors	0.001

Weinberg, 2019) as possible, the nominal RAMICES II model we present here is calibrated to replicate many features of its predecessor. The values which produce the greatest fit with RAMICES are presented in Table 4.2. Figs. 4.8 and 7.1 show the evolution of [Mg/Fe] and [Eu/Fe] in our nominal model.

Since we have tuned RAMICES II to be replicate RAMICES at the solar annulus, the differences in the cold gas evolution therefore arise solely from the different physics which we have incorporated into the model, the most important of which is the differing CGM prescription. RAMICES set the CGM to be an infinite reservoir of metallicity $Z = 0.1Z_\odot$, mimicking the individual elemental abundances of the half-mass annulus. This caused

unphysical artefacts at radii $\gtrapprox$ 14kpc (the point at which star formation was suppressed by Σ_c), as the abundance became dominated by the CGM prescription.

In RAMICES II, we treat the IGM/CGM as an evolving gas reservoir in its own right, polluted by both feedback and supernova events. This means that, unlike RAMICES which reacts instantaneously to the present day abundances of the galaxy, the CGM acts on a delay. During the SNIa 'knee', where [Mg/Fe] values fall rapidly in the disc, the CGM retains the prior imprint of the α-rich gas, and hence acts to oppose the change in [Mg/Fe]. We therefore see an increase in [Mg/Fe] in the outer galaxy compared to equivalent RAMICES models, in which [Mg/Fe] is almost perfectly flat across the galaxy, however we note that whilst the gradients in Figure 2.5 are consistent with a flat [Mg/Fe] in the galaxy, the signal indicates a weakly positive gradient, which the RAMICES II model replicates.

We are therefore confident that the cold-gas evolution of the RAMICES is consistent with that of RAMICES II, outside of the CGM-dominated region, and is in turn consistent with the observations of Figure 2.5.

4.8.1 Stellar Catalogue Synthesis

Having calibrated our cold gas abundances to those produced by the RAMICES model, we turn to the synthesised stellar catalogues. As per §4.7, this catalogue is meant to mimic the GALAH DR3 data displayed in Figure 2.2. Figure 4.9, however, shows a concerning deviation from the stellar catalogue we were targeting.

In particular, whilst Figure 2.2 had a strong secondary population at [Fe/H] ~ -0.5, almost directly above the thick disc, we see instead in Figure 4.9 a long tail off to low metallicities. When projected into an [Mg/Fe] histogram, we recover a similar bimodal population, however it appears that the thick disc in Figure 4.9 is significantly more extended in [Fe/H].

We note that this, might, however, be due to a quirk in the GALAH chemistry pipeline, or equivalently in our synthesis of the GALAH selection function. We note that, in Figure 2.3, we removed all stars with $\sigma_{\text{[Mg/Fe]}} > 0.2$dex, which had the effect of removing almost all stars with [Fe/H] < -1. This could be a strong indication that the GALAH chemical pipeline is simply ill-suited for accurately determining the chemical abundance of metal-poor stars, such that even though a star met the selection criteria to enter the survey, an accurate abundance determination was not made. As we do not attempt to mimic the effects of the GALAH spectroscopic pipeline – only the entry into the survey itself – this could explain why our model predicts far more low-metallicity stars than is observed in GALAH.

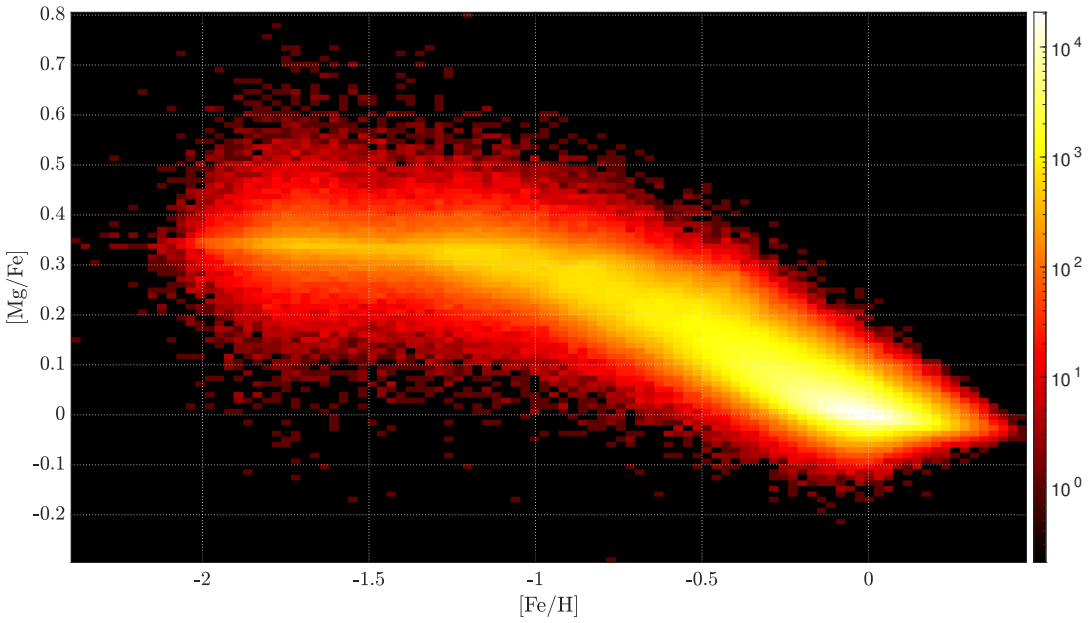

Figure 4.9: The stellar catalogue associated with the cold gas abundances of 4.8, following the prescription of §4.7. We note that this shows a significant deviation from the target data in Figure 2.2 in the thick disc.

We note, for example, that both Figure 2.4 (using data from G and F dwarfs of Bensby et al. (2014)) and Figure 5.1 (using data from SAGA, again primarily dwarfs) shows a similarly extended population in [Fe/H], which is again validated by Figure 4 of Schönrich & McMillan (2017), which is designed to mimic the Geneva-Copenhagen Survey (Nordström et al., 2004) (again, targeting dwarf stars), and demonstrates an almost identical extended-thick disc pattern to our own.

Therefore, whilst we consider the argument that the GALAH DR3 abundance pipeline preferentially omits stars with low metallicities, despite meeting the spectroscopic requirements to enter the survey likely, the question of the origin of this difference is undetermined – our synthesised catalogue is strangely consistent with catalogues we were not attempting to reproduce, and yet inconsistent with the target. We shall leave the question of why our thick disc differs so strongly to a future endeavour.

An exotic birthplace on its own is not informative of anything.

— Italo Calvino

5

The Origin of r-Process Nucleosynthesis

Contents

This chapter has been published, with minor alterations, as Fraser & Schönrich (2022).

5.1 Introduction

As discussed in Chapter 1, the landmark B^2FH paper set out the current paradigm of nucleogenesis, with primordial nucleosynthesis sourcing the lightest elements in the universe (H and He, Alpher et al. 1948), while the heavier elements are created in processes such as shell burning in the stellar interior (the α elements, including O, Si and Mg, Hoyle 1954), explosive nucleosynthesis during supernovae events (α elements, plus the iron peak elements: Fe, Ni, Co, Arnett & Clayton 1970) or cosmic ray spallation (Li, Be, B, Reeves et al. 1970).

The majority of the elements heavier than iron, however, are sourced from a variety of neutron-capture processes: the slow (s), intermediate (i) and rapid (r) neutron capture processes, where the implied 'speed' of these processes denotes the neutron capture timescale relative to the β-decay timescale of the intermediary elements. In the slow-process, β decays are able to occur between every successive neutron capture, resulting in synthesis up the β-stable track. In the rapid-process, neutron absorption occurs so quickly that no intermediary decays are able to occur, and the synthesis occurs up the so-called 'neutron drip lines'.

In all neutron capture chains, the process is interrupted at the magic number nuclei, P or $N = 50, 82, 126$ which are highly stable against further absorption. This results in a buildup of material at these magic numbers - the first, second and third peaks. For the semi-stable s-process elements, this results in a tightly grouped abundance peak around each of Zr, Ba and Pb, whilst the instability of r-process material results in β decay only after neutron flux has ceased, and hence in a much wider peak offset from the s-process, centred on Se, Te and Pt respectively.

Whilst the site of the s-process is well constrained (Clayton et al., 1961), and the i-process thought to contribute significantly to only a handful of isotopes (Côté et al., 2018), the debate about the astrophysical sites that can lead to sufficient r-process synthesis has remained an open and enduring question for many years, with several possible processes capable of igniting the r-process being identified.

Woosley et al. (1994) argue that neutrino heating in CCSN creates favourable conditions such that normal CCSN can be a source of r-process ejecta This model, however, is plagued by overproduction of certain elements, and the high entropy conditions required have been questioned in more recent studies (i.e. Fischer et al. 2012). In addition, the existence of ultra metal poor, but highly r-process enriched stars (such as that found in Sneden et al. 1996) indicates that the source of r-process nucleosynthesis must be a rare, high yield event. CCSN are therefore not considered a viable source of r-process enrichment.

The disruption of a neutron star by tidal interactions during a merger with a black hole (Lattimer & Schramm, 1974), or by a binary collision between two neutron stars (Symbalisty & Schramm, 1982; Freiburghaus et al., 1999; Rosswog et al., 1999) are also candidates for the r-process. The detection of the combined gravitational wave GW170817 and Short Gamma Ray Burst (SGRB) event GRB170817A, confirmed to arise from a NS-NS merger event (Abbott et al., 2017), and the subsequent detection of r-process material in the ejecta (Chornock et al., 2017; Rosswog et al., 2018) provided direct observational evidence that NSM produce r-process material.

Though the existence of NSM as an active r-process pathway is rarely called into question, it is seen as concerning that time-delayed nature of NSM formation would naïvely predict entirely different enrichment pathways in [Eu/Fe]-[Fe/H] space than is observed, leading to either the conclusion that NSM cannot be dominant r-process sources (Argast et al., 2004; Wanajo & Ishimaru, 2006), or the invocation of neutron star properties incompatible with their understood behaviour (Matteucci et al., 2014).

In an alternative approach, Fujimoto et al. (2006) combined an MHD jet method with the collapsar models of Woosley (1993), and demonstrated that this can produce a significant amount of r-process enrichment. Collapsars are a corollary to the existence of the 'failed supernovae' proposed in Bodenheimer & Woosley (1983) and since observed by Adams et al. (2017). 'Failed Supernovae' are the fate of stars which are so massive that the usual supernova mechanism is insufficient to prevent runaway radial collapse. However, progenitors with large angular momentum cannot collapse spherically. Instead, they collapse into a compact accretion disc around the growing black hole: a potential site for r-process nucleosynthesis and LGRBs.

Although LGRB are well-documented events and often tied to unusual forms of CCSN due to their formation in regions of rapid star formation (Bloom et al., 2002) and several closely tied observations of supernovae associated with LGRBs (Kulkarni et al., 1998; Mazzali et al., 2003; Sollerman et al., 2006), there is no direct evidence linking their formation with r-process synthesis. However, the model is widely favoured by sceptics of the NSM mechanisms, since the high progenitor masses imply a short lifetime and thus allow for very early r-process enrichment.

Other sources for r-process material have also been studied. For example, neutrino-driven winds (Wanajo et al., 2001) and electron-capture supernovae (Wanajo et al., 2011), have been suggested as sources, however Haynes & Kobayashi (2018) showed that these did not produce sufficient quantities of r-process material. Whilst in the case of the intermediary (i) process-producing White Dwarf binaries (Cowan & Rose, 1977; Denissenkov et al., 2017), it was predicted by Côté et al. (2018) that only specific isotopes are produced in significant quantities, with 45 per cent of solar Molybdenum predicted to be i-process in origin, but less than 10^{-2} per cent of solar Eu. For the sake of clarity and simplicity, we will therefore neglect these sources.

Confusing matters further, there is also evidence of an *incomplete* (or *weak*) r-process (Honda et al., 2006), in which the lighter r-process elements are synthesised, but not the heavier second and third peak elements. Magneto-rotational supernova discussed in Nishimura et al. (2017) (also referred to as 'hypernovae', though this phenomenological term can refer to collapsars) or the recently proposed Quark-Deconfinement Supernovae of Fischer et al. (2020) are thought to be good candidates. For our purposes, we use the term 'r-process synthesis' to refer to the *complete* r-process, in which all r-process material up to the third peak is synthesised.

It might feel natural to assume that multiple pathways actively contribute to the r-process enrichment, with collapsars providing the early time yield, and then NSM coming in later. However, this bears two problems: i) the enrichment profiles for [Eu/Fe] are poorly replicated in simple GCE models whenever NSM are significant contributors, and ii) Sneden et al. (1996)

demonstrated a common r-process fingerprint: a remarkable consistency of relative r-process abundances, with Sneden et al. (2008) (henceforth Sn08) extending this relationship. The relative abundances of r-process material in the metal-poor but highly r-process enriched star CS 22892-052 match those of the solar system, despite the fact that the high enrichment indicates very early enrichment from the unmixed ejecta of a single r-process event. The common 'fingerprint' with the solar system implies that the material is produced in the same ratio throughout galactic history that dominates today's abundances, a tension if one would like to assume that the dominant[1] channel switched from collapsars early on to NSMs today.

In this chapter, we will put quantitative limits on the relative contributions from both NSM and collapsars by comparing chemical evolution models with observed stellar abundances. This will show that under reasonable assumptions the relative contribution from collapsars is highly limited.

§5.2 will introduce some of the key aspects of collapsar formation as it pertains to chemical evolution, with §5.3 detailing the observational data that we will attempt to replicate in our models. § 5.4 introduces the analytical SACEM model, whilst § 5.5 outlines our attempt to eliminate regions parameter space, and §5.6 and §5.7 discuss our findings regarding the excluded regions of parameter space, and the required properties of our models.

5.2 Modelling Collapsars & Their Yields

5.2.1 Dependence on Initial Metallicity

The nature of collapsars necessitate a mass and core angular momentum cutoff for their formation: the star must be massive enough to defy normal supernova mechanisms, but have enough internal angular momentum to stave off radial collapse. There are also strong indications of a metallicity dependence on collapsar rates, both from models of how metallicity impacts the stellar interior (through variations in the opacity and angular momentum transport efficiencies), and from observational Long Gamma Ray Burst (LGRB) counts such as Perley et al. (2016).

Theoretical models of stellar evolution quote a strict cutoff in metallicity around $Z_c = 0.1Z_\odot$ (MacFadyen & Woosley, 1999; Yoon & Langer, 2005; Langer & Norman, 2006; Woosley & Heger, 2006), up to $Z_c = 0.3Z_\odot$ (Yoon et al., 2006). In these models, stars with $Z > Z_c$ undergo core-braking to below the critical threshold, and hence cannot be progenitors of collapsar events.

[1]Throughout this chapter we use the nomenclature that a *dominant* source is one which can be assumed to be the sole source, neglecting all others. A non-dominant source which produces more than 50 per cent of production is a *majority* source.

In contrast, whilst observational constraints of LGRB events find suppressed collapsar activity at higher Z, they instead find a much larger cutoff value, as high as $Z = Z_{\odot}$ (Wolf & Podsiadlowski, 2007), though they do find suppression beginning beforehand. Perley et al. (2016) propose that this discrepancy arises because there are both single-progenitor and double-progenitor pathways for GRB events, with the single progenitor pathway (isolated collapsars) dying away fastest.

Some proposed double-progenitor LGRB models (i.e. Podsiadlowski et al. 2010), however, source their energy from explosive detonation inside a common envelope rather than accretion disc/central engine interactions, making a significant contribution to r-process nucleosynthesis questionable. In addition, some observed LGRB events have no accompanying supernova detected (Fynbo et al., 2006), further indicating a diversity of origins for LGRBs beyond r-process-producing collapsars.

In summary, current observational evidence indicates that r-process producing collapsar events form a non-constant subset of all LGRB events. This is concerning as many studies (i.e. Si19) assume that the r-process nucleosynthesis tracks the GRB rate perfectly: such models thus might source r-process elements from collapsars long after they have stopped contributing in reality.

5.2.2 r-Process Yields

In principle, one may therefore continue to search for conditions which would lead to collapsar formation, and couple these with rotational, metallicity and stellar-mass distribution functions and yield tables to simulate the 'true' collapsar contribution. However, this approach is unfeasible for a number of reasons:

- The models of e.g. Heger et al. (2000) and Limongi & Chieffi (2018) show a complex/non-linear relationship between initial rotation and final core angular momentum, barring simple predictions for the likelihood of a star going collapsar from initial conditions.

- The physical processes occurring in the interior of massive stars are highly inaccessible. As a result, exact numerics of the yields from a given system must be considered poorly unconstrained.

Without rigorous observational or theoretical constraints, our models will therefore need to utilise an approximation for the collapsar yield function, $Y_{\rm coll}(M, Z, v_{\rm zams})$. We choose the simplest approximation: a constant yield which is suppressed at a metallicity $Z =$

Z_c. We justify this approach in §5.5.4 and conclude that it has negligible impact on the strength of our conclusions.

5.3 Patterns in the R Process Abundances

In this section we delineate the main features of the r-process abundances, both in the [Eu/Fe]-[Fe/H] plane, and as a function of time.

5.3.1 Europium as a tracer

It is convenient for us to use a proxy for the total r-process enrichment: Europium (Eu). Europium is chosen over similar r-process elements such as Gadolinium (Gd) and Dysprosium (Dy), since Europium is one of the purest r-process elements, being ∼98 per cent r-process in origin (Sn08). Europium has strong spectral lines in the optical spectrum, with well known oscillator strengths (Biemont et al., 1982) and a large amount of associated data.

The usage of a direct proxy might be called into question due to hints of differing trends between Eu, Gd and Dy (Guiglion et al., 2018), and the breakdown of the 'common fingerprint' for lighter r-process elements (see Sn08). Whilst the effects of the assumptions of Local Thermodynamic Equilibrium have been studied in Europium (i.e. Zhao et al. 2016), such corrections have not been calculated for Gd and Dy. It is therefore plausible that the reported differences in the trends between Eu, Dy and Gd vanish upon full consideration of NLTE and 3D atmospheric modelling.

On balance, Eu serves as a convenient proxy for r-process elements.

5.3.2 The Abundance Planes

Supplementing the GALAH data we used to produce our canonical model in §2.1, we require a sample of stars which have high-quality [Eu/Fe] measurements. Figure 5.1 shows a sample of 965 stars drawn from the SAGA database[2] which possess both upper and lower bounds for all of the elements of interest, and have [Fe/H] $>$ -2.5, since the low metallicity end of the distribution is dominated by stochastic processes, and by stars in the galactic halo which formed within dwarf galaxies accreted during the growth of the Milky Way. The low metallicity end therefore likely represents a superposition of the chemical histories of these dwarf galaxies (Ojima et al., 2018), rather than the in-situ history of the Milky Way. Due to the low mass of such accreted objects, their final impact on galactic chemistry is negligible, and this is therefore beyond the scope of this work to discuss.

[2] http://sagadatabase.jp

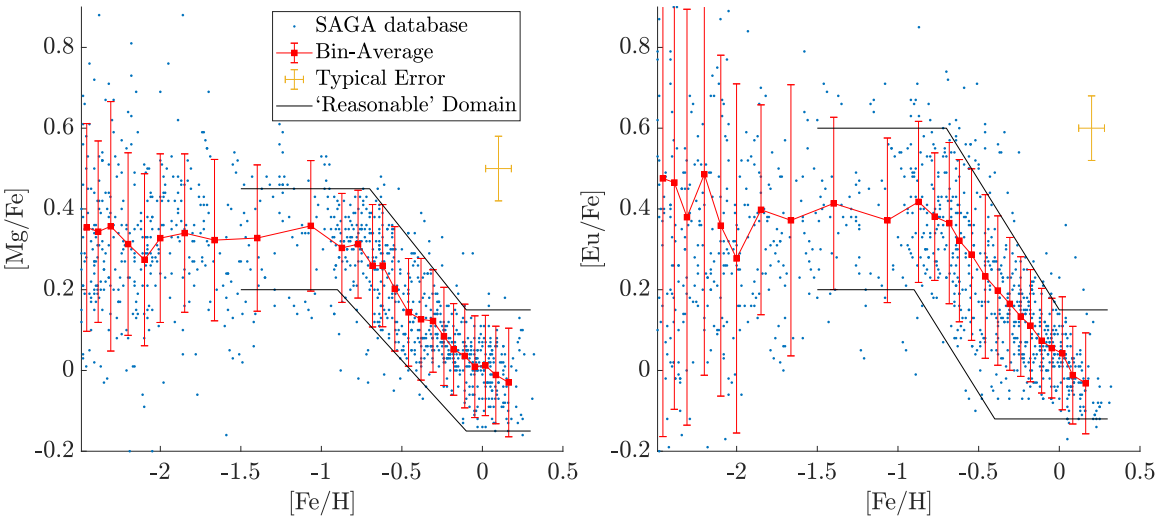

Figure 5.1: Chemical abundance plots for the 965 stars in the SAGA database with Eu, Fe and Mg measurements with both upper and lower error bounds, and [Fe/H]>-2.5. In both panels, blue dots show datapoints from this SAGA subset, and red squares show the average of bins containing 40 stars,with the error bars showing the intrinsic scatter of the bin. *Left Panel*: the [Mg/Fe]-[Fe/H] plane, *Right Panel*: the [Eu/Fe]-[Fe/H] plane. The black lines denote the 'limiting model domains', the regions which a theoretical model cannot leave for it to be considered consistent with this data, as discussed in §5.5.5.

The chemical data in Figure 5.1 are sourced from a variety of surveys, and encompasses studies of disc FGK dwarfs (Mishenina et al., 2013; Reddy et al., 2006), dwarfs in both the disc and the halo (Fulbright, 2000), as well as studies of both giants and dwarfs in the galactic halo (Sakari et al., 2018; Allen et al., 2012), thereby providing us with a wide sample of the enrichment of the galaxy both in physical and temporal space. Following the prescriptions of Bergemann et al. (2017) and Zhao et al. (2016), we performed a minor NLTE correction for both Mg and Eu for stars with [Fe/H] $<$ -2, though this affected our mean trends by less than 0.02 dex, and so did not meaningfully alter our results.

The mean behaviour of [Eu/Fe] and [Mg/Fe] for this sample is shown in red in Figure 5.1. At low metallicity ([Fe/H] < -1), we see that the bin-average shows behaviour consistent with an approximately flat curve in both planes with [Eu/Fe] $\sim$ 0.4 and [Mg/Fe]$\sim$ 0.3, corresponding to [Eu/Mg] $\sim$ 0.1.

5.3.3 Time Dependence

We leverage the data from Bensby et al. (2014) and Battistini & Bensby (2016) to constrain the time evolution of the galactic chemistry, in particular the relative change of [X/Fe] with stellar age. These papers provide age and abundance estimates for 714 G and F dwarfs in the

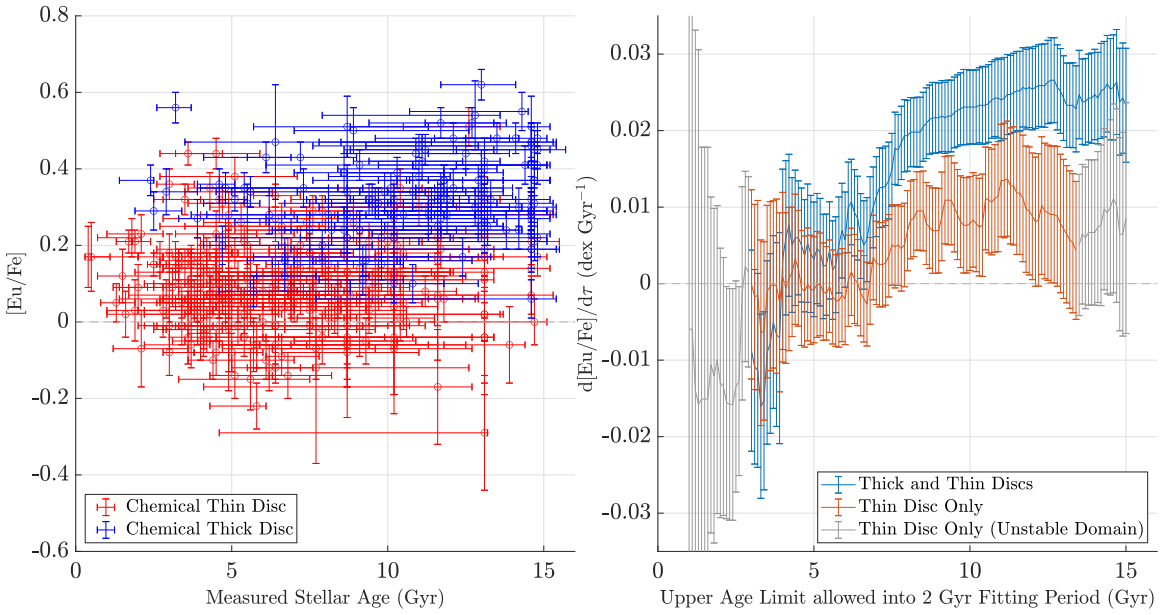

Figure 5.2: Temporal abundance data derived from Bensby et al. (2014) and Battistini & Bensby (2016). *Top:* The compiled data showing the inferred stellar ages and [Eu/Fe] abundances. *Bottom:* a plot of the calculated abundance-age gradients of each temporal subsample as a function of the maximum age permitted into the set. Bins shaded grey have fewer than 10 stars in them and so are liable to large errors.

Solar Neighbourhood. A subset of 339 stars has ages and abundances for [Eu/Fe] with both upper and lower limits. These datapoints are plotted in the first panel of Figure 5.2.

The dataset also assigns membership probabilities of each star to the Galactic thick and thin discs based on kinematics only. However, due to the large overlap of disc components in kinematics, we instead prefer a chemical cut in the [Mg/Fe]-[Fe/H] plane, while we also tested that our conclusions remain unchanged irrespective of this strategy.

We estimate the best-fit gradient of a given sample of stars, marginalising over unknown errors (the dataset provided only [Fe/H] uncertainties, not [Eu/H]). Using 2-Gyr long sampling periods, we build up a picture of how this gradient changes with the age of the stars, shown in the second panel of Fig 5.2.

We see that at early times the time derivative d[Eu/Fe]/dτ strongly positive – around +0.02 dex per Gyr for the combined sample and +0.01 dex per Gyr for the chemical thin disc. A positive gradient with stellar age is equal to a negative gradient with respect to forward time, so this shows that the [Eu/Fe] abundance was decreasing during this period.

Between 7 Gyr ago and the present day, however, gradient has decreased such that the average change in [Eu/Fe] is consistent with zero across this time period, indicating that chemical equilibrium was reached approximately 7 Gyr ago.

We note that the data of Figure 5.2 is derived under the assumption of LTE, and so both the age and abundance estimates may change under a full NLTE treatment. However, we note from Figure 15 of Zhao et al. (2016) that significant NLTE corrections for the elements in question occur (if ever) for stars with $[\mathrm{Fe/H}] \lesssim -1$, which implies that stars younger than 10Gyr are safe against this bias. Our conclusions about the recent equilibrium of the galaxy are therefore robust against the LTE approximation.

5.3.4 SFR

Observational studies show that the Milky Way has sustained a SFR until recent times which is no less than an order of magnitude below its maximal value. Though the exact value of the Milky Way SFR is contested, the general consensus is that it is in the region $1-2 M_{\odot}\ \mathrm{yr}^{-1}$ (Mor et al., 2019; Chomiuk & Povich, 2011; Robitaille & Whitney, 2010; Aumer & Binney, 2009).

5.3.5 Observational Summary

Usually models are fit to the data with the goal of optimising to a best-fit parameter set, which is then prone to systematic biases. Following the motivation presented in the introduction to Chapter 3, we instead walk a different path, where we instead try to falsify classes of models, based on their ability to reproduce a minimum set of constraints, which we draw from the observational evidence:

1. **[Eu/Fe] reaches ~ 0.4 dex at [Fe/H] ~ -1.**
   The average cannot exceed 0.5 dex or fall below 0.2 dex
2. **[Mg/Fe] reaches ~ 0.35 dex at [Fe/H] ~ -1.**
   The average cannot exceed 0.45 dex, or fall below 0.2 dex
3. **[Eu/Mg] is constant throughout.**
   Despite a large scatter, the average remains ~ 0.1 dex
4. **Galactic chemistry is (almost) in equlibrium**
   After 7 Gyrs, deviations from chemical equilibrium must be small.
5. **Star formation has continued until late times**
   A model must maintain enough cold gas for star formation

5.4 Models

In all our models we make the simplifying approximation that there are only 3 sources of nucleosynthesis for r-process elements: a small contribution from CCSN[3], NSM, and Collapsars.

[3]We use CCSN as a catchall term for all processes for which the yield rate is strongly correlated with the current SFR at all points in history, and hence, via our calibration procedure (§5.4.2), includes secondary sources such as magneto-rotational supernovae ('hypernovae').

Whilst other sources of r-process may exist and be important for explaining the abundance of individual stars, we expect the majority of such sources to be of subdominant importance in the case of the overall r-process trends in the galaxy. In the case of BH-NS mergers, we note that their formation mechanisms and timescales are so similar to NS-NS mergers that we can consider them part of the same process, though we note from Pannarale & Ohme (2014) that for $M_{\rm BH} \gtrsim 14 M_{\odot}$, the vast majority of NS companions will be swallowed without tidal disruption, and hence without r-process ejecta.

5.4.1 Simple Argument

Before deriving our analytical chemical evolution models, it is instructive to first discuss the qualitative appearance of Collapsar-dominated models in both the [Eu/Fe]-[Fe/H] and [Eu/Mg]-[Fe/H] planes.

In agreement with previous works (i.e. Si19), we parameterise r-process yields from collapsars (see eq. 5.2) with a constant synthesis rate for metallicity $Z < Z_{\rm c}$. This assumption guarantees that the dominant collapsar pathway produces a plateau in both the [Eu/Fe] and [Eu/Mg] planes. These approximations are justified by comparison with the data in Figure 5.1, which shows a plateau in both planes at early times.

Since the SNIa channel opens with a time delay and produces mostly iron, the plateau in the [Eu/Fe] plane will be interrupted, producing the sharp dip in the [Eu/Fe] abundance (the knee) as the iron production increases. The [Eu/Mg] abundance, however, will exhibit only minor changes (due to metallicity variations or the subdominant Mg yield from SNIa), matching the observations of Figure 5.1.

When the background metallicity reaches $Z \sim Z_{\rm c}$, the Eu synthesis rate will drop as collapsars cease to be formed. Iron and Magnesium production, however, undergo no such transition. This will manifest a 'knee' in both the [Eu/Fe] and [Mg/Fe] planes, with both abundance tracks dropping significantly from their previous values. However, observational evidence for such a feature is lacking. Figure 5.1 does not show a second knee in the iron plane, and nor is there evidence of a significant drop in [Eu/Mg] abundance – in fact, [Eu/Mg] remains almost flat across the entire space.

This simple argument already appears to put severe constraints on the dominance of collapsars in the production of r-process elements. Note, however, that whilst these naive constraints hold in general, they can be modified, e.g. by radial migration (Tsujimoto & Baba, 2019), the variations in the SFH, and the cooling of hot gas.

5.4.2 Analytical Model

Our purpose here is to explore the full possible parameter space of collapsar contributions to examine which regions are ruled out, not to point to some plausible solutions. This necessitates a rapid, streamlined model which can evaluate the chemical histories for billions of combinations of possible parameters.

It was for this purpose that SACEM, the Simple Analytic Chemical Evolution Model was developed. We use this model to track the enrichment of three elements representative of three key groups: the alpha elements (Mg), the iron peak (Fe) and r-process (Eu). The general form of SACEM is discussed in detail in Chapter 3, we discuss here only the modifications required for including collapsars into the prexisting process network.

Modelling r-Process Synthesis

Collapsars events are highly sensitive to the properties of the progenitor star (mass, M, metallicity, Z, and ZAMS rotation speed, v), hence Eq. (2.13) suggests that the rate at which collapsars synthesis r-process material is given by a weighted integral of the r-process yield $Y_r(M, Z, v)$ of each event from such a progenitor and the rate at which these progenitors undergo a collapsar event, $R_{\rm coll} = \mathrm{d}N_{\rm coll}/\mathrm{d}t\mathrm{d}M\mathrm{d}Z$

$$y_{\rm r}(t) = \int_0^\infty \mathrm{d}M \int_0^1 \mathrm{d}Z \int_0^\infty \mathrm{d}v \; Y_r(M, Z, v) R_{\rm coll}(M, Z, v, t) \tag{5.1}$$

In §5.2.2 we identified several problems barring us from formulating Y_r and $R_{\rm coll}$ from first principles, so we must instead opt for simple approximations to Eq. (5.1). Consistent with the approach of Si19 we assume that the yields from individual events are a constant, and that collapsars are a subset of CCSN events such that a fraction $\xi_{\rm coll} a(Z, Z_{\rm c})$ of all stars undergoing CCSN meet the criteria to become collapsars. The constant $\xi_{\rm coll}$ accounts for the mass and rotational constraints, whilst $a(Z, Z_{\rm c})$ is the metallicity suppression function which obeys $a(Z, Z_{\rm c}) = 0$ for $Z > Z_{\rm c}$. Therefore:

$$y_{\rm r}(t) \approx \xi_{\rm coll} \bar{Y}_{\rm coll} \times R_{\rm CCSN}(t) \times a(Z(t), Z_{\rm c}) \tag{5.2}$$

Here $\bar{Y}_{\rm coll}$ is the constant characteristic yield of a single collapsar event, $R_{\rm CCSN}(t)$ is the rate of CCSN events at the time t and $Z(t)$ is the metallicity of the star-forming gas at a time t.

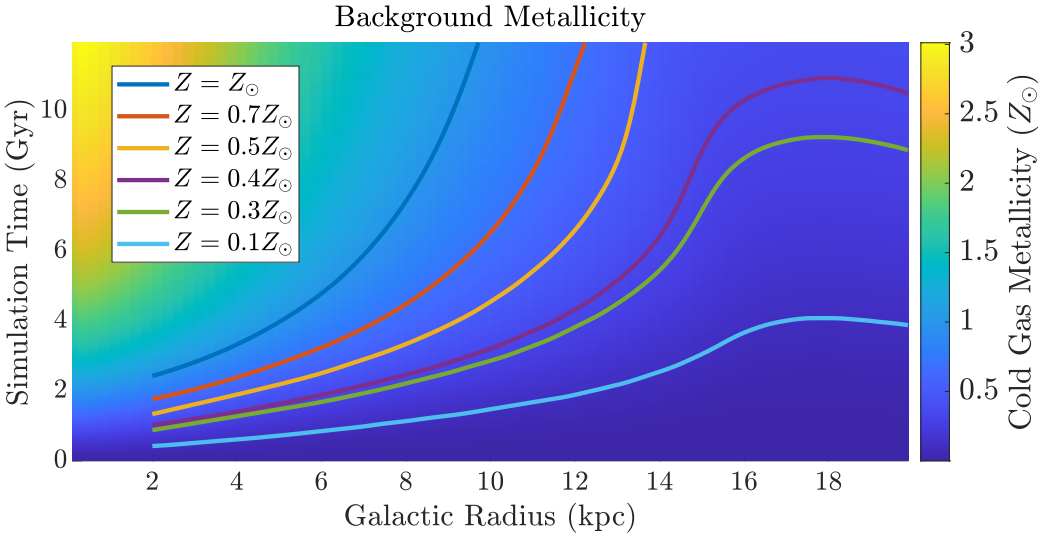

Figure 5.3: A plot showing the background metallicity evolution from the best-fit RAMICES II simulation. Superposed onto this are contours of constant metallicity, which are used to derive the critical cutoff time in SACEM

Metallicity Decoupling

SACEM tracks the abundance of individual chemical species, but not the overall metallicity. However, it is known that the yields of single elements can show strong and individual metallicity dependencies (Maeder, 1992; Chieffi & Limongi, 2004), making $Z(t)$ a functional of its past evolution which is impossible to evaluate analytically. Although iterative approaches are possible, this detracts substantially from the otherwise entirely analytical nature of SACEM.

For simplicity, we therefore fix $Z = \mathcal{Z}(t)$, where $\mathcal{Z}(t)$ is an externally imposed, monotonic function of time. Whilst unappealing, we note from Chapter 3, the only explicitly metallicity dependent part of our model is the yield of collapsar events, Eq. (5.2), which 'turns off' at $Z = Z_\mathrm{c}$. Utilising the monotonicity of $\mathcal{Z}(t)$, we therefore instead suppress collapsar contributions when $t = \tau_\mathrm{coll}$, such that:

$$\mathcal{Z}(\tau_\mathrm{coll}) = Z_\mathrm{c} \tag{5.3}$$

This is equivalent to encoding collapsar turnoff as occurring at a given *time*, instead of a given *metallicity*. We justify this approach further by noting from Casagrande et al. (2011), Schönrich & McMillan (2017) and Haywood (2008), we have a good understanding of the metallicity of the galaxy, and hence the relationship $\mathcal{Z}(\tau_\mathrm{coll}) = Z_\mathrm{c}$ is well constrained by observational data.

Since it is calibrated to precisely these observations, we choose $\mathcal{Z}(t)$ to be the function recovered from an instantiation the full RAMICES II simulation (§4) which, as per SB09, replicates many aspects of the historic enrichment profile. The evolution of $\mathcal{Z}(r, t)$ from

the best-fit RAMICES II model is shown in Figure 5.3. Since our data is mostly from the solar neighbourhood, we evaluate this function at the solar radius, which is well represented for $0 < t < 14$Gyr by the polynomial:

$$\mathcal{Z}(t \text{ Gyr}) = Z_{\odot}\left(0.000591t^3 - 0.0205t^2 + 0.255t - 0.00459\right) \tag{5.4}$$

Deviating from the Metallicity Enrichment History

Varying the properties of our model galaxies will cause the subsequent chemical history to diverge from that used to generate Eq. (5.4). This is potentially highly problematic for us, as we are now imposing $\tau_{\rm coll}$ as a model parameter, instead of the physically meaningful value $Z_{\rm c}$.

It is possible, therefore, that a model successfully replicates the r-process enrichment history of the Milky Way, and has a true cutoff metallicity Z_t such that $Z_t < 0.3Z_{\odot}$. However, because we have not updated $\mathcal{Z}$ for this model, we find that $\mathcal{Z}(\tau_{\rm coll}) > 0.3Z_{\odot}$, and we would reject this model as 'unphysical'.

However, as noted above, we have strong constraints on what a physically meaningful $Z(t)$ can be. Since RAMICES II (and hence Eq. 5.4) is calibrated specifically to reproduce this information, the observational evidence constrains how much $\mathcal{Z}(\tau_{\rm coll})$ and Z_t can differ. The 'falsely rejected' models above are explicitly those which lie in tension with this data.

The converse is also true: we may generate models with $Z_c(\tau_{\rm coll}) < 0.3Z_{\odot} < Z_t$. These models would pass the test of §5.2.1, but are unphysical. However, due to our focus on negative inference, accepting unphysical models merely limits our conclusions to upper bounds on the contribution of collapsars. Of far greater concern to us is rejecting physical models, which we must ensure we do not do.

Yield Calibration

As with any SACEM instantiation, we must choose the calibration points to determine the unknown prefactors discussed in §3.4. As we are generating a model capable of tracking 3 elements (Iron, Magnesium and Europium) across 4 processes (CCSN, Collapsars, NSM and SNIa), we nominally require 12 calibration points.

However, we make the simplifying set of assumptions that a) collapsars and NSM only produce Eu, b) SNIa only produce iron in meaningful quantities. We therefore require only 6 calibration points. We also make the choice to, wherever possible, calibrate the data against the present day values as this prevents overt systematic errors entering through our calibration efforts.

Table 5.1: The calibration points used in the SACEM to fix the yields of CCSN, SNIa, Collapsar and NSM events. The final constraint, χ, has no nominal value as it the parameter we wish to investigate: the contribution of collapsars to the total r-process budget.

Observable	Symbol	Nominal Value
[Fe/H] (at $t = t_{\rm today}$)	$\mathcal{F}$	0.1 dex
[Mg/Fe] (at $t = t_{\rm today}$)	$\mathcal{M}_\infty$	0 dex
[Mg/Fe] (at $t = \tau_{\rm SNIa}$)	$\mathcal{M}_0$	0.35 dex
[Eu/Fe] (at $t = t_{\rm today}$)	$\mathcal{E}$	0 dex
s Process fraction (at $t = t_{\rm today}$)	σ	0.02
Collapsar fraction (at $t = t_{\rm today}$)	χ	-

Table 5.1 shows the chosen calibration points. We note that the key parameter of interest for us is χ, the contribution of collapsars to the final time Europium abundance.

Resolving Contribution Ambiguity

χ, the collapsar contribution fraction, has two distinct definitions:

1. The mass ratio of collapsar-sourced europium to the total europium within the cold, star forming gas at a time t

2. The total mass ratio of all collapsar-sourced material ever produced to the total amount of europium ever synthesised (i.e. including that which has been subsequently locked up in stars, stored in the hot gas reservoir or folded into black holes)

These two definitions can diverge significantly: material produced early on is either lost to the CGM or locked up in stars/stellar remnants, so the definition (i) sets weights much more towards more recent enrichment. With that significant difference, we note that the qualitative picture and our inferred conclusions remains unaltered between the two choices and we will use option (i), defining χ as the collapsar contribution to the current europium in the cold gas phase. The s-process contribution, σ and NSM contribution, κ are defined similarly.

5.5 Search For Models

5.5.1 Varying $\tau_{\rm coll}$ in Analytical Models

As an initial experiment (and to confirm the intuition developed in §5.4.1), we observe the effects of modifying the collapsar cutoff time, $\tau_{\rm coll}$ on SACEM chemical histories.

The investigated model was calibrated against the iron and magnesium distributions – the only constraints placed on the Europium abundance is that, at simulation end, [Eu/Fe] = 0 and $(\chi, \sigma) = (0.98, 0.02)$, i.e. be 98 per cent collapsar in origin, with no NSM contribution.

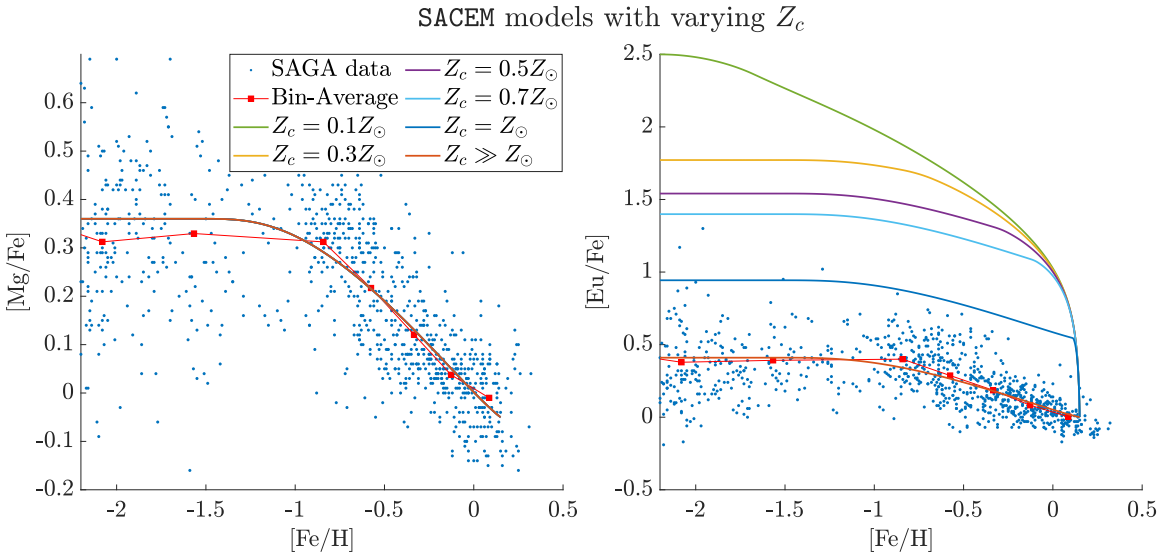

Figure 5.4: A series of chemical histories generated by SACEM for a collapsar-dominated galaxy, for various values of the collapsar cutoff time $T = \tau_{\rm coll}$ such that $\mathcal{Z}(\tau_{\rm coll}) = Z_{\rm c}$. The time/metallicity relationship is calculated from Figure 5.3.

Using the $Z_{\rm cut} - \tau_{\rm coll}$ relationship of §5.4.2, we generate 6 chemical histories corresponding to cutoff metallicities of 0.1, 0.3, 0.5, 0.7, 1 and 20 $Z_\odot$, with the final cutoff being deliberately large so as to take place at an infinite time in the future (we refer to this model as '$Z_{\rm c} \gg Z_\odot$').

The resultant chemical histories, shown by the solid lines in Figure 5.4, differ only in the normalisation of their Eu channels, so the behaviour in [Mg/Fe]-[Fe/H] (left panel) is unaffected. However, in the right panel we see the impact in the [Eu/Fe]-[Fe/H] plane: since the models are tethered to reach [Eu/Fe] ~ 0 at the end of the simulation, SACEM compensates for the truncated Europium production by increasing the collapsar effective yield $\bar{Y}_{\rm Coll,\ Eu}$, resulting in unreasonably large [Eu/Fe] values at early times: these models cannot be accurate depictions of our galaxy.

5.5.2 RAMICES II & Radial Structure

This is reinforced by Fig 5.5, in which we also run the same set of constraints on the RAMICES II simulation, verifying that these initial results hold up in the full multi-zone model. Unlike SACEM, the models used to produce Figure 5.5 are calibrated to match the early-time paths, and so we see these curves plunge below [Eu/Fe] = 0, rather than their early-time abundances shooting up.

We see that the RAMICES II models suffer from a greater differential between their early time abundances and their final values than SACEM did. This indicates that the approximations in SACEM make it more lenient with regards to gas depletion timescales, or late-time Eu

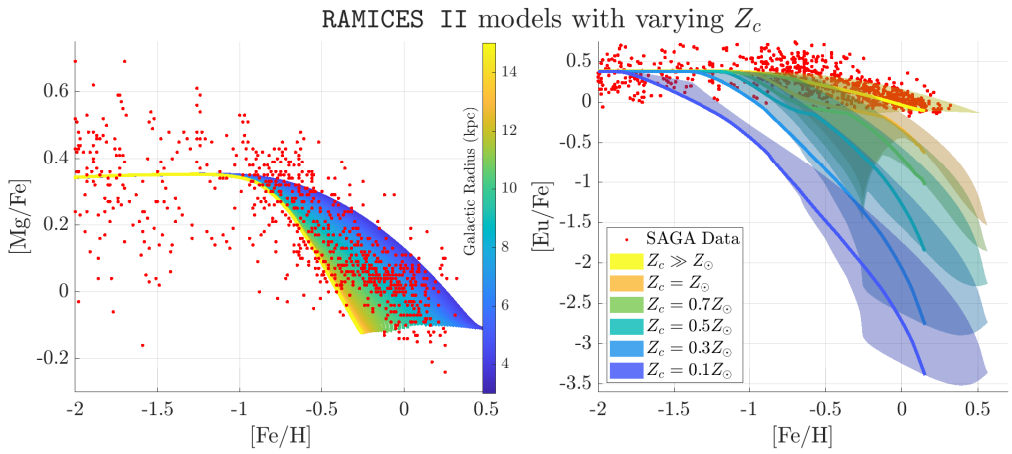

Figure 5.5: As with Figure 5.4, but generated using the RAMICES II simulation. In the left panel, the colour is used to denote the galactic radius, as varying Z_c does not alter the history of [Mg/Fe]. In the right panel, the shaded region shows the same radial distribution as in the left panel with the thick line tracking the chemical history at the solar radius.

abundances sourced from the hot gas phase. As a result, SACEM will provide less stringent constraints for the late Eu production than the full simulation. This works in our favour, as it means any regions of parameter space that we are able to exclude is likely to be a lower bound on the size of the excluded region.

As a multi-zoned model, RAMICES II allows us to investigate the guiding-centre radius (R_g) behaviour of the abundance distributions, shown in Figure 5.6. The rise in [Eu/Fe] seen from $R_g = 3$ to $R_g \sim 10$kpc would be what we expect from the impact of the galactic metallicity gradient: the inner galaxy hits $\mathcal{Z} = Z_c$ first, and so r-process abundances begin to plummet first at lower R. The inner radii have been Europium-deprived for longer than mid-disc radii, and so have a lower [Eu/Fe] ratio – hence the positive [Eu/Fe] gradient as you move out through the disc.

For $Z = 0.5,\ 0.7$ and $1Z_\odot$, this pattern in [Eu/Fe] is interrupted by an 'arch' at $R_g > 10$kpc – an effect induced by the inflow/accretion prescriptions of RAMICES II, namely that the inner disc outflows polluted the IGM with low-[Eu/Fe] gas. Whilst such an effect is at least partly physical, the exact position of the 'arch' and the resulting outer-galaxy gradients are a function of, for example, our IGM mass. In leiu of a full physical investigation of this behaviour, we consider this a quirk of the particular calibration used here. The behaviour inwards of the solar radius, however, is robust against the CGM prescription chosen and is a necessary consequence of the metallicity gradient of the galaxy. Though a lengthier discussion of the radial behaviour of [Eu/X] is beyond the scope of this work, we note that the radial distribution of europium, as inferred from cepheids in Luck & Lambert (2011) is inconsistent with a such a drastic

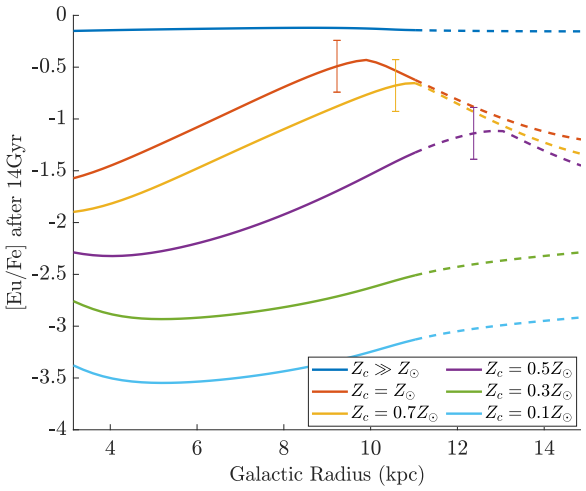

Figure 5.6: The final result of the simulations of Figure 5.5 at $t = t_{\rm end}$ explicitly projected into radial coordinates. The lines are dashed from the point at which the 50 per cent of the total metallicity budget was accounted for by the CGM. Where present within the galactic disc, the vertical lines denote the physical position of the 'critical radius': the point where $\mathcal{Z}(r, t_{\rm end}) = Z_{\rm c}$

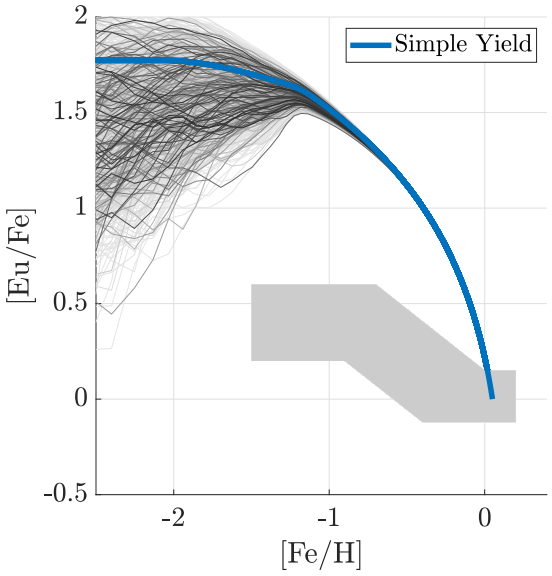

Figure 5.7: 10,000 SACEM instantiations in which the collapsar yield function is allowed to take on arbitrary forms are shown, relative the the grey 'allowed region'. All models have $\tau_{\rm coll} = 1.4$Gyr, $\chi = 1$ and otherwise use the best fit parameters of Figure 5.4. We see that within a fraction of a dex of [Fe/H] after collapsar suppression, all models converge on a single path through [Eu/Fe] space, determined solely by the depletion rate of Eu from the ISM.

increase in [Eu/Fe] with galactic radius. This further emphasises the conclusions drawn in §5.3.3: there is little evidence for a large change in galactic chemistry in the recent past – any changes to the galactic chemistry must have been far enough in the past for the radial mixing of the galaxy to smooth out the emergent patterns.

This could be leveraged to provide a much greater constraint on the success of our models – the models of Figure 3.1 undergo their collapsar cutoff during the thick/thin disc transition, and so we would expect to see the resultant radial patterns strongly rule out these model, even if properly calibrated. The restrictions imposed on our SACEM instantiations are far from the strictest we can generate, further emphasising that our conclusions are the upper bound of collapsar contributions.

5.5.3 Switching Dominance

In many of the models studied, the 'dominant' source of Eu changes drastically with time. Such a model was shown in Figure 3.1: even though the final contribution of collapsars to the Eu budget is $\sim 10\%$, at early times it accounted for more than 90 per cent of the synthesised Eu, of which more than 50 per cent was found within the hot gas phase. This collapsar-dominant

Table 5.2: A list of the named parameters within SACEM, their physical definitions and interpretations, and the bounds placed on them in the three primary constraint sets. Note that the mixed set, SET M, uses the constraints from the Viable set with the exception of the 'SFR Parameters', which it draws from the unconstrained set. For an algorithmic 'recipe' for how to incorporate these parameters into SACEM, see §3.5

	Quantity	Description	Units	Unconstrained		Weak Constraints		Viable Constraints	
				Lower	Upper	Lower	Upper	Lower	Upper
SFR Parameters	M_0	Initial Cold Gas Mass present in galaxy	$10^9 M_\odot$	0.001	200	0.01	10	0.1	1
	M_1	Primary Infall Mass (thick disc)	$10^9 M_\odot$	10^{-6}	20	1	10	1	10
	M_2	Secondary Infall Mass (thin disc)	$10^9 M_\odot$	10^{-6}	200	10	70	20	70
	β_1	Primary Infall Frequency	Gyr^{-1}	0.1	100	0.1	10	0.333	10
	β_2	Secondary Infall Frequency	Gyr^{-1}	0.001	1	0.02	0.1	0.033	0.1
	ν_{SFR}	SFR frequency	Gyr^{-1}	0.01	8	0.05	5	0.05	5
	μ	Stellar Death Frequency (rate = μM_*)	Gyr^{-1}	10^{-4}	1	0.001	0.5	0.001	0.1
	δ	Outflow/heated mass per stellar mass formed	-	10^{-5}	5	0.01	2	0.01	1.5
	M_c/M_*	Final ratio of cold mass to stellar mass	-	10^{-5}	1	0.05	0.25	0.05	0.15
Thermal Parameters	f_{CCSN}	Fraction of CCSN mass put into hot phase	-	0.001	0.999	0.6	0.999	0.7	0.999
	f_{NSM}	Fraction of NSM mass put into hot phase	-	0.001	0.999	0.3	0.999	0.3	0.999
	f_{SNIa}	Fraction of SNIa mass put into hot phase	-	0.001	0.999	0.6	0.999	0.7	0.999
	f_{Coll}	Fraction of Collapsar mass put into hot phase	-	0.001	0.999	0.6	0.999	0.7	0.999
	λ_{CCSN}	Cooling frequency for ejecta from CCSN	Gyr^{-1}	0.04	10	0.4	2.5	0.5	1.5
	ℓ_j	3 independent parameters. $\lambda_j = \lambda_{\mathrm{CCSN}}(1+\ell_j)$ for $j \in$ {CCSN, SNIa, NSM}	-	-0.999	0.999	-0.2	0.2	-0.1	0.1
Temporal Parameters	τ_{SNIa}	SNIa delay time	Gyr	0.001	2	0.05	1	0.1	0.6
	τ_{NSM}	NSM delay time	Gyr	10^{-5}	2	10^{-4}	0.6	10^{-3}	0.1
	ν_{SNIa}	SNIa DTD Decay frequency	Gyr^{-1}	0.05	50	0.05	25	0.1	15
	ν_{NSM}	NSM DTD Decay frequency	Gyr^{-1}	0.05	50	0.05	25	0.05	25
	Δ	Collapsar turnoff width	Gyr	0.1	15	0.1	15	0.1	15
Chemical Parameters	F_{mod}	Lockup-SFR modification factor	-	0.3	1	0.3	1	0.3	1
	X	Hydrogen Fraction at simulation end	-	0.5	0.9	0.65	0.75	0.68	0.72
	$\mathcal{F}$	Final [Fe/H] value	dex	-0.5	1	0	0.5	0.05	0.3
	$\mathcal{M}_0$	Thick disc value for [Mg/Fe]	dex	0.2	0.6	0.3	0.5	0.3	0.4
	$\mathcal{M}_\infty$	Final value for [Mg/Fe]	dex	-0.3	0.2	-0.1	0.1	-0.1	0
	$\mathcal{E}$	Final value for [Eu/Fe]	dex	-0.3	0.2	-0.1	0.1	-0.1	0.05
	σ	Fraction of europium produced through the s process at simulation end	-	10^{-8}	0.2	10^{-8}	0.1	10^{-8}	0.05
	κ	Fraction of europium produced through NSM $\kappa \equiv 1-\sigma-\chi$	-	-	-	-	-	-	-
	τ_{coll}	Time of final collapsar event	Gyr	0	16	0	16	0	16
	χ	Fraction of Europium produced in collapsars at simulation end	-	0	1	0	1	0	1

early time contribution has been observed even as $\chi \ll 0.01$. Even if collapsars are negligibly responsible for r-process synthesis at late times, Figure 3.1 shows that they may have been important contributors at early times. This model class would be disfavoured due to the Sn08 fingerprint indicating a monolithic r-process source, however, if we suggest that the high-[Eu/Fe] stars sampled to form the 'fingerprint' may not be representative of the bulk of stars at this early time (which did source their Eu from collapsars), this tension is alleviated.

5.5.4 Impact of the Simple Yield Approximation

In the absence of any observational constraints on the collapsar yield function $Y_{\rm coll}$ we follow the literature (i.e. Si19) by adopting the simplest possible form: that of Eq. (5.2). One could argue that this functional form is inappropriate, or otherwise corrupts any conclusions we might draw from our models. However, several factors justify that approach.

The behaviour of Figs. 5.4 and 5.5 shows that even if the yield functions are not well-approximated by a simple function, the overall range in [Eu/Fe] is a necessity that cannot be mended by introducing an additional variation into $Y_{\rm coll}$. Once collapsars stop forming, the [Eu/Fe] ratio drops at a fixed rate which depends only on the gas depletion timescale, and independent of any chosen form of the yield function. The only way to ensure that [Eu/Fe] ≈ 0 at simulation end is by adding more Eu before collapsars die off. However, because collapsars are only active for a short time, no matter the functional form of the yield, increasing the Eu abundance pushes [Eu/Fe] above our 'reasonable domain', and the model would be ruled a failure. We demonstrate this principle in Figure 5.7, in which we replace our simple yield function with a positive random-walk function $y_{i+1} = |y_i + R|$ where R is a random number in the range [-1,1], allowing for arbitrary functional forms of $y_{\rm coll}(Z)$. We see that no matter the functional form, after the collapsar suppression begins, all of the yields converge rapidly to a single path through [Fe/H]-[Eu/Fe] space. This path is determined solely by the required amount of Eu in the galaxy at the time that collapsars are suppressed, and hence is independent of the collapsar yield function: i.e., the models care how much Eu there is in the galaxy at $\tau_{\rm coll}$, not how it got there.

The post-turnoff behaviour of these models would therefore be very well represented by a 'simple'-type model which was tuned to produce the same Eu mass at collapsar turnoff as the varying model. For a given set of model parameters if a simple-type yield would cause the model to be ruled incompatible with the observed evidence at any time after collapsar suppression, then so would the corresponding model with a varying yield function.

By eliminating model classes based on their simple-yield approximations, we cannot be falsely eliminating any physically meaningful models, so we conclude that our results are robust against the impacts of the approximation of the simple form of $y_{\rm coll}$.

5.5.5 Expanding the Search

In these initial forays we were altering a single parameter in a single 'best fit' model, so one could argue that the search missed the right combination of parameters that allows for a high collapsar contribution despite a reasonable cut-off metallicity.

The challenge is that such a set of parameters must not simultaneously make, for example, the [Mg/Fe] histories unrealistic: we must ensure that any alterations produce a simultaneously realistic model in multiple chemical planes.

'Success' & Viable Models

Following up on our summary in §5.3.5 we now define the criteria for a viable chemical evolution model. We emphasise that, in line with our efforts to falsify model classes rather than find a best fit model, the criteria developed here do not imply that a model is fully physical if it fulfils the criteria, but models that breach the criteria are clearly in contradiction to the empirical evidence.

We define a galactic chemical history as 'viable' if the resulting Tinsley curves for [Mg/Fe], [Eu/Fe] and [Eu/Mg] lie between the two black two curves in the corresponding panels of Figure 5.1. Whilst a good model should also reproduce the observed distributions not just pass within its range, in keeping with our approach of negative inference, we instead choose to have room to be generous and yet still constrain the models. All three planes must simultaneously meet this criterion for the model to count as successful.

We make the intentional choice that the constraints only apply for [Fe/H] > -1.5. Though this weakens our constraints, it gives us numerous advantages: we limit ourselves to the strictly non-stochastic regime, avoid halo contamination and as per 5.5.4, we avoid eliminating classes of models which have high variability at low metallicities.

In addition, we study the effect of introducing a constraint on the time evolution of the models in concordance with the observational constraint in Figure 5.2. Such Gradient-Constrained models are only considered successful if:

$$\left|\left\langle \frac{\mathrm{d[Eu/Fe]}}{\mathrm{d}t} \right\rangle\right| \leq 0.01 \text{ dex Gyr}^{-1} \tag{5.5}$$

where this is measured over the final 2 Gyrs of the simulation, and is in agreement with Figure 5.2.

Finally, a model galaxy must be capable of a sustained rate of star formation, even at late times. In GCE models it is common to constrain this through the star formation efficiency, represented by (up to a factor of order $\nu_{\rm sfr}$) the quantity $M_{\rm c}/M_{*}$. For the Milky Way this value can be measured (i.e., McKee et al. 2015), indicating that we should constrain this value to ~ 0.1.

We reemphasise that we have left the acceptance criteria intentionally generous: we want to find which models are excluded by failing to meet even these lax criteria.

5.5.6 Variable Selection for SACEM

We perform a Monte-Carlo exploration of the high-dimension parameter space by randomly and independently selecting the model parameters (except χ, the collapsar contribution ratio, and $\tau_{\rm coll}$, the collapsar cutoff time) from between bounds determined by the imposed constraints (see Table 5.2). For each random set of parameters, we then perform a sweep of χ-$\tau_{\rm coll}$ space, such that each instantiation is evaluated uniformly across this space. Note that the enforcement of the temporal gradient constraint (Eq. 5.5) plays a special role, as we perform all simulation runs both with and without this constraint.

We call the different constraint sets in Table 5.2 the UNCONSTRAINED set (denoted U), WEAK (W) and VIABLE (V). As indicated by the name, in Set U, parameters are allowed to take almost any value, with essentially no constraint or regard for their physical reality. Naturally, the 'unconstrained' set is not truly unconstrained. Instead we mean that the boundaries imposed reflect our desire to search 'interesting' regions of parameter space within finite computation time, and do not meaningfully eliminate any kinds of physically interesting galaxies we might care to consider. We include this set to act as a null hypothesis – that without physical constraints, models can be made to fit to the data – and hence that any region of parameter space that is eliminated is due to the imposition of physical constraints.

The WEAK set of parameters (W) adds an amount of physical intuition – the order of magnitude of the galactic mass is determined, cooling timescales are of an order close to what we might expect, and so on. These constraints encapsulate the behaviour of most galaxies, but do not uniquely identify the Milky Way. In contrast, the VIABLE set imposes the minimum requirements to match the properties of the Milky Way.

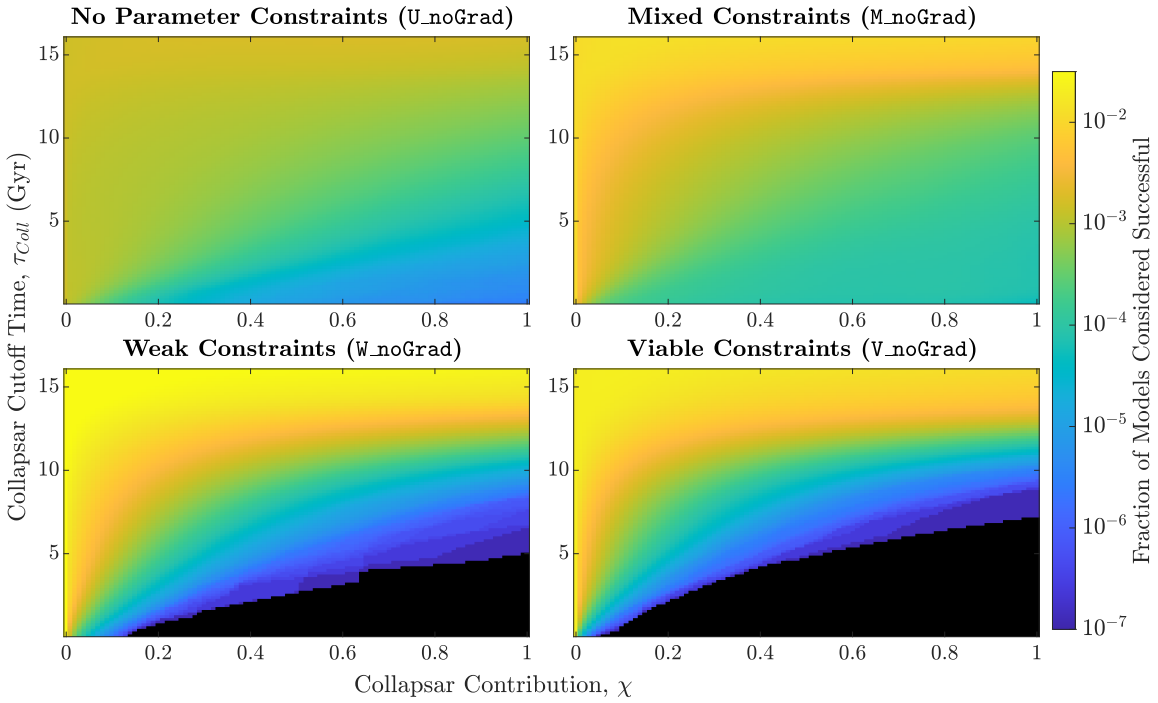

Figure 5.8: The results of the SACEM gridsearch for models where no gradient check was performed, and where the fraction of europium, χ, was measured from the present-day cold gas reservoir. Black regions are those in which no successful models (as determined by the success conditions of §5.5.5) could be found. Note that the $\tau_{\rm coll}$ axis does not extend exactly to $\tau_{\rm coll} = 0$ (the region where collapsars are never active), as $\chi \neq 0$ is nonsensical if collapsars are inactive for all history. Therefore $\tau_{\rm coll} \geq 0.16$Gyr.

A fourth variant, the MIXED model (M), was also studied, which uses the VIABLE constraints for all parameters except those relating to the SFR, which are treated as UNCONSTRAINED.

5.6 Results

5.6.1 SACEM Search

For each of the constraint sets of §5.5.6 we randomly generated 10^7 parameter sets, which were then each evaluated across a 101×101 grid in $\chi - \tau_{\rm coll}$ space, for a total of $\sim 10^{11}$ models.

Due to the high density of generated models, we may therefore be confident that the regions in which no models could be found (the 'exclusion zone') are genuine forbidden regions of parameter space. This inference is strengthened if the exclusion zone are contiguous: whilst statistical fluctuations may alter the position of the boundary, the contiguous nature implies that regions a significant distance from this boundary are robustly excluded.

Without Temporal Gradient Constraints

We first consider those models which did not use the condition of Eq. (5.5) to evaluate the success of a model. The density of successful models is shown in Figure 5.8, with regions coloured black denoting the points where no successful model could be found.

As expected, for the Unconstrained (U) and Mixed (M) constraints, we are able to find allowed models for all values of the collapsar contribution χ and the cutoff time $\tau_{\rm coll}$. However, the density variation for models with large-χ but low $\tau_{\rm coll}$ (i.e. lower right of each panel) shows that such models are disfavoured even with these lax constraints: in both the Unconstrained and Mixed investigations, NSM-dominated models (left of each panel) were favoured by a factor ~ 500 over the corresponding high-χ-low-$\tau_{\rm coll}$ models.

With the introduction of the only Weak constraints, we see the formation of an exclusion zone in the high-χ, low-$\tau_{\rm coll}$ region: the Weak constraints fail to find any models which are collapsar dominated ($\chi \sim 1$) for $\tau_{\rm coll} < 6$Gyr, and limits the Milky Way to $\chi > 0.65$ for $\tau_{\rm coll} < 3$Gyr. The exclusion zone is not qualitatively changed by the introduction of the Viable constraints, but the size increases such that collapsar dominated models are prohibited for $\tau_{\rm coll} < 7$ Gyr,$\chi > 0.5$ is prohibited for $\tau_{\rm coll} < 4$Gyr and $\chi > 0.2$ is prohibited for $\tau_{\rm coll} < 2$Gyr.

This immediately brings us into direct tension with the theoretical models of collapsars. As per §5.2.1, no theoretical models yet allow for events above $Z = 0.3Z_{\odot}$, but following the discussion in §5.4.2 and the work of i.e. Casagrande et al. (2011), we know the Milky Way was rapidly enriched early in its lifetime. In other words, $\tau \sim 2$Gyr is the approximate time coordinate that we should associate with the 'maximum theoretical collapsar time'.

These initial results therefore strongly indicate that, under Weak and Viable constraints, the r-process enrichment of the Milky Way cannot be dominated by collapsars events constrained to occur before $Z = 0.3Z_{\odot}$.

However, in §5.4.2, we cautioned against inferences which rely on coupling $\tau_{\rm coll}$ and $Z_{\rm c}$ too tightly, as there is a small amount of flexibility in the relationship due to our decoupling of the background metallicity evolution from the properties of galaxies. The Viable exclusion zone in Figure 5.8 has a boundary at $\sim$ 4Gyr, so between the Monte Carlo search and the decoupling of $Z_{\rm c}$ and $\tau_{\rm coll}$, it is plausible that our results might still allow a collapsar dominated galaxy which we failed to detect: more work is yet needed.

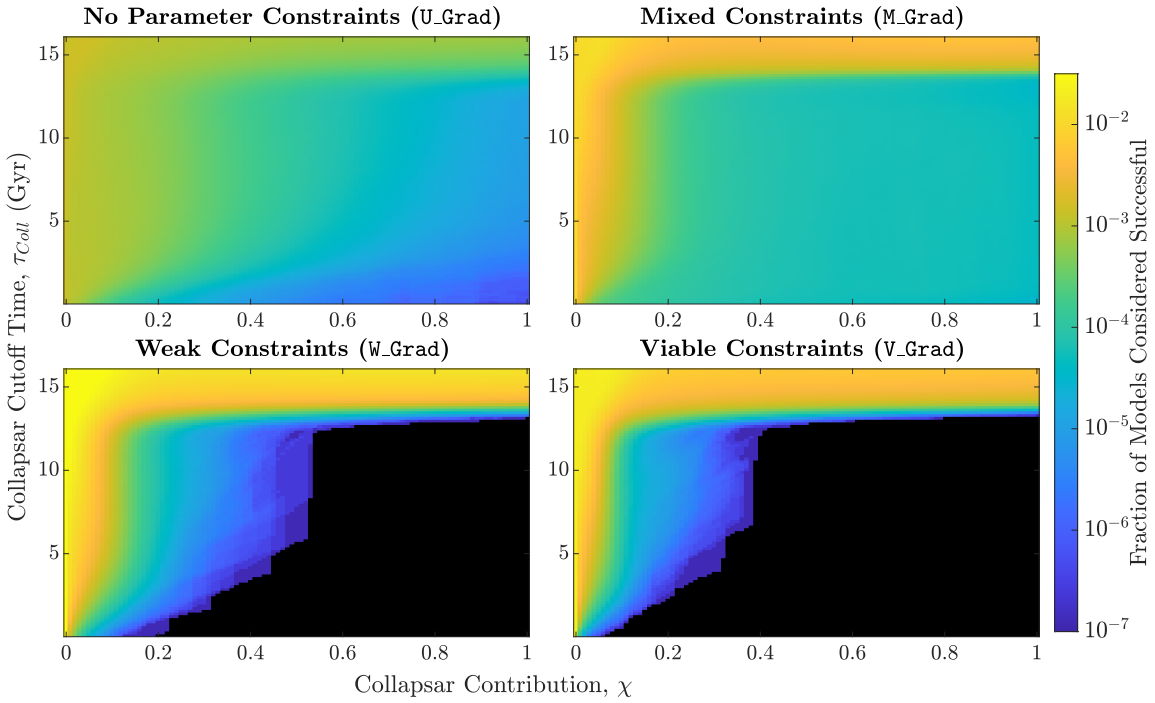

Figure 5.9: As with Figure 5.8, with the inclusion of the constraint that only models with $\left|\left\langle \frac{\mathrm{d[Eu/Fe]}}{\mathrm{d}t} \right\rangle\right| \leq 0.01$ dex Gyr^{-1} can be considered successful.

Temporally Constrained Models

In Figure 5.9 we examine the same set of constraints as Figure 5.8, but with the additional constraint that successful models must obey $\left|\left\langle \frac{\mathrm{d[Eu/Fe]}}{\mathrm{d}t} \right\rangle\right| \leq 0.01$ dex Gyr^{-1} across the final 2Gyr of evolution. We note from Figure 5.2 that this is still a generous constraint.

We see that, once more, models which use Unconstrained and Mixed constraints excluded no regions of parameter space, though as before they favoured low-χ or high-τ_{coll} models by more than a factor of 10^2 and 10^3 respectively.

However, for Weak and Viable models, we see that large regions of space have been declared non-viable. Comparing models W_noGrad and W_Grad (the lower left panels of Figs. 5.8 and 5.9 respectively) we see that the addition of the temporal gradient constraint has expanded the exclusion zone. The region $\chi > 0.6$ (previously excluded for τ_{coll} <6Gyr) is now excluded for all times $\tau_{\mathrm{coll}} < 12$Gyr. This makes the exclusion zone in this region extremely resistant to both statistical errors and errors arising from the imposed metallicity function, as such errors would need to cause errors on the order of 5Gyr, which we do not consider a reasonable expectation.

Similarly, with the addition of the temporal constraints, all successful Viable models were limited to $\chi < 0.4$ for $\tau_{\mathrm{coll}} < 12$Gyr and $\chi < 0.3$ for $\tau_{\mathrm{coll}} < 4$Gyr. This eliminates all models

for which collapsars are either the majority or dominant source of r-process nucleosynthesis, as there is substantial theoretical and observational evidence that collapsars cannot have been occurring more recently than 2Gyr ago.

Again, we emphasise that the vertical expansion of the exclusion zone observed in models W_GRAD and V_GRAD makes these conclusions extremely robust to statistical and $Z_c - \tau_{\rm coll}$ induced errors. We are therefore confident that these regions are truly excluded portions of parameter space.

5.6.2 RAMICES II Search

We performed a similar search on the RAMICES II simulation, though due to computational limitations, the number of models is much smaller. In addition, RAMICES II is more physically motivated than the semi-empirical approach of SACEM, such that there are fewer parameters which we may alter without undermining some other observable property of our galaxy. In particular, the SFR (which is coupled to gas infalls and radial gas motion) is chosen to replicate features of the solar neighbourhood (see SB09). The remaining free parameters which are not fixed to replicate observable properties of the galaxy are therefore:

- The hot-gas injection fraction $f_{\rm CCSN}$ and $f_{\rm NSM}$
- The delay time of SNIa processes, and the fraction of white dwarf remnants which can undergo SNIa events
- The fraction of synthesised material lost to the CGM
- The effective r-process yields $\bar{Y}_{\rm Eu}^{\rm s\text{-}process}$, $\bar{Y}_{\rm Eu}^{\rm NSM}$ and $\bar{Y}_{\rm Eu}^{\rm Coll}$
- The value and width of the metallicity cutoff.

Because the RAMICES II model is inherently much more tightly constrained, we widened the success condition to prevent overfitting of the model. The only success condition applied to the RAMICES II models is that, for the solar radius, the final [Eu/Fe] value must lie within the 'viable domain' of Figure 5.1 (the final [Fe/H] value already being calibrated to the solar neighbourhood).

Note that, unlike SACEM, this simulation does use metallicity for the collapsar cutoff, rather than a metallicity-inferred time. We also note that due to the properties of the RAMICES II simulation, we cannot target a specific final collapsar-contribution fraction χ; we have to generate models with a given set of europium yields and determine χ at simulation end. Hence,

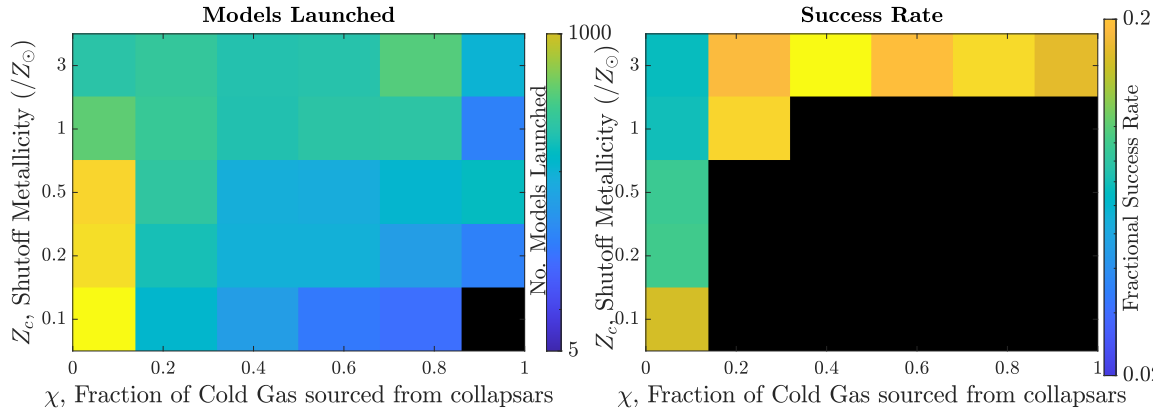

Figure 5.10: The equivalent success plot to those shown in figures 5.8-5.9, generated instead from the RAMICES II simulation. The top panel shows the density of all models which were launched (RAMICES II measures χ at simulation end). The lower panel shows the ratio of successful models to those launched.

unlike the models of the previous section, the models are not produced on a uniform grid of χ, and due to the computational constraints it proved somewhat difficult to even generate a meaningful number of models with $\chi \sim 1$ at low Z_c. In particular, we were not able to generate a single model for which $Z_c = 0.1 Z_\odot$ but $\chi > 0.85$. However, given the contiguous (and large) nature of the final exclusion zone, we do not expect this to impact our final conclusion.

We ran 2000 iterations of the RAMICES II simulation, with the subset of variables drawn from the WEAK sample shown in Table 5.2. The results of this search are plotted in Figure 5.10. Despite the fact that the random parameters were generated with the WEAK model, with no explicitly included gradient consideration and with an even more generous acceptance criteria, RAMICES II produced a grid with an even larger exclusion zone than the temporally-constrained VIABLE model of SACEM.

It seems that the initial results we inferred from Figure 5.5 – that a multi-zone model with more physically motivated parameters suffers a significantly larger drop in [Eu/Fe] than the corresponding SACEM model – continues to hold, thereby indicating that the SACEM results should be interpreted as an upper bound on the maximum collapsar contribution.

5.7 Properties of Successful Models

Figs. 5.8 and 5.9 demonstrate that the metallicity dependence of collapsar events necessarily limits them to sub-dominant contributors to the galactic r-process budget, yet it is more instructive to examine in more detail those models deemed "successful".

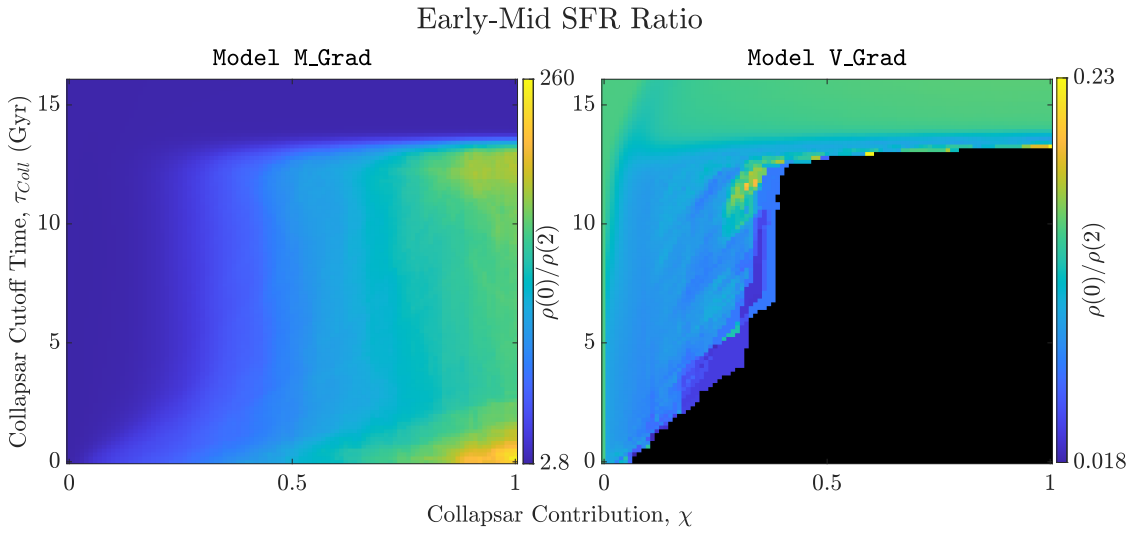

Figure 5.11: The mean ratio $\rho_{\rm sfr}(0)/\rho_{\rm sfr}(2{\rm Gyr})$, a measure of the 'peakiness' of the early time SFR as measured across the M_GRAD and V_GRAD models.

SACEM is unique in its low computational footprint, so we can for the first time systematically scan the full parameter space. This allowed us to entertain non-physical sets of parameters in our search for the minimum required set of constraints: the models generated from the UNCONSTRAINED and MIXED sets, in particular.

This was, primarily, an aesthetic choice: we wished to impose as few additional constraints onto our models as possible. However, by examining the behaviour of the global parameters from the less-constrained models within the $\chi - \tau_{\rm coll}$ plane, we may infer how the successful models fulfil the imposed requirements. This provides three benefits: i) comparing how past studies have been able to match chemical observations, and if this lies in tension with other observables ii) examinations into the limitation to these approaches, i.e. how far we may 'tweak' the parameters of our model whilst remaining physically viable and iii) examining if there are other ways for a model to fit the data beyond the canonical understanding.

5.7.1 Star Formation

The most straightforward way for a model to produce a required chemical history within the specified bounds is with an extremely tightly controlled star-formation rate. Within SACEM, if a model has a highly-peaked early-time SFR, then the large collapsar population could generate a vast amount of Eu over a short period of time. In such models, a true GCE model would also accumulate an equally vast number of metals during this period due to the high SFR, such that collapsars should shut off. However, because the imposed function $\mathcal{Z}(t)$ is insensitive to the particular parameters of our model, the 'true metallicity' of the model

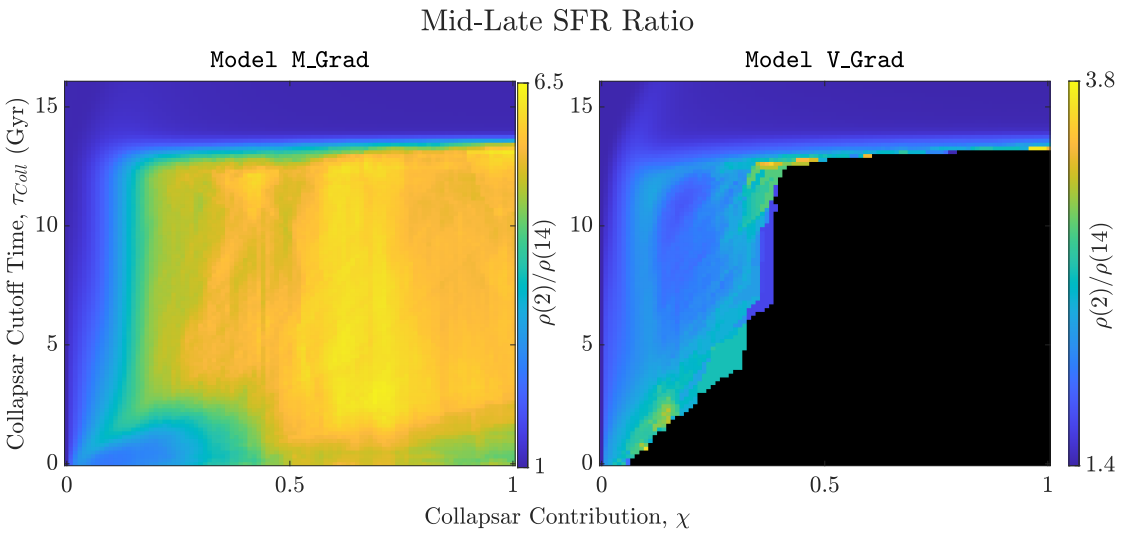

Figure 5.12: As with Figure 5.11 but with $\rho_{\rm sfr}(2)/\rho_{\rm sfr}(14)$: a measure of how quickly the star formation in the galaxy dies down.

diverges significantly from the nominally assumed model, and so collapsars will continue contributing long past the time they should have died out.

This therefore enables models with even tiny values of $\tau_{\rm coll}$ to be considered viable, where a more physically coupled SFR and metallicity evolution would discard these models. We suspect that the vast majority of $\tau_{\rm coll} - \chi$ space accepted under the MIXED regime, but rejected under WEAK and VIABLE models are achieving their success through this method.

Figs. 5.11 and 5.12 justify this claim, they respectively depict the ratio of the initial SFR and after 2Gyr: $\rho_{\rm sfr}(0)/\rho_{\rm sfr}(2{\rm Gyr})$ and the corresponding ratio between 2Gyr and the end of the simulation, $\rho_{\rm sfr}(2{\rm Gyr})/\rho_{\rm sfr}(14{\rm Gyr})$, for two classes of models. The properties of the models in the high-χ-low-$\tau_{\rm coll}$ region of M_GRAD are striking: the search preferentially found models where the initial SFR was more than 2 orders of magnitude greater at simulation start than it was just a few gigayears later, and where the mid-time SFR was on average a factor of 6 greater than the SFR at simulation end, such that there is a total of more than 3 orders of magnitude between the $t = 0$ SFR and the final time SFR. These ratios are orders of magnitude outside the constraints measured for the star formation history of the Milky Way, which typically show a decline in SFR by less or equal to an order of magnitude (Mor et al., 2019; Schönrich & Binney, 2009b; Aumer & Binney, 2009).

By comparison, the V_GRAD model shows no particular bias in either of Figs. 5.11 or 5.12, indicating that large-scale constraints such as the total mass of the Milky Way are reasonable methods to eliminate these models.

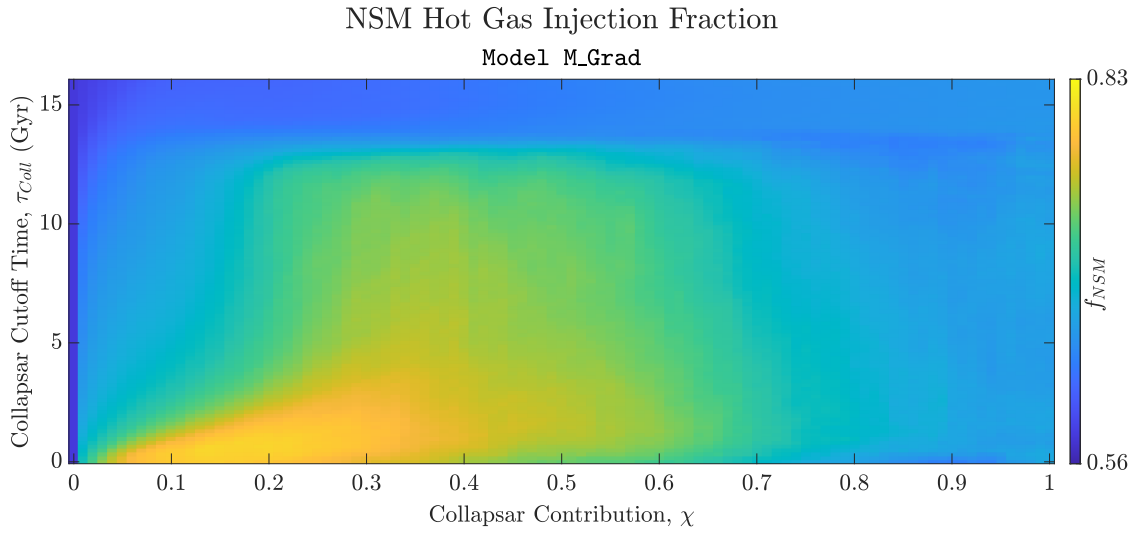

Figure 5.13: The mean value of $f_{\rm NSM}$, the fraction of NSM ejecta placed into the hot gas reservoir for successful M_GRAD models.

5.7.2 Hot Gas

The inclusion of a thermal gas phase is a relatively rare feature of GCE models, though recent work has highlighted its importance (Khoperskov et al., 2021; Schönrich & Weinberg, 2019). Within SACEM and RAMICES II we implement the thermal phase in a relatively crude fashion: a distinct hot and cold phase, with the hot phase being fed with a fraction f of the enriched gas, and then 'cooling' into the cold phase with a decay frequency λ such that $\dot{M}_{\rm cool} = \lambda M_{\rm hot}$.

Figure 5.13 shows the variation of the parameter $f_{\rm NSM}$ in the successful models, which displays an interesting pattern, which we argue both vindicates the inclusion of the hot phase, and is central to our heuristic understanding of the interplay between NSM events and collapsars.

At very low χ ($\ll 0.1$), this value is favoured to be low, around 0.6 or below (for comparison, the corresponding CCSN, Collapsar and SNIa values exceed 0.9 almost everywhere) – this was argued for in Schönrich & Weinberg (2019) as a way for the thermal properties to permit NSM to be viable contributors at early times, due to the correspondingly higher immediate contribution to the star forming reservoir. As collapsars become more prominent at early times (higher χ and lower $\tau_{\rm coll}$), it is no longer necessary to invoke the arguments of Schönrich & Weinberg (2019) to explain the early time [Eu/Fe] values, and hence the need for low values of $f_{\rm NSM}$ to meet the chemical criteria at [Fe/H] = -1.5 vanishes, leading to the observed increase in $f_{\rm NSM}$. At even higher values of χ, the value of $f_{\rm NSM}$ decreases again – notably to a value of 0.65, which is the exact midpoint of the permitted range: hence we are seeing ambivalence to the value of $f_{\rm NSM}$ when determining if a model is successful.

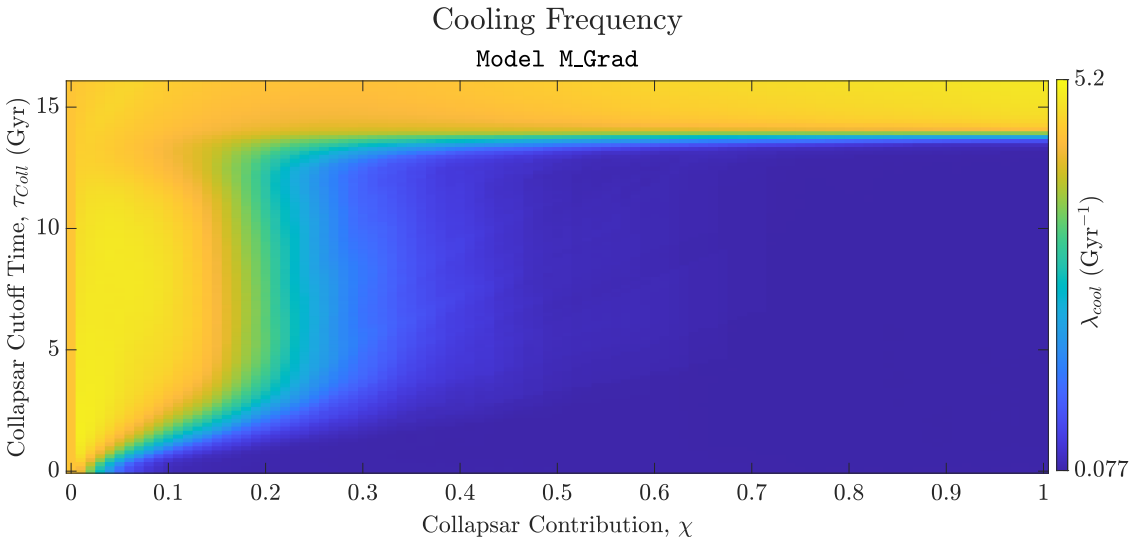

Figure 5.14: The mean ratio λ_{cool}, the hot gas cooling rate, for successful models constrained with M_Grad. Note that the midpoint of the permitted range for λ_{cool} is 5 – models close to this value are likely to be insensitive to the value of λ.

However, for $\chi < 0.7$ there are clear signs that the thermal properties of NSM are very important for determining the success or failure of a model.

Given that the thermal properties have been shown to be important, we now consider the case where the cooling rate is small. If the hot-gas injection fractions are non trivial, then a large portion of the enriched gas can be secreted away in the hot gas phase, preventing over-enrichment at early times, and providing a 'source' of enriched gas into the star-forming phase long after the gas itself was actually produced. Whilst this is desirable to an extent – it is reasonable to expect exactly this to happen – Figure 5.14 shows that the successful models with a large collapsar contribution were almost overwhelmingly those which abused this property to the extreme.

In Figure 5.14 we see that for $\chi > 0.3$ and $\tau_{\text{coll}} < 13$Gyr, the mean value of λ is ≈ 0.077 Gyr^{-1}, an order of magnitude below the commonly used value of ~ 1 Gyr^{-1}. These models cool almost no gas into the cold gas phase, allowing the hot gas phase to become hyper-enriched relative to the cold gas phase. This shows that these successful models favoured extreme thermal fractionation as an explanation for the chemical history of the galaxy – whilst we do not doubt that a hot gas population is important for understanding GCE, the cooling rates indicated here strongly suggest that these models are highly unphysical.

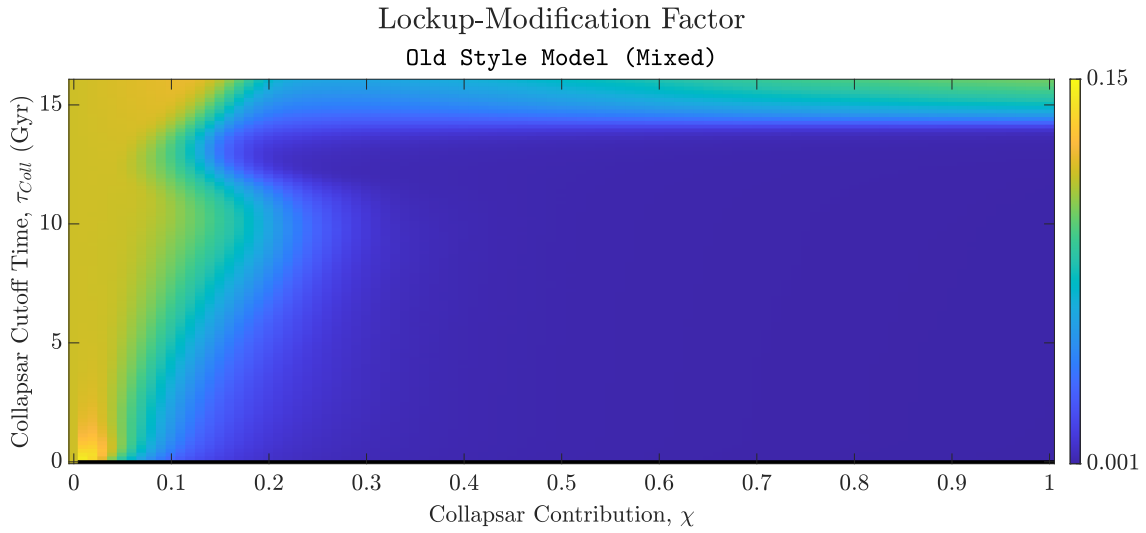

Figure 5.15: The mean value of $F_{\rm mod}$, the lockup-modification factor, for successful models at each point in $\tau_{\rm coll} - \chi$ space for the prototype (`Old-Style`) simulations in which $F_{\rm mod}$ is bounded by $0 < F_{\rm mod} \leq 1$. We note that the black strip at $\tau_{\rm coll} = 0$ is the nonsensical combination $\tau_{\rm coll} = 0, \chi > 0$. As per Figure 5.8 this was omitted it all other plots.

5.7.3 Lockup Modification

The lockup modification factor, $F_{\rm mod}$, encodes the rate at which synthesised material is removed from the cold gas reservoir by star formation, such that the lockup rate is $\propto F_{\rm mod}\rho_{\rm sfr}$ rather than the zeroth-order approximation $\propto \rho_{\rm sfr}$. $F_{\rm mod}$ differs from 1 since, though normal CCSN events may not synthesise new Eu material, they can recycle pollutant metals back to the ISM, thereby reducing the effective lockup rate of synthesised material.

We must be aware, however, of the extreme case $F_{\rm mod} \approx 0$ in which no synthesised material can be locked up. Star formation would continue to deplete the unenriched gas mass, such that the stars are preferentially forming from primordial material: the cold gas phase becomes chemically fractionated, and becomes ever more enriched without any additional synthesis events.

Collapsar models with $\tau_{\rm coll} < 4$Gyr could exploit this property to 'hide' a reservoir of Eu which persists until late times, when by all physical reasoning, it should have been depleted by the lockup of the continual star formation. This was made evident in a prototype set of simulations (termed the Old-Style simulations), in which the lower bound on $F_{\rm mod}$ was set to be 0 for all simulations. Figure 5.15 shows the behaviour of $F_{\rm mod}$ of successful models in the prototype simulation. We note that the prototype simulations differ from the final ones in several ways that make a direct comparison difficult, but we can see that the lower right section of the Figure 5.15 is overwhelmingly dominated by models for which $F_{\rm mod} \sim 0.001$, the case

where there is no chemical lockup. A similar pattern was observed for models U, W and V: an overwhelming bias towards extremely small values of $F_{\rm mod}$ in the region of high-χ and low-$\tau_{\rm coll}$.

Such models are evidently unphysical, and so in the final suite of simulations, we bounded $F_{\rm mod}$ from below by 0.3. This value was motivated by comparison with a reasonable IMF, noting that $\int_0^1 M\zeta(M)\mathrm{d}M > 0.35$ for all commonly used IMFs (Salpeter, 1955; Chabrier, 2003). Since stars with $M = M_\odot$ have lifetimes $\sim$10Gyrs, at least this fraction of the gas must be locked-up for long timescales.

When bounded in this fashion, the bias in the values of $F_{\rm mod}$ vanished almost entirely, and so we surmise that we have closed off this unphysical route to achieving 'success'. We note that some values near the boundary of the exclusion zone of both W_Grad and V_Grad did show a small bias towards smaller values of $F_{\rm mod}$, indicating that if we were to impose stricter and more physical constraints on the value of this parameter, these models would similarly be eliminated. However, in the spirit of our attempt to find the minimum possible set of constraints, we leave these potentially problematic models unchallenged, as the size of the exclusion zone is already sufficient to draw our conclusions.

5.7.4 Improperly Coupled Models

In the above discussion, we saw a number of traps which models can fall into: failing to properly lock up their materials, poor treatment of hot gas phases and associated cooling rates and torturing the SFR until it allows you to replicate your desired features. In utilising these 'traps', the models seemingly satisfy all chemical constraints which were placed upon them. However, on closer inspection these models were only able to reproduce the chemical properties due to an unphysical assumption elsewhere in the model.

The general theme of these assumptions was that they allowed the chemical reservoirs to become separated or fractionated in some way, such that the evolution of the SFR, cooling, chemical enrichment and subsequent lockup of the reservoirs was not functioning properly.

We note, for example, that Si19 use an SFR which is decoupled from the present gaseous or stellar mass in their model, which we suggest falls into a similar camp of improperly coupled models, and explains why their findings are in such strong tension to our own, despite seeming to meet all of the observable chemical criteria.

We also suggest that, due to the expansive sampling of parameter space, we would have been able to notice if some unusual combination of physical properties were able to replicate both the chemical data, and not fall into one of the four 'traps' outlined above – in fact, we observed no such signal. Given that this is the case, it must be true that any successful

GCE model which replicates a collapsar-dominated galaxy must strongly deviate from the physics encapsulated in the core physics of SACEM.

5.8 Conclusions

In this chapter, we have performed the first comprehensive chemical evolution study which examines the multi-dimensional parameter-space associated with the origin and evolution of galactic r-process material. In this extensive analysis, we could find no viable model with collapsars as the dominant source for today's r-process element budget.

In this work, we have introduced our newly developed Simple Analytical Chemical Evolution Model (SACEM). SACEM is an analytical framework, which incorporates the relevant physics (star formation histories, inflow, outflow of gas, yields to both a hot and cold gas phase, cooling of material from the hot phase, star-forming ISM, and different temporal and thermal properties for different sources of yields), but at the same time has run-times of fractions of a second, i.e. orders of magnitude faster than existing chemical evolution codes. Although SACEM relies heavily on some simplifying approximations (namely single-zone space with instantaneous mixing, empirical fixing of yields, and does not consider the lifetime of stellar populations), we found good agreement with the full RAMICES II model, which does not make such approximations.

Where there are divergences between the models, we found that SACEM was more generous to collapsar-dominated models than RAMICES II, though we note that both SACEM and RAMICES II do not directly consider tertiary sources of r-process material (such as the i-process or magneto-rotational supernovae), the impact of dwarf galaxy accretion, and assume a simple form of the collapsar yields. Although these omissions might limit the generality of our results, we have justified their long-term impacts on the abundance patterns as negligible, or already encapsulated in part by features of our models. Hence, our results are robust against these approximations.

A central problem holding chemical evolution studies back has been a reliance on costly models in a high-dimensional parameter space, which has forced prior studies to operate with exploratory modelling of a small number of models. SACEM's performance allowed us to run $> 10^{11}$ models, mapping out the full parameter space of r-process chemical evolution with both collapsars and NSM, and allowing us to pursue an entirely different strategy: instead of trying to find models that match some observational constraints, we drew up a full set of "minimal consensus" observational constraints which models must replicate, and look for those models which fail to reach even these lax conditions: rough 'bounding boxes' that

our chemical tracks in [X/Fe]-[Fe/H] must pass through (Figure 5.1), chemical equilibrium imposed across the final 2Gyr of evolution, and a sustained rate of star formation. Our search through parameter space was bounded by imposed conditions on the allowed parameter values – the one which best represented the Milky Way (whilst still allowing for large degrees of variation) we termed the VIABLE set of constraints. We also explored the parameter space with unphysically lax constraints on parameter values (the UNCONSTRAINED and MIXED sets) – comparing these with the VIABLE constraint set reveals and analysis how classes of models with dominant r-process contributions discussed in the literature appeared to satisfy observational constraints. However, with the VIABLE constraints, we have found that:

1. no SACEM model could be found where collapsars contribute more than 30 per cent of the modern r-process budget, as long as collapsars were suppressed as in MacFadyen & Woosley (1999). NSM were always required to be dominant (Figure 5.9).

2. A significant collapsar contribution at early times was not eliminated: Many models in which NSM are responsible for > 99 per cent the Eu abundance at late times had > 50 per cent collapsar contributions at $t < 1$Gyr (Figure 3.1).

3. The RAMICES II code shows that the remaining parameter space allowed by SACEM still contains models that are in stark contradiction with the data. In particular the metallicity dependent cut-off can introduce a radial [Eu/Fe] increase in the galactic disc (Figure 5.6), which starkly contrasts with observations: our limit of <30 per cent is likely still too high, and can be refined further.

We deliberately chose constraints which were overly generous, and our results should be seen as the maximum possible contribution of collapsars to the modern r-process budget. We leave further discussion regarding how far the constraints can be pushed and how far the allowed parameter space can be further shrunk to future studies, preferring to keep our argument simple: we have shown that even a minimal set of constraints permits no models with collapsars as dominant source of r-process elements, and thus leaving by exclusion (Sn08) only NSM as a dominant source.

In connection with the numerous celestial spheres, however, there are many Motions.

Nicolas Copernicus

6 Constraining Stellar Evolution Models

Contents

6.1 Introduction

The study of the evolution of the chemical content of our universe and that of the pre- and post- supernova nucelosynthesis of massive stars are inextricably linked. The vast majority of elements heavier than helium are expected to have been synthesised in the many nucleosynthetic pathways of a stellar life-cycle, so to understand the composition of a galaxy, it is key to understand the production and destruction channels in the hearts of its component stars.

The publication of new nucleosynthesis grids – such as those of Karakas & Lugaro (2016), or Limongi & Chieffi (2018) – are therefore often followed by series of studies integrating these grids into models of GCE (for example, Kobayashi et al. (2020) and Prantzos et al. (2018), respectively). Such studies of stellar and supernovae processes are the lifeblood of GCE as an academic field.

There have, however, been notable instances of chemical evolution models constraining the nucleosynthesis models: Hoyle (1954) famously predicted the presence of the 7.7MeV resonance line in ^{12}C, showing that its boost to helium fusion was necessary to match the carbon abundance in the solar neighbourhood.

To the present day, GCE has retained this diagnostic appeal as timescale and resolution problems, inaccessible measurements in both nuclear and plasma physics (Bailey et al., 2015), as well as vast and weakly constrained parameters for the internal structure of massive stars leave us with an uncertain and coarse understanding of the evolution of stars as a function of their initial properties. This results in unphysical assumptions and impositions, like the assumption of fixed ^{56}Ni yields in core collapse supernovae, setting a rather artificial 'mass cut' to 1D supernova models, i.e. the imposed boundary between the matter falling back onto the proto-neutron star and the ejected material (Limongi & Chieffi 2018, henceforth LC18). While such a limit made sense within the framework of 1D/1D-piston models, both 3D models (Vartanyan et al., 2018) and observations (Chevalier & Kirshner, 1979; Hwang & Laming, 2012) now show that convection translates into plumes from different levels escaping the collapse.

The aim of this chapter is to develop and formalise a set of tools that allow the backpropagation of GCE-derived constraints into the study of stellar properties, and to elevate the use of GCE models for constraining model assumptions from anecdotal use to numerically quantifiable constraints. In particular, we focus on the ability to rapidly and efficiently diagnose a proposed nucleosynthetic yield table as being incompatible with the observed properties of our galaxy.

6.2 Data

6.2.1 Chemical Abundances

To tether theoretical models of stellar evolution to the real world, we use the observational data from GALAH DR3 (Buder et al., 2021), depicted in Figure 2.2. This dataset includes abundances for 588,571 stars across 30 elemental species - 11 of which include corrections for NLTE effects. We choose to study the abundances of the alpha elements O, Mg and Si and Ca, and use Al to deselect halo stars. All used abundances are NLTE-corrected in GALAH DR3.

We limit the sample to stars with `flag_sp_fe_h` and `flag_sp_x_fe = 0` for each element `x`, removing stars which have only upper bounds on abundances. We also impose a quality cut on abundances: $\sigma_X < 0.1$ dex. Halo stars carry the signature of their accreted galaxy's evolution, rather than illuminating the in-situ Milky Way signal. To remove these extragalactic contaminants, we follow Hayes et al. (2018), excising all stars for which [Al/Fe] $<$ 0 dex at [Fe/H] $<$ -1, and [Al/Fe] $<$ -0.1 dex for [Fe/H] $<$ -0.5.

The chemical thick/thin disc dichotomy is most clearly visible in the [Mg/Fe] plane. We follow the usual nomenclature, calling high-α/low-metallicity stars the thick disc and

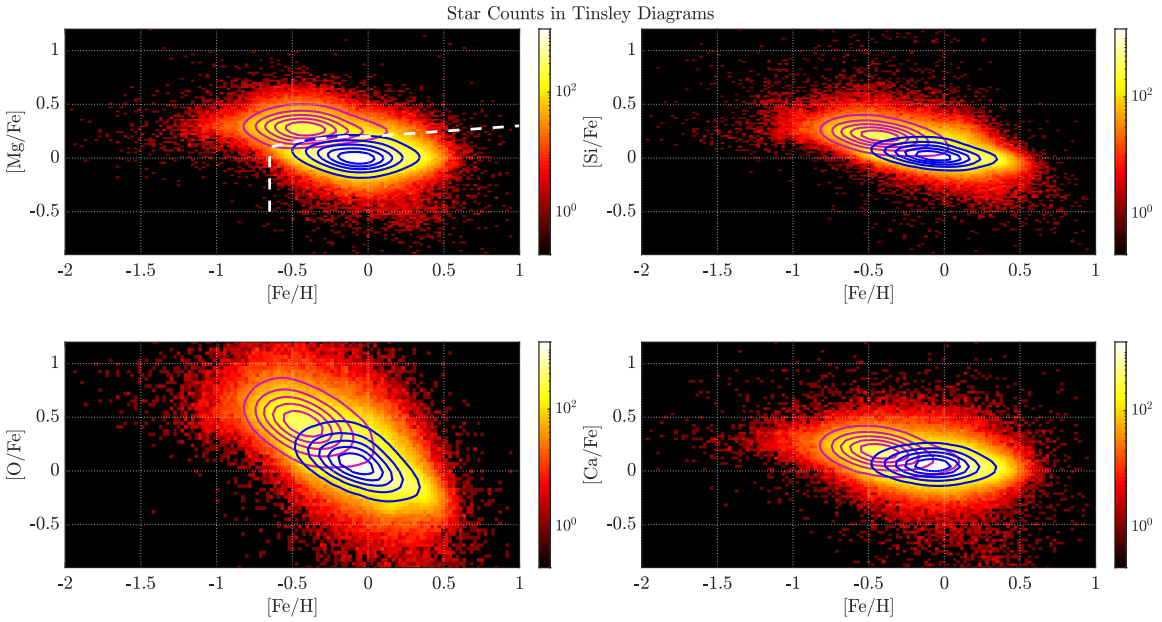

Figure 6.1: The data from Figure 2.2 overplotted with contours of Eq. (6.1) for the thick disc (*purple*) and thin disc (*blue*), as determined by the white dashed line in the [Mg/Fe] plot.

low-α/high-metallicity stars the thin disc, with the boundary between these regions denoted by the dashed line in the upper-left plot of Figure 6.1.

Following the discussion of §2.3.1, we are interested in the mean abundance ratios for the early disc populations. For the sake of simplicity, we recover these values by assuming a Gaussian errors to produce estimated population density contours in Figure 6.1 using the following predictor:

$$\mathcal{S}(f,x) = \frac{1}{N}\sum_{\mathrm{i}=1}^{N} \frac{1}{\sqrt{2\pi\sigma_{F,i}\sigma_{X,i}}} \exp\left(\frac{(f-F_i)^2}{2\sigma_{F,i}^2} + \frac{(x-X_i)^2}{2\sigma_{X,i}^2}\right). \tag{6.1}$$

Where (F_i, X_i) are the coordinates in the [Fe/H]-[X/Fe] abundance plane for the i^{th} star in the sample, and σ_i are the associated errors. We note that this is not necesarily indicative of the true underlying populations (which would require an involved deconvolution), but it is sufficient for us to recover the mean abundance values.

We follow the traditional nomenclature of thin/thick disc (§2.2), though we note that for the present work we have no stake in the various interpretations of dichotomies or lack thereof and different formation scenarios (Gilmore & Reid, 1983; Chiappini et al., 1997; Schönrich & Binney, 2009a; Loebman et al., 2011; Bovy et al., 2012), instead it is merely important for us identify the early populations in a relatively homogeneous group of stars, which the thick disc constitutes. The overdensities separate best in [Mg/Fe], so we use the top left figure to separate stars along the white line, drawing then separate contours in

Table 6.1: The extracted 'thick disc ridge' values [X/Fe] obtained from Figure 6.1, corresponding to the low metallicity, early-time abundances of the cold star-forming gas in the Milky Way.

Species	Lower Bound	Mean	Upper Bound
[Mg/Fe]	0.25	0.29	0.4
[Si/Fe]	0.15	0.26	0.35
[O/Fe]	0.45	0.56	0.65
[Ca/Fe]	0.1	0.20	0.3

all plots. To reduce contamination from the thin disc and the abundance knee, we limit conclusions about the thick disc to [Fe/H] $<$ -0.6.

These mean [X/Fe] for the "thick" disc with generous error bars are given in Table 6.1, and will be used as model constraints throughout this work.

6.2.2 Yield Grids

We will demonstrate our techniques on the yield grids of LC18. Their yield grids feature pre- and post-supernova chemical yields for 54 different elements up to Pb (including all elements below $Z = 42$) on a grid of 9 masses between 13 $M_\odot$ and 120$M_\odot$, four metallicities of [Fe/H] = -3, -2 , -1 and 0 dex and three ZAMS equatorial rotation velocities $v_{\rm zams}$ = 0, 150 and 300km/s.

LC18 deserve special status, as they have the only fully published and comprehensive yield grid in the literature after Woosley & Weaver (1995) with a clear policy of transparency and full publication of all necessary data. Their model results are also of particular interest, since no ad-hoc calibrations of nuclear yield rates have been done to meet chemical evolution models. This follows their tradition since Chieffi & Limongi (2004), which published the full range of mass cuts. The avoidance of human intervention is particularly useful, since it allows one to challenge the nuclear reaction rates and stellar physics directly against available data via chemical evolution models. In this sense, our study presented here should by no means be understood as a criticism of their work. By contrast, their principles allow us to learn more about physics when their models fail to reproduce data.

LC18 published four different yield tables, corresponding to different classes of models denoted `F, I, M` and `R`, distinguished by their use (or not) of the the 'Mixing Fallback Scheme' pioneered by Umeda & Nomoto (2002). In addition, some models undergo 'total collapse', absorbing the entire stellar mass into the compact remnant: only the stellar wind is ejected.

- **Set F:** No 'mixing fallback', no total collapse
- **Set M:** Includes 'mixing fallback', no total collapse

- **Set I:** No 'mixing fallback', includes total collapse
- **Set R:** Includes 'mixing fallback', includes total collapse

LC18 note their preference for Set `R`; however, we will test all four sets equally against the constraints of galactic chemistry.

6.2.3 Stellar Rotational Distributions

Stellar rotation impacts stellar evolution (and hence the subsequent yields) through changes to the internal structure and rotational mixing (Sweet, 1950; Meynet & Maeder, 2000; Leitherer et al., 2014).

Large-scale spectroscopic surveys such as VLT-FLAMES (Pasquini et al., 2002) have attained measurement accuracies and number statistics sufficient to allow deconvolution the effects of inclination and so infer the the present distribution of stellar equatorial rotation velocities.

We note, however, that the current rotation speed of evolved stars is different from the rotation speed when these massive stars have entered the main sequence (the parameter required for the LC18 evolution models). In most formation models the disc accretion spins young stars to an extreme $\sim 50\%$ of break-up velocity at birth. This fast rotation is later slowed by stellar expansion and/or angular momentum loss through (magnetic) stellar winds (Parker, 1958; Ekström et al., 2012). This process can be seen via the well-known dependence of stellar rotation on age and mass/effective temperature (Kraft, 1967; Matt et al., 2015; Healy et al., 2021). Therefore, whilst the ZAMS rotation distribution function is not directly observable, the current measured rotation distribution places a lower limit on what this distribution might be.

Zorec & Royer (2012) studied low mass stars ($M \leq 3M_{\odot}$) and found at most 20% of stars are classified as slow rotators (decreased to $< 1\%$ for the youngest populations) with $v_{\rm eq} < 50$km/s, whilst the remainder populations were fast rotators with mean rotation speeds in excess of 150km/s. Ramírez-Agudelo et al. (2013) (RA13) studied O-type stars (tightly peaked around $M/M_{\odot} = 20 \pm 5$, per Schneider et al. 2018) and found that the distribution was best fit by the sum of a Gamma distribution component and a Gaussian component, giving a long tail to extremely fast rotators (> 500km/s). We note that their predicted distribution has $p(v_{\rm eq} > 300{\rm km/s}) = 0.2$, and thus a significant subset of the data lies outside the regime covered by LC18.

In §6.3.3 we use these distributions to place numerical expectations on our integral weighting parameters.

6.3 Model

In Chapter 3 we developed a powerful, general, and intentionally weakly constrained tool in the form of SACEM. Eq. (3.12) provided us with an analytic expression for the rate of production of an element through arbitrary sources j,

However, in this chapter we focus on early times, elements dominated by core collapse supernovae yields and on the thick-disc abundances. This permits the following further simplifications:

- All synthesis is dominated by CCSN events (especially, SNIa channel has not yet opened)
- We may ignore thermal phasing (assuming no element-dependence on the contribution to different gas phases)
- We can make the instantaneous recycling approximation
- Accreted gas is almost entirely pristine

Under these approximations, we introduce a metallicity-weighting distribution function $\omega(Z)$, and hence derive the following expression for the thick disc abundance at early times:

$$\left[\frac{A}{B}\right] = \log_{10}\left(\frac{\int_{M_{\rm min}}^{\infty}\zeta(M)\mathrm{d}M\int_0^Z\omega(Z)\mathrm{d}ZY_A(M,Z)}{\int_{M_{\rm min}}^{\infty}\zeta(M)\mathrm{d}M\int_0^Z\omega(Z)\mathrm{d}ZY_B(M,Z)}\right) - \mathcal{R}_{\odot}(A,B). \tag{6.2}$$

Here $\mathcal{R}_{\odot}(A,B)$ is the solar normalisation for these elements.

6.3.1 Introducing Rotation

Whilst Langer et al. (1999) showed it was possible to account for the effects of stellar rotation in the evolution of massive stars, most nucleosynthetic yield grids are restricted to only M and Z dependence, for simplicity. LC18, however, published comprehensive yield grids for a large number of elements, covering a large part of (M,Z) space for three choices of equatorial rotation speed, and thus we can, for the first time, fully include the effects of stellar rotation in GCE models.

Assuming that rotation does not significantly change stellar lifetimes in chemical evolution, the instantaneous recycling approximation holds. We further neglect rotational stochasticity, averaging over the rotation distribution, and extending Eq. (6.2) to:

$$\left[\frac{A}{B}\right] = \log_{10}\left(\frac{\iiint\mathrm{d}M\mathrm{d}Z\mathrm{d}v\zeta(M)\omega(Z)f_{M,Z}(v)Y_A(M,Z,v)}{\iiint\mathrm{d}M\mathrm{d}Z\mathrm{d}v\zeta(M)\omega(Z)f_{M,Z}(v)Y_B(M,Z,v)}\right) - \mathcal{R}_{\odot}(A,B) \tag{6.3}$$

Here $f_{M,Z}(v)$ is the ZAMS equatorial velocity distribution function.

The question now arises as to which rotational distribution to use. We first note that, since we are using finite sampled yield grids (§6.3.2), any distribution function is merely an analytical way to assign weights to the finite sum:

$$\int f(v)Y_X(M_k, Z, v)\mathrm{d}v := \sum_i \mu_{ik} Y(M_k, Z, v_i) \tag{6.4}$$

The values of μ_i are constrained only by the requirement that $\sum_i \mu_i = 1$ and $\mu_i > 0$. Should we so choose, we could manually assign their values, as was done in Prantzos et al. (2018). For simplicity, We have assumed here that although the distribution function can change with mass, it is not a function of metallicity.

In the case of an observationally motivated model, we note from §6.2.3 that the young, fast rotating tail is well fit by a Maxwellian in the dimensionless parameter $\beta = v_{\rm rot}/v_{\rm BU}(M)$, where $v_{\rm BU}$ is the ZAMS break-up velocity:

$$f(\beta)\mathrm{d}\beta = \mathcal{N}\beta^2 \exp\left(-\frac{\beta^2}{2\sigma^2}\right)\mathrm{d}\beta. \tag{6.5}$$

Here, σ is the rotational 'temperature' of the distribution, approximately equal to the mean rotational speed of stars (in units of the breakup velocity), and $\mathcal{N}$ ensures that the function is normalised over the domain $\beta \in [0, 1]$, since stars cannot rotate faster than their breakup velocity. Assuming a simple power-law for the ZAMS radius (Demircan & Kahraman, 1991), $R = R_\odot(M/M_\odot)^\kappa$, with $\kappa = 0.5 - 0.7$, then we obtain the distribution function in $v_{\rm rot}$: a Maxwellian distribution with temperature

$$\tilde{\sigma}(M) = \sigma \times v_{\rm BU} = \sigma\sqrt{\frac{GM_\odot}{R_\odot}}\left(\frac{M}{M_\odot}\right)^{\frac{1-\kappa}{2}} \tag{6.6}$$

6.3.2 Yield Weightings

To evaluate [A/B] we need to integrate the provided yield grids over M, Z and v. Whilst lacking a full functional form, LC18 provides a sample of 9 masses ($13M_\odot \leq M \leq 120M_\odot$), three ZAMS rotational velocities (0km/s, 150km/s and 300km/s), and 4 metallicities ([Fe/H] = -3,-2,-1 and 0).

We employ a linear interpolation between the points to compute the integrals, formally defined in Appendix A. We note again that this equivalent to assigning weights to the finite number of points:

$$\iiint \mathrm{d}M\mathrm{d}Z\mathrm{d}v\zeta(M)\omega(Z)f_{M,Z}(v)Y_A(M, Z, v) \approx \sum_{i,j,k} W_{ijk}Y(M_j, Z_k, v_i) \tag{6.7}$$

The approximate nature of this relationships arises from the finite number of sampled points. Equivalently, one could manually assign the weightings W_{ijk} as was done, for example, for the rotational weights in Prantzos et al. (2018). In such a 'manual' case, we assume that each of the distribution functions is independent[1] of the other parameters, such that we may associate a set of weights with each integral:

$$\int \mathrm{d}M\zeta(M)q(M) \approx \sum_k \nu_k q(M_k) \tag{6.8}$$

$$\int \mathrm{d}Z\omega(Z)p(Z) \approx \sum_j \epsilon_k p(Z_k) \tag{6.9}$$

$$\int \mathrm{d}v f(v)r(v) \approx \sum_i \mu_i r(v_i) \tag{6.10}$$

Therefore Eq. (6.7) becomes:

$$\iiint \mathrm{d}M\mathrm{d}Z\mathrm{d}v\zeta(M)\omega(Z)f(v)Y(M,Z,v) := \sum_{ijk} \mu_i\epsilon_k\nu_k Y(M_k,Z_j,v_i) \tag{6.11}$$

The assumption that the weights are derived from good DFs provides the constraint $0 \leq \mu_i \leq 1$ and $\sum \mu_i = 1$. We will use this strategy in §6.4 to test the results independent of our chosen rotation and metallicity DFs.

Extrapolation

In the non-manual assignment case, we have a fundamental problem that the allowed parameter space over which the integrals Eq. (6.3) are perfomed extends beyond the grid of sample points, necessitating either a truncation of the integral, or extrapolation of the data. We expect this to only be meaningful at the lower end (both in mass and rotation), since the declining PDFs imply low weights at the top end. We will investigate three extrapolation schemes, the `Truncated`, `Flat` and `Limited Linear` cases, demonstrated in Figure 6.2.

Within this scheme, the mass integral has 4 hyperparameters: $M_\pm$, the upper and lower bounds of the mass integral; $\delta_\pm^M$ the amount of extrapolation allowed above and below the mass data points; whilst the rotation integral introduces 2 (as the $v = 0$km/s datapoint is a true physical boundary): v_+, the upper limit on the rotation integral; and δ_+^v, the maximum amount of extrapolation allowed above the rotational data points, leaving us with a total of 6 extrapolation hyperparameters of a given model.

We note that each scheme in Figure 6.2 may be successively mapped onto each other - scheme 3 with $\delta_\pm = 0$ is scheme 2, and scheme 2 may be mapped to scheme 1 by sending $M_\pm$ to the first and last datapoints.

[1]Note that this assumption is violated when we use the 'thermal' assignment, as we assume $f(v) = f(M, v)$. We include both 'manual' and 'thermal' models to ensure this assumption does not lead us astray.

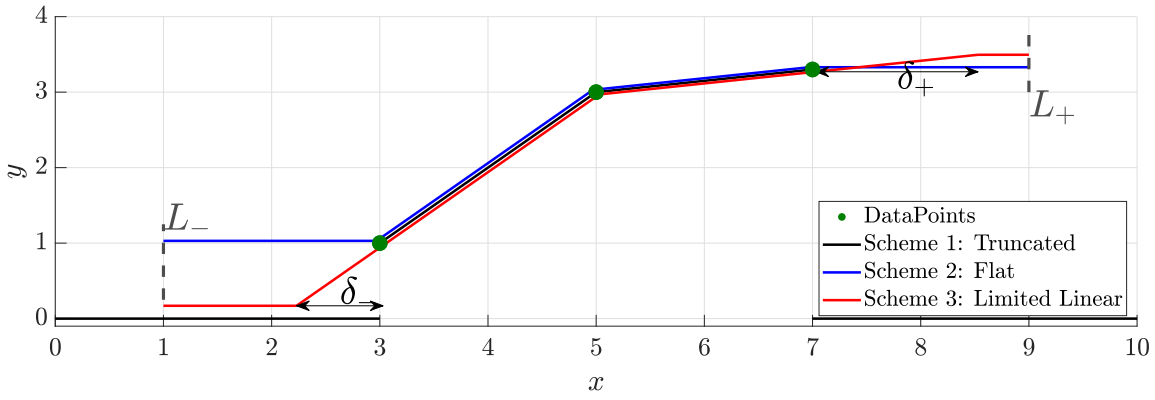

Figure 6.2: A mockup of the three extrapolation techniques used to explore the datasets. The `Truncated` scheme is zero outside of the provided data, whilst `Flat` remains stationary outside of the data range, befre being set to zero at $L_\pm$. The `Limited Linear` scheme allows a small amount of linear extrapolation, for a distance $\delta_\pm$, before following the `Flat` scheme.

6.3.3 Observational Constraints on Rotation

Under a manual assignment of the rotational weightings, the RA13 distribution would imply that the $(0, 150, 300)$ km/s grids would be weighted in the ratio $(0.192, 0.474, 0.34)$ respectively.

We also note that the fast-rotating tail of the RA13 distribution is equally well fit by a Maxwellian distribution (Eq. 6.5) with rotational temperature $\sigma = 0.25 v_{\rm breakup}$, corresponding to a mean rotational speed of 40% of breakup[2]. We therefore adopt a Maxwellian ZAMS distribution, expecting the spin-temperature to be on the order of $0.25 v_{\rm breakup}$.

6.4 Distribution-Independent Viability

The previous assumptions imply a large allowed parameter space is required to explore the permitted values of $\left[\frac{\rm X}{\rm Fe}\right]$. Utilising negative inference, we claim that a yield grid can be shown to be incompatible with observational evidence if all possible parameterisations fail to reproduce the observed data. In particular, we will show that we can provide such a demonstration independent of any assumed stellar or metallicity distribution function.

6.4.1 Method

Using Eq. (6.11), we rewrite Eq. (6.3) as

$$\left[\frac{A}{B}\right] = \log_{10}\left(\frac{\sum_{ij}\mu_i\epsilon_j f_{A,ij}}{\sum_{kl}\mu_k\epsilon_l f_{B,kl}}\right) - \mathcal{R}_\odot(A,B) \tag{6.12}$$

[2]When calculating breakup velocities, we assume the stars are rigid, spherical rotators and have a mass-radius relation given by a simple power law, following i.e. Demircan & Kahraman (1991)

Where

$$\begin{aligned} f_{X,ij} &= \int_{M_{\min}}^{\infty} \mathrm{d}M \zeta(M) Y_X(M, Z_j, v_i) \\ &= \sum_k \nu_k Y(M_k, Z_j, v_i). \end{aligned} \tag{6.13}$$

Ideally, μ_i would be calculated from the rotational DF (and appropriate interpolation/extrapolation), ϵ_j from the time-dependent metallicity distribution, and ν_k from the IMF. The latter is well-determined by the relatively well-populated mass grid and well-understood IMF (the Salpeter 1955 function with $\alpha \sim 2.3$), however the grids in v and Z are much sparser, and the distribution functions more poorly constrained.

To curb these uncertainties, we show in Appendix A.3 that the properties of μ_i, ϵ_k allow us to bound the ratio by:

$$\log_{10}\left(\min\left[\frac{f_{A,ij}}{f_{B,ij}}\right]\right) \leq \left[\frac{A}{B}\right]_{\text{early}} + \mathcal{R}_\odot(A,B) \leq \log_{10}\left(\max\left[\frac{f_{A,ij}}{f_{B,ij}}\right]\right) \tag{6.14}$$

It follows that no possible combination of μ_i, ϵ_k (and hence no meaningful distribution function) can cause the ratio to exceed the range of yield values at each single grid point.

This allows a simple diagnostic from the single grid points of what chemical abundances can, under any distribution, be attained by a galaxy with these stellar yields. Hence, if the grid-ratios prohibit values of $\left[\frac{\mathrm{A}}{\mathrm{B}}\right]$ consistent with the data, we can exclude these models as non-physical.

Alternatively, we might concede that the underlying model is sound, but that the sampled datapoints are too sparse, or missed strongly varying regions of parameter space. The difference between a badly sampled (but physical) model, and an unphysical model, whilst conceptually important, is irrelevant for us: the signal is clear that more stellar modelling is required.

6.4.2 Model Scan

We applied the technique outlined above to our chosen yield tables, targeting Oxygen, Magnesium, Silicon and Calcium, and across each of the four modelling techniques used in LC18 (the 'sets' F, I, M and R). We allowed all three of the rotation grids to be included in the analysis, but excluded the [Fe/H] = 0 yield grids, as the chemical thick disc occurred for [Fe/H] $\lesssim$ -1.

There are a number of hyperparameters which may be invoked to alter the value of our bounds, be they associated with physical parameters, such as the Salpeter parameter α, or associated with the hyperparameters used to invoke extrapolation to properly cover the integration bounds. We must therefore scan this multidimensional space in order to ensure that any viable combinations can be discovered.

For each element and set of hyperparameters, we are able to generate an estimate of the thick disc ridge using Eq. (6.14). A given estimate is counted as 'viable' if the estimating bounds overlaps with the observational error bars of the observed thick disc ridge found in Table 6.1. If a single viable estimate is discovered, we report the element as being viable within the provided yield grid.

We note that even if two or more elements from the same dataset are considered successful, the bounding nature of the test means that they could have used entirely different rotational distribution functions to each other, and hence we have no guarantees that a true, *simultaneous* chemical evolution model arises from this data. This test is only able to tell us what is evidently unphysical, not what *is* physical.

Scanning Bounds

The chosen bounds for the scans are:

- Salpeter parameter, $2 \leq \alpha \leq 2.6$ (all Schemes)
- Extended upper and lower bounds of the mass integral (Schemes 2, 3a and 3b)
  - $10 \leq M_{-} \leq 13$
  - $120 \leq M_{+} \leq 500$
- Upper and lower extrapolation distances for the mass integral (Scheme 3a and 3b only)
  - $0 \leq \delta_M^{+} \leq \min(100, M_c^{+})$
  - $M_{ab} \leq \delta_M^{-} \leq \min(3, M_c^{-})$
  - Where $M_{ab} = \begin{cases} 0 & \text{Scheme 3a} \\ 1.5 & \text{Scheme 3b} \end{cases}$

$M_c^{\pm}$ is a critical limit on the extrapolation parameter: it is the point at which the extrapolation would cause $Y(M, Z)$ to change sign (i.e. going from synthesis of material to destruction or vice versa) - we strongly prohibit such an eventuality unless the provided grid already oscillates between destruction and production. For the purposes of the following analysis, all of the used yield grids are synthesis-only, and hence we can always assume our yield functions are positive, even when extrapolated.

The integral parameters associated with the rotation are not included, as we are bounding the integral rather than trying to calculate it.

Table 6.2: Distribution-Independent Viability Analysis of 4 elements across the four datasets provided in LC18. A ✓ means that the model was deemed viable, whilst × means the model failed. This analysis shows that for some elements in each dataset, there exists no rotational distribution functions which would allow Eq.(6.3) to replictae the data observed in Figure 6.1. Notably, no single set is able to simultaneously satisfy the conditions on these four elements.

		Si	Mg	O	Ca
Scheme 1	**Set F**	×	×	×	×
	Set I	×	✓	✓	×
	Set M	×	×	×	✓
	Set R	×	✓	✓	✓
Scheme 2	**Set F**	×	✓	×	×
	Set I	×	✓	✓	×
	Set M	×	✓	×	✓
	Set R	✓	✓	✓	✓
Scheme 3a	**Set F**	×	✓	×	×
	Set I	×	✓	✓	×
	Set M	×	✓	×	✓
	Set R	✓	✓	✓	✓
Scheme 3b	**Set F**	×	×	×	×
	Set I	×	✓	✓	×
	Set M	×	×	×	✓
	Set R	×	✓	✓	✓

6.4.3 Results

The result of our Distribution-Independent Viability Analysis is shown in Table 6.2, listed for the each of the elements, LC18 data sets and extrapolation schema used.

Sets F, I and M always failed to match observations in the full set of elements, even given our extremely generous viability criteria. These three models are unable to generate acceptable GCE histories, and we must consider them unphysical.

Set R (which includes total stellar collapse above $25M_{\odot}$, and the improved mixing/fallback scheme of Umeda & Nomoto 2002) was considered viable in Schemes 2 and 3a, meaning that we cannot discard it as unphysical as we did the other sets. Although this validates LC18's statement that Set R was their preferred model, we issue some notes of caution.

Firstly, Set R was only able to replicate the observed data in [Si/Fe] if extrapolation beyond the provided data set was allowed: we see that Scheme 1 models failed to predict the correct thick disc ridge. It is only by explicitly extrapolating beyond the data's regime that it was considered 'viable'.

Secondly, we see that Set R *did* fail in Scheme 3b - the difference between 3a and 3b is that 3b forces at least $1.5M_{\odot}$ of extrapolation, whilst 3a allows the models to choose not to extrapolate (i.e. become Scheme 2 models). That 3b differs from Scheme 2 and 3a implies

that the model is relying very heavily on increasing the weighting of the $13M_{\odot}$ data point, whilst ignoring any surrounding trends (which would be magnified by higher extrapolation values). This indicates that the extrapolation required is potentially unphysical.

Finally, we remember that although the complete set of Si, Mg, O and Ca were able to successfully generate appropriate thick disc ridges, they were explictly allowed to generate their own metallicity and rotation distribution functions, independent of those used by the other elements, in order to do so. In reality, we know that all elements should have (to first-order) the same ϵ_i and μ_i values.

Therefore, whilst we have not yet completely ruled out Set R as being able to replicate features of the galactic chemical abundance, we have reason to suspect it, and therefore subject it to further testing.

We note that the analysis presented here was very quick to run: Table 6.2 was generated from scratch in under 10 seconds on a standard desktop PC, with the majority of the time being used to parse the LC18 datafiles.

6.5 Parameter Estimation

Although the above analysis indicated that the Set R from LC18 was 'viable', and hence the best candidate for reproducing the thick disc abundance ratios for all 4 of our target elements, we have no indication whether the required distribution functions (or equally, the set of rotational and metallicity weightings $\{\mu_i, \epsilon_j\}$) are physically reasonable, or to what extent we are able to *simulataneously* match the abundance ratios of the elements. In this section, we therefore attempt to isolate the rotational distribution functions which the provided yield grids would require, in order to reproduce the thick disc ridge under the current understanding of GCE.

6.5.1 Likelihood Model

The likelihood that a given model represented the data for some combination of elements $\{[\mathrm{X/Y}]_i\}$ given model parameters $\boldsymbol{\theta}$ would take the simple form:

$$\mathcal{L} \propto \left[\prod_i \exp\left(\frac{(\mathcal{D}_i - \mathcal{P}_i(\boldsymbol{\theta}))^2}{2\Delta_i^2}\right)\right] \times \mathrm{Prior}(\boldsymbol{\theta}) \tag{6.15}$$

Where $\mathcal{D}_i$ is the observational value of the thick disc abundance for the element combination $[\mathrm{X/Y}]_i$ = {[O/Fe], [Mg/Fe], [Si/Fe], [Ca/Fe]} , as given in Table 6.1, Δ_i is the associated error and $\mathcal{P}_i(\boldsymbol{\theta})$ is the predicted value of the thick disc abundance for a given model.

Parameters

A model is defined by the set of parameters $\boldsymbol{\theta}$ used to generate a theoretical estimate the value of the thick disc abundance, $\mathcal{P}(\boldsymbol{\theta})$. These parameters control the computation of the integral weightings of Eq. (6.11).

As before, we calculate the mass weightings $\{\nu_k\}$ analytically, allowing for mass-extrapolation, and we allow the metallicity weightings $\{\epsilon_i\}$ to be arbitrary, with the restriction that $\sum_i \epsilon_i = 1$ and $\epsilon_i > 0$.

For the assignment of the rotation weightings, $\{\mu_i\}$, we use two different paradigms: the 'manual' assignment follows the same prescription as ϵ_j, allowing the functional form to be arbitrary, whilst we also use a 'thermal' assignment, which weights the distributions according to Eq. (6.5), and hence introduces additional extrapolation parameters:

- Rotation parameters
  - For the 'thermal' models:
    - Population temperature, σ
    - Mass/Radius exponent, κ
    - Rotational extrapolation parameters v_+ and δ_V^+
  - For the 'manual' models:
    - Manual weightings: $\{\mu_i\}$
- Salpeter Alpha ($\zeta(M) \propto M^{-\alpha}$)
- Mass Extrapolation parameters: $M_\pm$, $\delta_M^\pm$
- Metallicity weightings, ϵ_i

Priors

The majority of our parameters are associated with the extrapolation and/or weighting of the various yield grids, and hence we have no *a priori* information on their values.

We therefore use uniform priors on the most of our parameters, with hard cuts at the imposed boundaries, which are physically motivated (i.e. $0 < \sigma < 1$ as the mean rotation speed cannot be less than 0, or greater than breakup, $M_- > 10 M_\odot$ due to the CCSN limit), or imposed to prevent runaway extrapolation.

For our two physically motivated parameters: the Salpeter (1955) α parameter and the Mass-Radius scaling exponent κ, we use simple Gaussian priors, $\mathcal{N}(\mu, \sigma)$ with mean μ and

Table 6.3: The priors used on the model parameters for Eq. (6.15). If a parameter exceeds the stated bounds, $p(\boldsymbol{\theta}) \to -\infty$, imposing a hard cut. $\mathcal{N}(\mu, \sigma)$ is the normal distribution with mean μ and standard deviation σ. All symbols are as defined in §6.5.1.

Parameter	Prior	Bounds
σ	Uniform	$0 \leq \sigma \leq 1$
κ	$\mathcal{N}(0.6, 0.3)$	-
v_+	Uniform	$300 \leq v_+ 600$
δ_v^+	Uniform	$0 \leq \delta_v^+ \leq 50$
μ_i	Uniform	$\mu_i > 0$, $\mu_n = 1 - \sum_i^{n-1} \mu_i$
ϵ_i	Uniform	$\epsilon_i > 0$, $\epsilon_n = 1 - \sum_i^{n-1} \epsilon_i$
α	$\mathcal{N}(2.3, 0.3)$	-
M_-	Uniform	$10 \leq M_- \leq 13$
M_+	Uniform	$120 \leq M_+ 200$
δ_M^-	Uniform	$0 \leq \delta_M^- \leq 3$
δ_M^+	Uniform	$0 \leq \delta_M^+ \leq 50$

standard error σ. Following i.e. Chabrier (2003) we take $p(\alpha) = \mathcal{N}(2.3, 0.3)$, and following Demircan & Kahraman (1991) we take $p(\kappa) = \mathcal{N}(0.6, 0.1)$.

We list the form of the priors imposed and the parameter bounds in Table 6.3. The prior on σ and $\{\mu_i\}$ is discussed in Appendix B

6.5.2 Model Evaluation

We utilised an MCMC search using the Python `emcee` package of Foreman-Mackey et al. (2013) and the `corner` package of Foreman-Mackey (2016) to produce the posterior PDFs for our models. Figure 6.4 shows the full corner plot from the thermally-weighted models, whilst the red curve in Figure 6.3 shows a subselection of the posterior PDFs for both the manually weighted and thermally-weighted models.

We first note that our assertion that the upwards extrapolation was negligible due to the decaying nature of the PDFs has been validated: Figure 6.4 shows no overall trend or correlation in M_+, δ_M^+, v_+ or δ_v^+ : this is consistent with these parameters being irrelevant for a sufficiently decaying PDF.

It is perhaps surprising, however, that the same is true for δ_M^-. Although there is a clear signal indicating a preference for no extrapolation (i.e. $M_- \to 13 M_\odot$), there appears to be no noticeable correlation with the extrapolation distance allowed. We note, however, that when we repeated this exercise with a different set of elements (i.e. just Si), a strong signal appeared in δ_M^-, whilst the posteriors for the upwards extrapolation parameters remained uniform: the

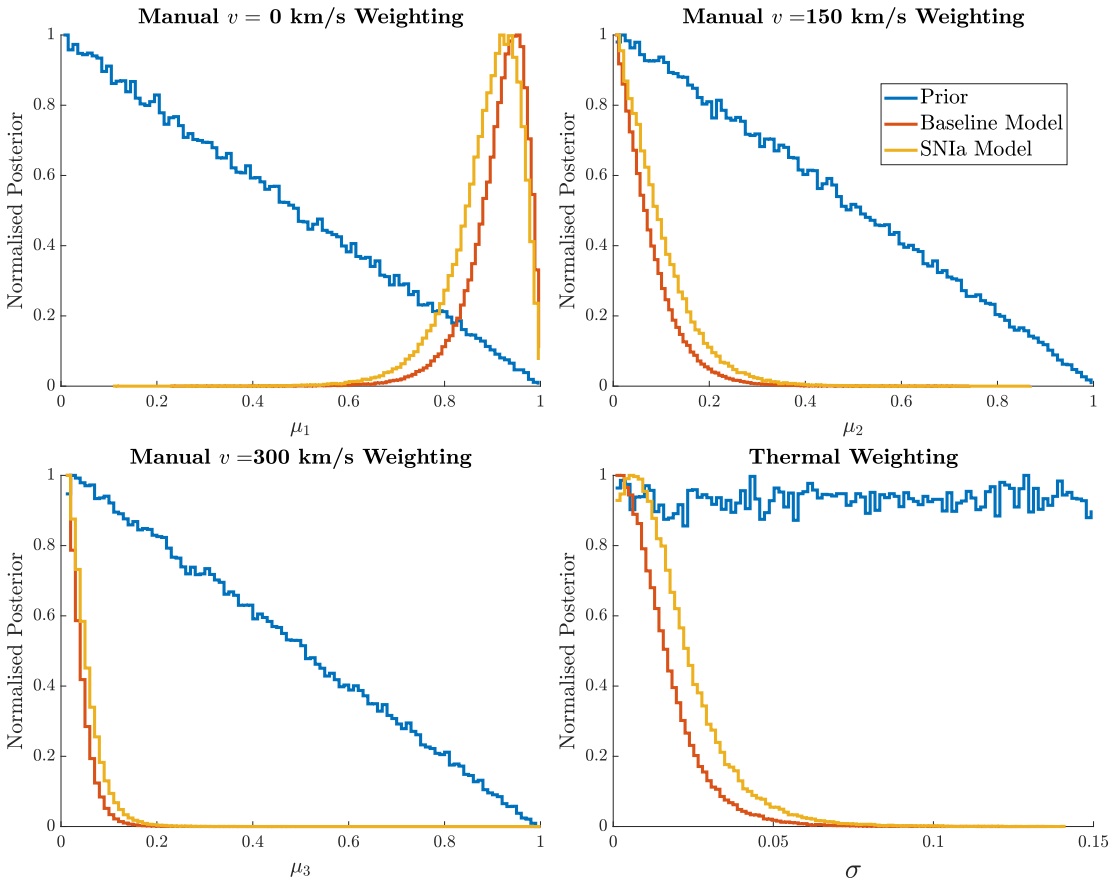

Figure 6.3: The posterior distributions of the rotation-weighting parameters for various models. The Baseline model for the σ-weighted models corresponds to Figure 6.4. The posterior for μ_3 is generated implicitly from the cross correlations between μ_1 and μ_2 and the condition that $\sum_i \mu_i = 1$.

lack of signal seen for δ_M^- appears to be a function of the simultaneous fitting of our four elements, whilst the lack of signal in the upwards extrapolation is a more general feature.

The key result from both plots is that the models can only be made to fit the data if they are *overwhelmingly* dominated by the zero-rotation grid: the "Baseline Model" row in Table 6.4 shows that the thermal model requires $\sigma = 0.02^{+0.01}_{-0.01}$, and the manual model requires $(\mu_1, \mu_2, \mu_3) \approx (0.91, 0.07, 0.04)$ i.e. these results indicate that the majority of stars must have no significant rotation for these yield grids to give us a thick disc ridge closest to the observed data.

This lies in strong tension with the data discussed in §6.2.3, for which the fast-rotating stars fit our thermal model with $\sigma \approx 0.25$, and the whole distribution required a weighting $(\mu_1, \mu_2, \mu_3) = (0.1923, 0.4740, 0.3400)$ - far above the yield-fitted values of $\sigma \ll 0.1$ and $\mu_1 \gg 0.5, \mu_2 < 0.1$ of our results here, and we recall again that these values were derived from the lower bound imposed by our observations being of current rotation state, instead of the ZAMS rotation state.

Table 6.4: A summary of the inferred values of rotational distribution function parameters for all of our models. Quoted values are the mean of the pdf, and the 16 and 84% confidence intervals. 'Baseline' models are those discussed in §6.5.2, 'SNIa' models are those introduced in §6.5.3 and the three final models are the subject of §6.5.4.

Model	**Thermal Model**	**Manual Model**		
	σ	μ_1	μ_2	μ_3
Observational Expectation	≈ 0.25	≈ 0.2	≈ 0.5	≈ 0.3
Imposed Prior	$0.5^{+0.25}_{-0.25}$	$0.34^{+0.25}_{-0.25}$	$0.34^{+0.25}_{-0.25}$	$0.34^{+0.25}_{-0.25}$
Baseline Model	$0.02^{+0.01}_{-0.01}$	$0.91^{+0.06}_{-0.05}$	$0.07^{+0.03}_{-0.05}$	$0.04^{+0.00}_{-0.03}$
SNIa Model	$0.02^{+0.01}_{-0.02}$	$0.88^{+0.07}_{-0.07}$	$0.08^{+0.05}_{-0.07}$	$0.05^{+0.01}_{-0.04}$
Basline Without Mg	$0.04^{+0.02}_{-0.03}$	$0.84^{+0.10}_{-0.10}$	$0.13^{+0.08}_{-0.10}$	$0.05^{+0.01}_{-0.04}$
SNIa Without Mg	$0.07^{+0.04}_{-0.04}$	$0.63^{+0.22}_{-0.23}$	$0.29^{+0.21}_{-0.21}$	$0.10^{+0.04}_{-0.08}$
Si Only	$0.04^{+0.03}_{-0.03}$	$0.72^{+0.18}_{-0.19}$	$0.21^{+0.15}_{-0.18}$	$0.09^{+0.04}_{-0.07}$

The top panel Figure 6.6 shows the actual value of the derived thick-disc ridges which result from the thermal likelihood estimation: we see that even though this set of parameters maximised $\mathcal{L}$, it still results in thick disc ridge values which lie outside of our acceptable abundance domain: Oxygen lies within the acceptable bounds for $\sigma < 0.007$, and Calcium for $\sigma < 0.005$, whilst Magnesium and Silicon are never within the acceptable bounds.

This shows that even the model parameters which maximise $\mathcal{L}$ are unable to simultaneously satisfy the constraints from the observational galactic properties. Put simply, we either have a fundamental flaw in our understanding of the thick disc - or perhaps more likely the Limongi & Chieffi (2018) yield grids are not capable of reproducing the observed properties of our galaxy under realistic galactic distribution functions.

6.5.3 Accounting For Additional Sources

In §6.3, we made the simplifying approximation that CCSN events from stars with $M > M_-$ are the only source of metals, which allowed us to ignore many of the complicating factors of GCE models, and which may have led us to prematurely reject the models.

We now make an attempt to account for this, by considering the effects that both AGB stars, and SNIa pollution would have on our inferred yields. We therefore propose a modification to our predicted yield abundances $[\mathrm{X/Fe}]_{\mathrm{predicted}} = \mathcal{P}_X$:

$$\mathcal{P}'_X(\boldsymbol{\theta}, S_X) = \mathcal{P}(\boldsymbol{\theta}) - S_X \tag{6.16}$$

Here $S_X > 0$ is an additional free parameter for each abundance [X/Fe]. The parameter is required to be positive since from Karakas & Lugaro (2016), AGB stars produce material

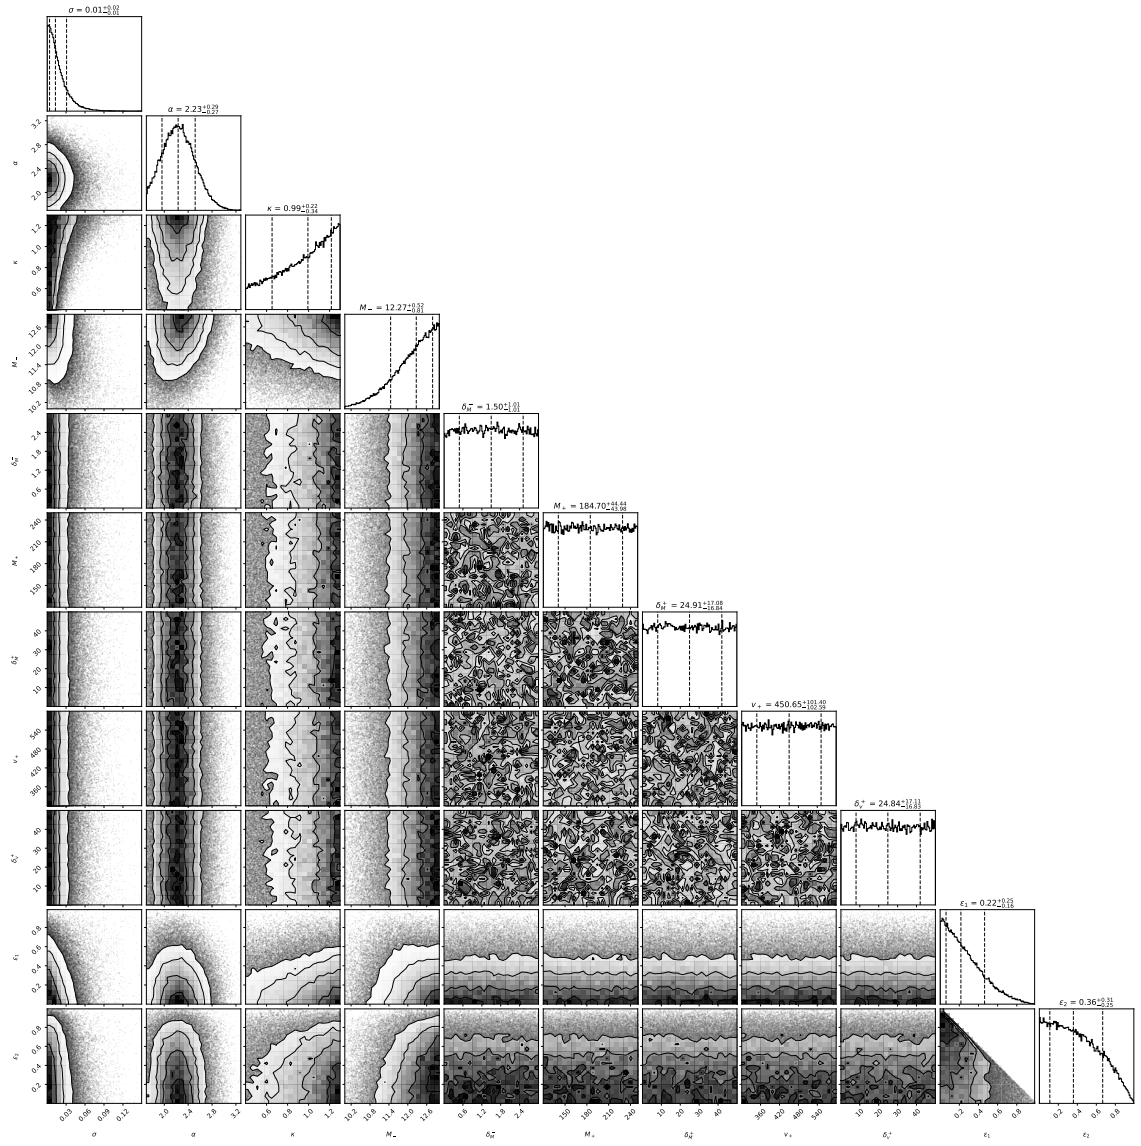

Figure 6.4: The corner plot showing the posterior distributions and cross-correlations of the baseline model parameter inference using the pseudo-thermal distribution function Eq. (6.5) to generate rotational weightings.

with [X/Fe] ≈ 0 for the elements of interest to us and from the W70 model of Iwamoto et al. (1999), SNIa produce almost pure iron-peak elements ([X/Fe] $\ll 0$). Therefore, the action of the AGB and SNIa contributions would be to *reduce* the final ridge value from the pure-CCSN abundance ratio $\mathcal{P}_X(\boldsymbol{\theta})$ derived in Eq. (6.3).

As SNIa are expected to be a larger contributor to the bulk chemical evolution than AGB (see the conclusions of Karakas & Lugaro (2016)), we shall assume that their properties will be the primary driver of any as-yet ignored behaviour.

Table 6.5 shows the [X/Fe] ratios from the Iwamoto et al. (1999) of the ejecta from

their W70 model (the literature standard) of a SNIa detonation. This data shows that some elements are produced in significantly higher quantities in SNIa than others: in particular [O/Fe], [Mg/Fe] $<$ -1, whilst [Si/Fe], [Ca/Fe] $>$ -0.5, meaning that we expect SNIa to have a far greater impact on the [O/Fe] and [Mg/Fe] abundance than [Si/Fe] or [Ca/Fe].

More modern SNIa yield predictions (i.e. Gronow et al. 2021) produce similar conclusions, implying that this is a general feature of SNIa modelling. Our prior on S_X is therefore based on this feature. We assign elements to one of two categories:

- 'Large Impact Elements' (LIE) with [X/Fe] ejecta < -1
- 'Small Impact Elements' (SIE) with [X/Fe] ejecta > -1

The prior is then uniform, with a hard cut at any point where a SIE offset is greater than a LIE offset, thereby ensuring that our fitted model of SNIa yields has approximately the same abundance pattern as those of the theoretical models:

$$p(S_X, S_Y) = \begin{cases} 0 & \text{if } S_X < S_Y \text{ and } S_X \in \{\text{LIE}\}, S_Y \in \{\text{SIE}\} \\ 1 & \text{else} \end{cases} \tag{6.17}$$

We emphasise that this is an incredibly weak constraint, and were we to rely on a more physical parameterisation, we would end up with more meaningful (but more model-dependent) results. We also note that Figure 6.6 would naively imply that this is a doomed endeavour: SNIa necessarily *lower* our predicted ridge values, but [Mg/Fe] already lies below the target region, however, this was a function of the best-fitting without the induced offset, to truly check if this works, we need to fully re-run the likelihood analysis and then make the comparison.

Figure 6.5 shows the inferred posterior PDFs for the chosen offsets. We note that although there are clear signals in the S_X parameters, there was no major alteration in the remainder of the inferred parameters, so we remain beholden to the low-rotation grid: the yellow curve in Figure 6.3 shows the posteriors on the spin parameters which we would then infer and the 'SNIa Model' row of table 6.4 shows the maximal values - whilst there is a slight increase in the allowed rotation speeds, we remain in strong tension with the observable data.

The dashed lines on Figure 6.6 shows the impact of the re-fitted models on the predicted thick disc ridges: as there was no major change in the fitted values, the functional shape of the lines remains mostly unaltered, with the only difference being the offsets. We see that Oxygen is now viable up to a larger temperature, though Silicon still remains over-abundant and the fit for magnesium has, as we naively predicted simply become worse. That the best-fit offset

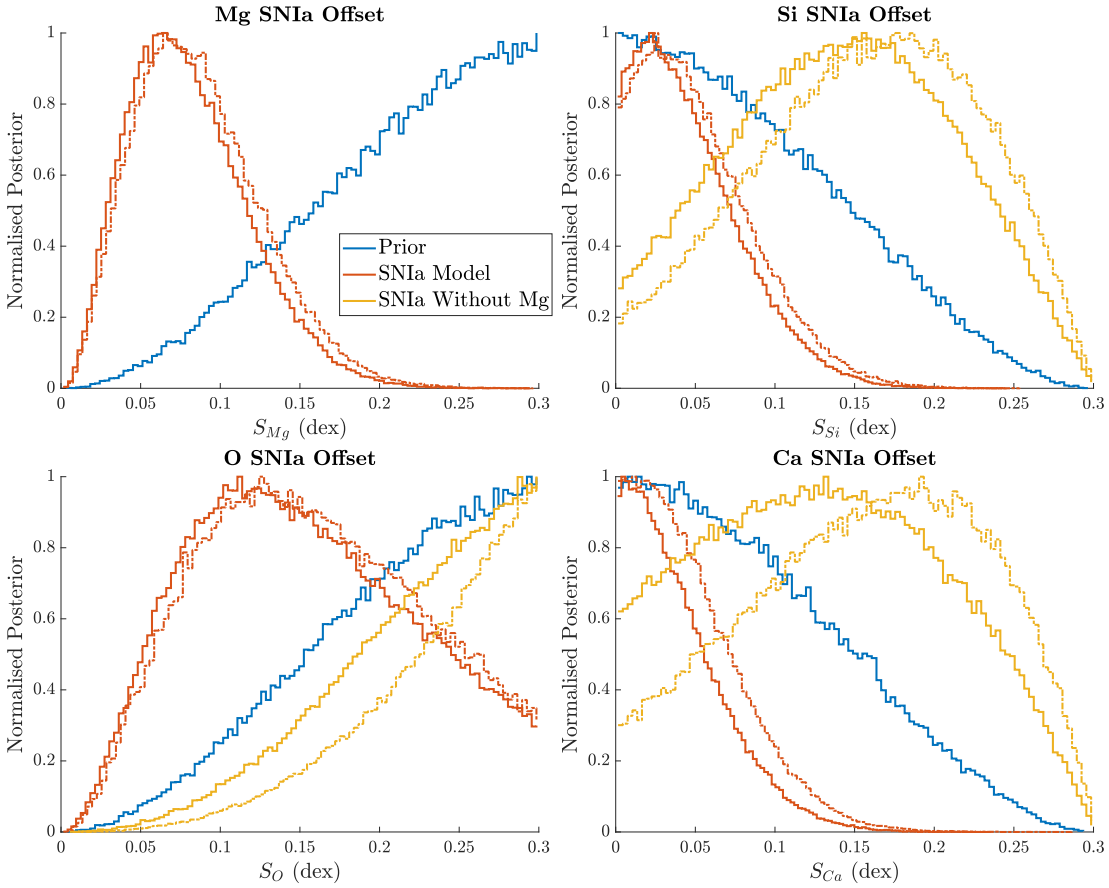

Figure 6.5: The posterior distributions of the required SNIa-induced chemical offsets, S_X, inferred for the simple models described in Eq. (6.16). Full curves are those derived with the thermal weighting, dashed curves are those from manually assigned rotation weightings.

Table 6.5: The [X/Fe] abundance of ejecta from SNIa following the W70 model of Iwamoto et al. (1999). Solar abundances are taken from M22.

Species	Synthesis ($M_\odot$)	Solar Density ($m_X^\odot/M_H^\odot$)	SNIa [X/Fe]
Fe	0.77	1.8×10^{-3}	–
O	0.133	7.8×10^{-3}	-1.41
Mg	0.0158	9.6×10^{-4}	-1.42
Ca	0.0181	8.8×10^{-5}	-0.32
Si	0.142	9.1×10^{-4}	-0.44

for Mg is strongly greater than 0 is a function of our prior and the fact that Mg is a 'Large Impact Element'. If $S_{\rm Mg}=0$, we would also require the offsets for SIE elements to be zero, however $S_{\rm Si}$ and $S_{\rm Ca}$ have a strong non-zero signal, leading to a tradeoff in the [Mg/Fe] value.

Even accounting for these additional sources of uncertainty, we are unable to reconcile the LC18 yield grids with the observed properties of our galaxy.

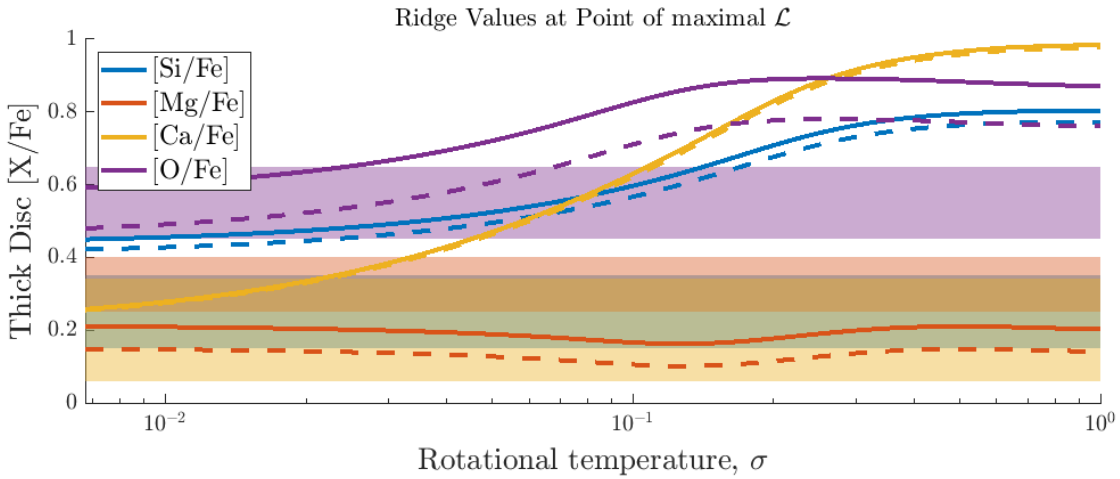

Figure 6.6: *(Solid)* The thick disc ridge values for the four studied elements as a function of rotation temperature σ. All other parameters are set to those which maximised $\mathcal{L}$ as per Figure 6.4. *(Dashed Lines)* as above, but with optimised offsets to model contributions of AGB stars or SNIa events - the model was refit to these new observational constraints.

6.5.4 Further Model Testing

In addition to the incorporation of SNIa/AGB offsets, we tested the optimisation routine on a number of other variant models, including omitting elements. In particular, motivated by the observation of the previous section that Mg seemed to be the limiting factor (i.e. we had [Mg/Fe] $<$ observed abundances, so SNIa offsets would only make the fit poorer), we ran the optimisation in both the original non-SNIa case, and including the SNIa offsets, without including Mg. The results of these additional tests are found in the bottom rows of Table 6.4.

Whilst omitting Mg did permit a slightly faster rotational distribution, and with substantially larger error bars, the resulting distributions were still significantly more slowly rotating than the observational constraints allow. We find that, contrary to our initial expectations, Si impacts the rotational inference just as strongly as Mg, which we verify in the case where we performing the optimisation only over Si, and find an equally slowly-rotating distribution.

6.6 Conclusion

In this chapter we developed a number of lightweight, yet powerful computational tools which allow rapid comparisons between a proposed explosive nucleosynthetic yield grid, and the observational abundances of the galactic thick disc. We propose that these tools can act as independent diagnostics for the underlying stellar evolution models from which the yield grids are derived.

Whilst it is not uncommon to for chemical evolution models to be used as constraints on stellar evolution, such models are often used in a rather anecdotal manner: to test physical inputs, a set of only very few chemical evolution models is run to produce some model chemical history. The drawing of conclusions is then treated rather as an art, where an expert needs to compare these trajectories to data and then judge if a better model can be found, and how well the model and data might fit or be in tension. By contrast, the low computational cost of our tools allows full exploration of the available parameter space and thus stringent negative inference, and the construction of powerful, quantative constraints on the input models. We tested these tools on the new rotating yield grids of Limongi & Chieffi (2018). These yield grids were of great interest to us as they had the potential of allowing us to replace the heterogeneous yield grids of the RAMICES II model (§4.6) with state-of-the-art modern yield grids which incorporated the new stellar initial condition of rotation.

However, we found that of the four yield grids provided by LC18, three of them (sets 'F', 'I' and 'M') failed to reproduce observed thick disc abundance values even when we left the rotational distribution function and the galactic enrichment history (hence hence $Z(t)$) almost entirely unconstrained.

We subjected the remaining yield grid (set 'R') to additional scrutiny and discovered that even under the addition of the weak constraint that the distribution functions must be the same for all elements, the best-fitting rotational distribution function was highly peaked around $v_{\rm zams} = 0 {\rm kms}^{-1}$, in tension with the evidence that stars are born with very high break up speeds. This conclusion was unaltered even when we relaxed our canonical assumption (§2.2) that the thick disc structure arises from pre-SNIa, pre-AGB enrichment.

We conclude that the LC18 yield grids are incompatible with the observed galactic chemistry, which must in turn mean that either our understanding of the origin of the thick disc is irreconcilably flawed, or that the underlying models used to generate the LC18 yields is flawed. On the balance of evidence, we are confident that the flaw lies with the stellar evolution models.

We again emphasise that this should not be taken as a criticism of the methods of LC18, their first-principles method and open publication policy is what allows us to use the tools of galactic chemistry as an independent diagnostic. Given their prior policy (Chieffi & Limongi (2004)) of publishing the full range of yield grids based on different mass cuts, we suggest further investigation of how their choice of a singular mass cut in these models impacts their yield grids.

The RoSY tool developed here is publicly available[3], and we anticipate its future use to constrain these models in a rapid and flexible fashion, and thereby allow stellar modellers access to galactic level constraints, and thereby improve the accuracy of stellar modelling in future.

[3] https://github.com/JTFraser/RoSY

Life itself is an exercise in exceptions.

Jean-Luc Picard

7

Hoops, Loops & Chemical Evolution

Contents

Having developed the powerful RAMICES II tool in Chapter 4, we now utilise the model to briefly discuss the nature of 'gas caustics'. Caustics are sets of trajectories in the gas phase Tinsley Diagrams which, when composited together, produce global trends not seen in the individual paths, and hence can cause the resulting stellar Tinsley Diagrams to look very different to the gas phase counterparts.

7.1 Caustics & NSM

In Chapter 5, we focused in detail on the nature of r-process nucleosynthesis, and demonstrated conclusively that collapsar events cannot contribute meaningfully to the total Eu budget. This leaves the potential for a significant gap in our understanding of the source of r-process material, since collapsars are often invoked precisely because the other major model – NSM – produce chemical evolution tracks which (seemingly) do not match with the observational evidence. In this section, we demonstrate that this argument is at least partially flawed.

7.1.1 The Problem: Gas-Hoops

Since NSM, by necessity of their formation mechanism, require a binary pair to undergo supernova (and survive) and then inspiral via gravitational wave emission, there is a delay between the time when the first CCSN occur, and when the majority of the associated NSM events occur. It therefore follows that a galaxy which sources its r-process material entirely by

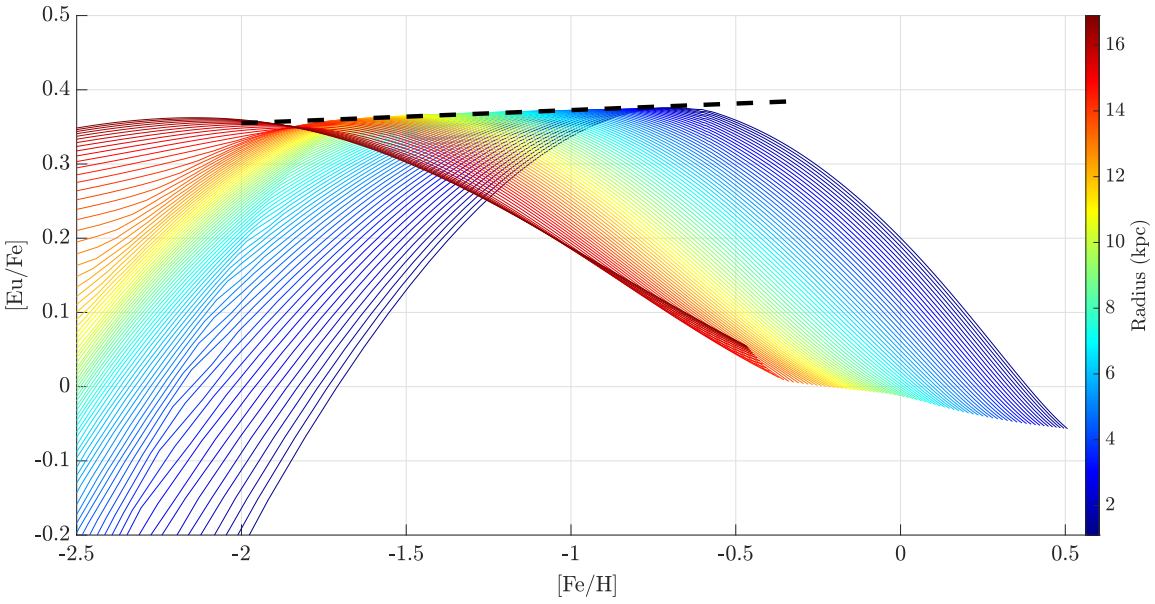

Figure 7.1: Evolution of [Eu/Fe], where NSM are the sole r-process source, derived from the nominal RAMICES II model. Distinctive hoop-structures are observed, forming a 'gas caustic', denoted by the black dashed line.

NSM would start out in a very low state of [Eu/Fe], which would then increases to a maximum as the NSM 'catch up' to the CCSN, before [Eu/Fe] decreases again at the SNIa turn on.

This profile is clearly seen in Figure 7.1, the Tinsley diagram for [Eu/Fe] from our nominal RAMICES II model. These 'hoops' are distinctive indicators of NSM-dominated enrichment, and are widely seen as evidence that NSM-dominated models are incompatible with galactic chemistry, since, as shown in Figure 5.1, no such hooped-structure is visible in the observations.

7.1.2 The Solution: Caustic-Induced Flattening

However, whilst each individual annulus has a strong hooped structure, we see that the continuum metallicity gradient of the galaxy causes the peaks of each individual hoop to form an almost-flat line in [Eu/Fe]. We term this confluence of many different chemical evolution tracks into a single cohesive structure a 'caustic', in parallel with the familiar structure in optics.

Figure 7.2 shows how this caustic translates into the synthesised stellar catalogue. Despite the heavily 'arched' nature of the individual annuli observed in 7.1, the stellar catalogue shows a remarkably flat value of [Eu/Fe]$\sim$ 0.4 between [Fe/H] = -1.75 and -1, due to the coherent addition of stars at different annuli along the caustic, rendering the model a far better fit to the observable data than it would otherwise have seemed to be.

The length of the stellar caustic is a function of the birthplace-distribution of stars included in the survey, incorporating both radial migration, and also the volume covered by the survey.

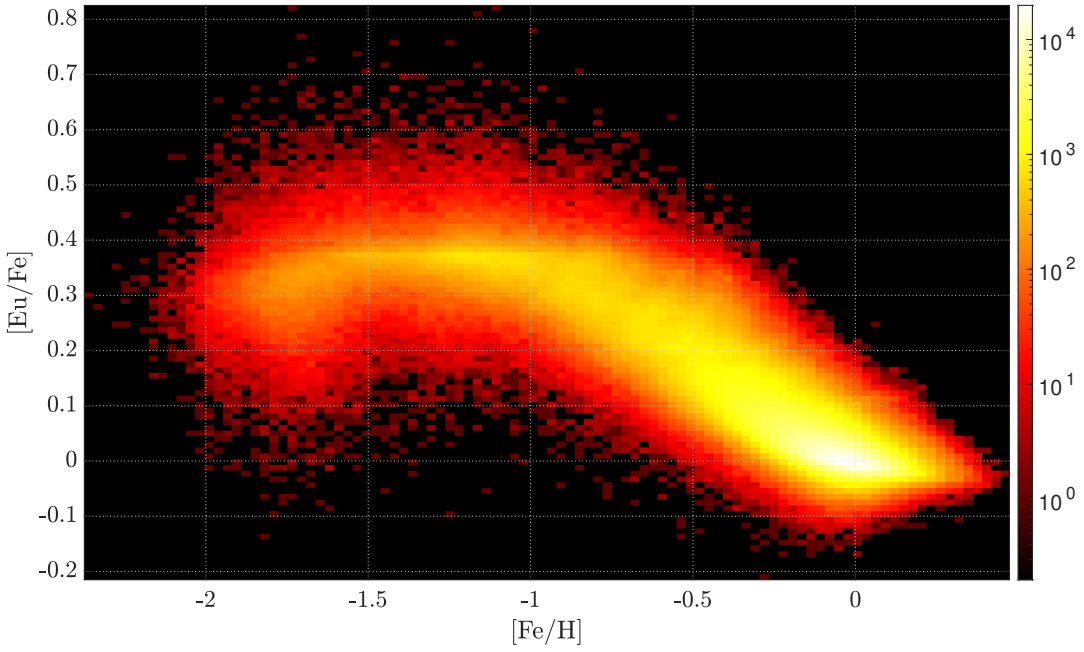

Figure 7.2: The stellar catalogue associated with the cold gas abundances of 7.1. We note that, although the strongly 'arched' gas patterns in Figure 7.1 would seem to contradict the observed evidence of a flat [Eu/Fe] curve at early times (see Figure 5.1), the associated stellar catalogue shows a remarkable consistency with the observed data.

If we were able to limit our survey only to stars born exactly at the solar annulus, we would recover the hooped structure. Of course, limiting the survey in this fashion is exactly the same as studying only the gas phase diagram. Conversely, if we were able to extend our survey to include every star in the model, we would not see the stellar caustic extend to [Fe/H] ~ -2.5, as the density of stars at radii >12kpc is too low.

7.1.3 Summary

When one considers that a) the galaxy has a metallicity gradient and b) even the smallest volume-limited survey includes stars born outside the precise solar annulus, the presence of a gas caustic can drastically alter the predicted stellar distribution.

Importantly, whilst it would have seemed self-evident that Figure 7.1 would produce a solar neighbourhood stellar catalogue with [Eu/Fe] structures incompatible with observed evidence, Figure 7.2 showed that a physical stellar catalogue which includes finite survey volume and stellar migration renders such a conclusion false.

We emphasise that this is *not* proof that NSM are dominant sources of r-process nucleosynthesis, it merely serves as a rebuttal to the argument that the observed [Fe/H] < -1

behaviour of [Eu/Fe] strongly prohibits NSM from being vibale candidates.

7.2 Dilution-Driven Evolution

In recent years, a new class of models has gained favour with some GCE groups (Grisoni et al., 2018; Palla et al., 2020; Spitoni et al., 2020, 2021). Under these models, the two infall prescription, Eq. (2.15), is modified such that there is a distinct time delay before the secondary onfall occurs:

$$\dot{M}_{\rm onfall} = a \exp\left(-\frac{t}{\tau_1}\right) + b\Theta(t - t_{\rm merger}) \exp\left(-\frac{t}{\tau_2}\right) \tag{7.1}$$

Here $\Theta(x)$ is the Heaviside step function. The values of a, b, $t_{\rm merger}$, τ_1 and τ_2 are then found by parameter optimisation searches, with Spitoni et al. (2021) finding $b/a \approx 5.6$, $\tau_1 = 0.1$Gyr, $\tau_2 = 4.1$Gyr, $t_{\rm merger} = 4$ at the solar annulus. Under such models, the early-time evolution proceeds not too differently from our own, however, at the moment of the turn-on, a fresh influx of gas rushes into the galaxy. As the material is chosen to be more metal poor than the galaxy (see Spitoni et al. (2021), where they choose [Fe/H] of the onfall to be -0.5), this causes a rapid dilution of [Fe/H], which is followed by a rapid burst of star formation. Such a rapid burst of star formation causes CCSN events to outstrips SNIa for a brief time, and hence drive [Mg/Fe] values up towards their thick disc values. The SNIa then 'catch up', and drive the galaxy back to higher [Fe/H] and lower [Mg/Fe]. In the extreme case this can cause a self-intersecting 'Matteucci Loop', demonstrated in our models in Figure 7.3, and qualitatively very similar to Figure 3 of Spitoni et al. (2021).

We note that, although the 'loop' structure is highly unusual, the final [Mg/Fe] and [Fe/H] spatial gradients that the model recovers are consistent with those of Figure 2.5, seeming to indicate that the models are plausible. We also see that the peak of the 'loops' exhibits a similar caustic structure to the 'hoops' of §7.1, albeit demonstrating a significant radial gradient.

We do issue a slight note of caution in comparing the models in Figure 7.3 and those in i.e. Spitoni et al. (2021), insofar as their models derive the entire radial gradient by calibrating the values of τ_i, $t_{\rm merger}$ etc. within three spatial bins (2-6kpc,6-10kpc and 10-14kpc), with each bin assumed to be spatially uniform and non-interacting. The models in Figure 7.3, however, use only a singular value of τ_i etc., and keep the usual RAMICES II radial inflow and migration prescriptions. However, given the magnitude-limited nature of our GALAH selection survey, the majority of the stars in a given synthetic catalogue are from the solar neighbourhood. By taking the solar neighbourhood values of the τ_i, $t_{\rm merger}$, we will therefore synthesise stars within the appropriate radial ranges, and hence permit a valid comparison.

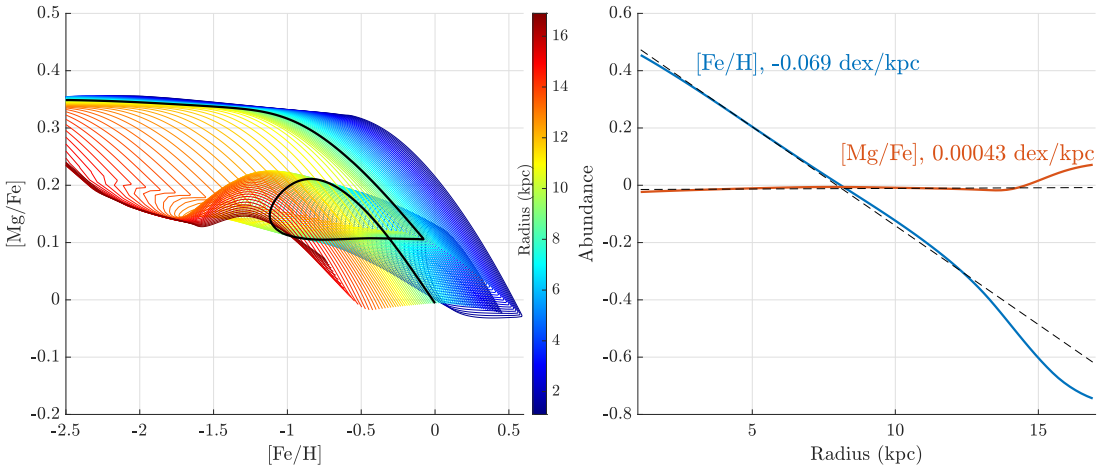

Figure 7.3: The Magnesium Tinsley diagram for the cold gas abundances resulting from a dilution-driven evolution model with the parameters from Spitoni et al. (2021). For clarity, the solar annulus path is highlighted with the solid black line. The model is otherwise our best-fit model demonstrated in Figure 4.8

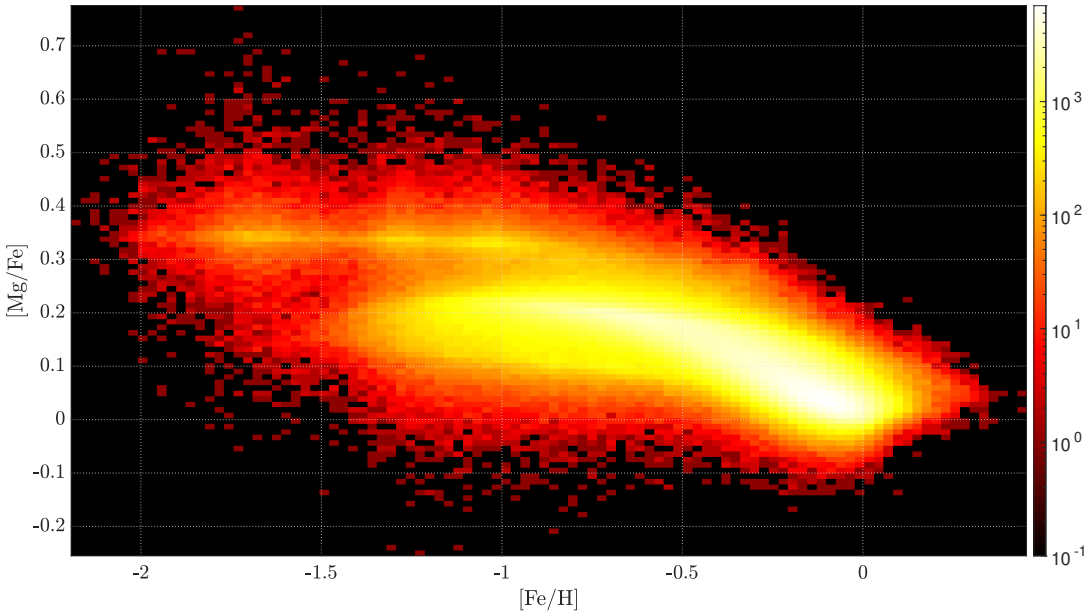

Figure 7.4: The associated stellar population from Figure 7.3.

Figure 7.4 shows the stellar distribution we infer from the gas abundances in Figure 7.3, and we immediately see a significant problem: a secondary ridge has formed in [Mg/Fe] ~ 0.2, which is, simply, nowhere to be seen in the observational data.

The papers which favour this model argue that because the galaxy passes through this dilation/loop phase rapidly, leaving only a negligible imprint on the stellar catalogue, with the stellar distribution only meaningfully affected by the subsequent downwards knee. The effect of the loop is merely to increase the scatter around the mean value within the thin disc structure (Grisoni et al., 2018; Palla et al., 2020; Spitoni et al., 2020, 2021).

However, Figure 7.4 shows that this is not reconcilable with a high-resolution model since the presence of a gas caustic has gives rise to a structure which is not supported by the observational evidence. This structure arises because all of the annuli pass through the peak of the loop coherently: one SNIa-delay time after the infall starts, with [α/Fe] values approximately halfway between the thick and thin disc values, but at slightly different values of [Fe/H], due to the metallicity gradient of the galaxy.

We speculate that the reason this was missed in the prior studies of these models is due to the extremely coarse spatial binning (4kpc per bin), significantly larger than the scale on which the abundance gradients are manifest. As we have seen, gas caustics imprint themselves upon stellar distributions at points where 'loops' reach peak values which are very slightly offset in metallicity at different radii. By removing these 'nearby radii', the effect of the caustic is suppressed.

7.3 Caution & Caustics

In the two brief studies presented above, we demonstrated that radial abundance gradients between nearby galactic annuli can cause gas caustics to imprint structures into the stellar abundance gradient which are not seen in any individual annulus. In both cases, this drastically altered the conclusions one would draw from the gas phase diagram alone. Of course, this is not a particularly novel phenomena, it is well known that the stellar sample we observe is formed from a convolution of many different gas reservoirs, with varying densities and stellar lifetimes. This was, in fact, exactly why the stellar synthesis portion of the RAMICES II was constructed: to enable direct comparisons between stellar samples.

It is for reasons of simplicity than many models of GCE (SACEM included), opt to study only the gas phase abundances: tracking the full stellar populations and their abundances is a computationally intensive problem. In the majority of cases, gas caustics are not

present, and hence the comparison between the gas phase abundances and the stellar abundances holds no such surprises.

All of the gas-phase comparisons presented in this thesis have subsequently been tested against the presence and impact of gas caustics, and the results – outside of those presented in this chapter – were found to be negligible. In the majority of models, the evolution of [X/Fe] is smooth and hence no gas caustics form, it is only in unusual edge cases where the chemical gradient of the galaxy allows such coherent behaviour.

However, the existince of these edge cases shows that it is nevertheless important to ensure that any future studies are aware of the impacts that gas caustics can have on stellar distributions, and that such studies are tested to ensure they are not falling prey to such errors.

The stars don't look bigger - but they do look brighter

— Sally Ride

8 Concluding Remarks

Contents

8.1 Conclusions

This thesis demonstrates several new tools and techniques for studying the GCE of our galaxy. Our analytical model, SACEM, was introduced in Chapter 3 and allows us to scan the high-dimensional GCE parameter space with unprecedented resolution. The model is designed to be highly flexible and minimally constrained, with the view that merely replicating subsets of the observational data introduced in Chapter 2 is a trivial exercise given the number of highly unconstrained parameters present in such an analytical GCE model. For this reason, SACEM is most useful as a tool of negative inference: where data cannot be replicated for any set of parameters in a highly flexible model, we must infer that the model is truly incompatible with the data.

Chapter 4 contrasted this approach with the more realistic RAMICES II, an update to the RAMICES model of SB09, with increased performance and allowing us to probe deeper into the internal physics. RAMICES II fully simulates the evolution and lifetime of the galaxy and the stars within it, including radial migration over the course of the stellar lifetime. We also attempt to place physically meaningful constraints on the behaviour of the circum-galactic medium, by modelling it as a gas reservoir surrounding the galaxy, enriched by material ejected during supernova and feedback events.

We then utilised these models to investigate a number of open questions, in both the field of GCE and beyond. Chapter 5 focused on the origin of r-process material, showing that the necessary metallicity dependence on the formation of collapsars prohibits them from contributing meaningfully to the global r-process budget at late times. This was achieved by running more than 10^{12} iterations of the SACEM code, and more than 5000 combined simulations in both RAMICES and RAMICES II, making the exploration of model-space extremely robust, and the conclusions overwhelming in their nature. Although the excluded region of parameter space incorporated all models in which collapsars were the dominant source of r-process material, we did not find this to be true for NSM. In the light of the observational evidence of r-process ejecta surrounding GW170817, we therefore concluded that NSM were the likely source of the majority of r-process material in the universe.

In Chapter 6 we utilised the data of GCE to identify flawed models of stellar evolution. By utilising the early-time, CCSN-dominated portion of SACEM, we were able to place extremely strong constraints on the ability of all four of the proposed rotating stellar models used by Limongi & Chieffi (2018), despite the constraints on these models being (again, intentionally) incredibly lax. This powerful tool, ROSY, is publicly available and designed to allow future stellar evolution models to be similarly constrained. In doing so, we have used GCE as a powerful, independent diagnostic on the internal nuclear and plasma physics in the hearts of stars.

Finally, Chapter 7 briefly used the RAMICES II to demonstrate the phenomena of 'gas caustics', in which strongly varying, but metallicity-offset paths in the gaseous abundance plane produced global structures in the stellar abundance plane through a process of coherent addition. We used this phenomena to show how NSM-dominated models could show unusually flat trajectories in [Eu/Fe], and therefore strongly support our argument from Chapter 5 that NSM are the dominant r-process source, as well as addressing why the 'Matteucci Loop' cannot be considered physical.

8.2 Future Work

The work presented here is by no means the end of this avenue of investigation: rather, it is simply the first few steps utilising the new tools we have demonstrated here. There are many future paths forward from the work we have presented here.

8.2.1 Refinement of RAMICES II

Whilst RAMICES II has already proven invaluable in the study of GCE, we note that there are several key improvements that could be made to the simulation, before we release it for general use.

Firstly, upon comparing our synthesised stellar catalogue with the GALAH data of Figure 2.2 it is clear that our model is not fully replicating the observed abundances of stars. Whether this is because the fundamental underlying mechanisms and calibration of the model is flawed in some way (though we note that our gas-content evolution is similar to other literature values, so this is unlikely), or because the mechanisms for selecting stars belonging to our pseudo-GALAH catalogue are flawed, is as yet-undetermined. Expanding our observational catalogue to include GCS (Nordström et al., 2004), or APOGEE (Ahumada et al., 2020) would allow us to make this determination, and hence represents a vital step in improving the simulation.

Secondly, for consistency of comparisons, we used the same set of nucleosynthetic yield grids as the original SB09 work. This is in spite of more than a decade of advancement in our understanding of stellar structure and supernova physics. The original work which spawned Chapter 6 was intended to replace these yield grids with the state-of-the-art, in the form of the Limongi & Chieffi (2018) yield grids. However, as we showed in that chapter, this particular yield cannot be accurate, and so the original yield grids were retained.

Although the LC18 yields proved incompatible, other modern yield grids, such as Karakas & Lugaro (2016) (which provide AGB yields for an extensive set of elements, including the s-process, at high metallicities) may provide additional insight into the processes which drive late-time GCE. With the ROSY tool in hand, we are also in a position to rapidly identify any newly published yield grids as viable candidates, and hence can rapidly incorporate them into our models as and when they become available.

8.2.2 Looking Inward: Nuclear Stellar Disc

One of the motivating factors for the development of RAMICES II was the inflexibility of prior iterations. The study radical and unusual portions of the Milky Way, such as the kinematically cold Nuclear Stellar Disc in the central few parsecs of the Milky Way, which is nevertheless

intimately connected to the global evolution of the galaxy through its inflow-dominated gas-dynamics, was beyond the scope of traditional models to handle.

RAMICES II was designed specifically with this kind of flexibility and reparameterisation in mind. The work in this thesis therefore lays a very strong groundwork for studying the chemical evolution of this unusual laboratory at the heart of the galaxy.

8.2.3 Looking Upward: The CGM & Galactic Halo

The nature of the interaction between the galaxy, its gaseous halo and the IGM is poorly constrained. In Chapter 4, we placed a number of constraints on where, and in what quantity and composition, new gas fell into the galaxy. However, whilst this prescription was motivated to model physical observables, the underlying mechanics are not physically motivated.

There is therefore vast scope for exploring the space of galaxy/CGM interactions through the lens of GCE, a field which has often been left unexplored on a galactic scope for being too complex, and poorly resolved on a cosmological scale for being too small-scale.

8.2.4 Looking Outward: Dwarf Galaxies

Much of the work in this thesis has solely focused on the Milky Way. This is partly because the Milky Way is our cosmological home, and hence of special interest when it comes to understanding the origins of the material and the process which drives its evolution, but also because the highest quality spectroscopic abundance data is available for the Milky Way.

Modern observational evidence, however, has hinted at intriguing trends in, for example, [Eu/Fe] in dwarf galaxies, which may be an indicator of differing interactions between the hot gas phase and the CGM which depend upon the mass of the dwarf galaxy. Both SACEM and RAMICES II are ideal candidates for studying this question, which may also allow us to disentangle parts of the canonical models from general principles of GCE, and those which are a property unique to the Milky Way.

Appendices

A

Mathematical Proofs & Definitions

A.1 Stellar Observation Integral

In §4.7, we introduced the integral Eq. (4.15), which must be computed in order to determine the number of stars within a given RAMICES II annulus which satisfy the cuts for a given survey selection function.

Computing Eq. (4.15) is nominally an arduous task requiring a 3 dimensional numerical integral (r, θ, ϕ), which must itself be computed for each annulus R_i, and over a large number of magnitudes, M_V and ages τ. However we note that the computation can be simplified, as the vertical integral can be analytically computed. We split the integral into components above and below the plane, such that:

$$\mathcal{I}(r, \varphi, M_V, \tau) = \int_{-\infty}^{\infty} \mathrm{d}z S\left[d(\rho, \varphi, z), M_V\right] n_i(\rho, \varphi, z, \tau) \tag{A.1}$$

$$= \frac{1}{4\pi R_i w_i \bar{z} \tau} \left(\int_{a_-}^{a_+} \mathrm{d}z e^{-|z|/\bar{z}(\tau)} + \int_{b_-}^{b_+} \mathrm{d}z e^{-|z|/\bar{z}(\tau)} \right) \tag{A.2}$$

$$= \frac{1}{4\pi R_i w_i} \left(e^{-a_-/\bar{z}} + e^{-b_-/\bar{z}} - e^{-a_+/\bar{z}} - e^{-b_+/\bar{z}} \right) \tag{A.3}$$

Here $a_\pm$ and $b_\pm$ are the maximum and minimum distances above ($a_\pm$) and below ($b_\pm$) the plane which the photometric and angular cuts allow into the integral. With $d_p = \sqrt{d_\odot^2 + r^2 - d_\odot r \cos(\varphi)}$ the in-plane distance to the star, $d_+ = 10^{\frac{4-M_v}{5}}$ kpc the maximum distance at which the target star is in the magnitude range and $d_- = 10^{\frac{2-M_V}{5}}$ the corresponding

minimum distance, the values of $a_\pm$ and $b_\pm$ are found from:

$$z_{\rm sc} = d_p \tan(b(\ell)) \tag{A.4}$$

$$a_- = \sqrt{\max\left[d_-^2 - d_p^2, d_p \tan(10^\circ)\right]} \tag{A.5}$$

$$a_+ = \max\left[a_-, \min\left[z_{\rm sc}, \sqrt{\max\left[d_+^2 - d_p^2, 0\right]}\right]\right] \tag{A.6}$$

$$b_+ = \sqrt{\max\left[d_+^2 - d_p^2, 0\right]} \tag{A.7}$$

$$b_- = \min\left[b_+, |\min\left[-a_-, z_{\rm sc}\right]|\right] \tag{A.8}$$

Here the value of $\tan(b(\ell))$ used for $z_{\rm sc}$ is given by Eq. (4.17). Whilst unpleasant to write down in a formal mathematical sense, these terms are in fact much more intuitive in a computational sense - for example, the statement $\sqrt{\max\left[d_+^2 - d_p^2, 0\right]}$ ensures that if $d_p > d_+$ (i.e. the star is further away than the maximum visible distance), then the integral is not computed, as there is no height above or below the plane at which it would be a valid survey target.

A.2 Interpolated Integrals

In §6.3.2 we introduced the interpolation/extrapolation scheme necessary to integrate the LC18 yield grids over the full, continuous domain of mass, metallicity and rotation. Here we derive the method for generating the weights of Eq. (6.7). We suppose that we have a function $f(x)$ which performs our inter/extrapolation method on the dataset of points $\{(x_i, y_i)\}$ labelled from $1, ...N$ such that $x_1 < x_2 < .... < x_N$:

$$f(x) = \begin{cases} 0 & \text{if } x > L_+, x < L_- \\ f(x_1 - \delta_-) & L_- \leq x < x_1 - \delta_- \\ y_1\left(\frac{x_2 - x}{x_2 - x_1}\right) + y_2\left(\frac{x - x_1}{x_2 - x_1}\right) & x_1 - \delta \leq x < x_2 \\ y_i\left(\frac{x_{i+1} - x}{x_{i+1} - x_i}\right) + y_{i+1}\left(\frac{x - x_i}{x_{i+1} - x_i}\right) & x_i \leq x < x_{i+1},\ 2 \leq i < N \\ y_{N-1}\left(\frac{x_N - x}{x_N - x_{N-1}}\right) + y_N\left(\frac{x - x_{N-1}}{x_N - x_{N-1}}\right) & x_{N-1} \leq x < x_N + \delta_+ \\ f(x_N + \delta_+) & x_N + \delta_+ \leq x \leq L_+ \end{cases} \tag{A.9}$$

The integral $\int f(x)\zeta(x)\mathrm{d}x$ will therefore simply to a weighted sum of the y_i terms, with the weighting terms only depending on the x_i sampling, and the zeroth and first moments of $\zeta(x)$. To that end, we define the functionals:

$$\alpha[\zeta, y_1, y_2] = \int_{y_1}^{y_2} \zeta(x)\mathrm{d}x \tag{A.10}$$

$$\beta[\zeta, y_1, y_2] = \int_{y_1}^{y_2} x\zeta(x)\mathrm{d}x \tag{A.11}$$

We now assume that the bounds of the integral always extend beyond $L_{\pm}$, and that the extrapolation constants $\delta_{\pm}$ are truncated such that $L_{-} \leq x_1 - \delta_{-}$ and $L_{+} \geq x_N + \delta_{+}$. Defining the following intermediary quantities as follows:

$$\ell_N = x_N + \delta_{+} \tag{A.12}$$

$$\ell_1 = x_1 - \delta_{-} \tag{A.13}$$

$$J_i = \begin{cases} 0 & i = 1 \\ \frac{\beta[\ell_1, x_2] - x_1 \alpha[\ell_1, x_2] - \delta_{-} \alpha\left[L_{-}, \ell_1\right]}{x_2 - x_1} & i = 2 \\ \frac{\beta\left[x_{N-1}, \ell_N\right] - x_{N-1} \alpha\left[x_{N-1}, L_{+}\right] + \ell_N \alpha\left[\ell_N, L_{+}\right]}{x_N - x_{N-1}} & i = N \\ \frac{\beta\left[x_{i-1}, x_i\right] - x_{i-1} \alpha\left[x_{i-1}, x_i\right]}{x_i - x_{i-1}} & \text{else} \end{cases} \tag{A.14}$$

$$K_i = \begin{cases} \frac{x_2 \alpha[L_{-}, x_2] - \beta[\ell_1, x_2] - \ell_1 \alpha[L_{-}, \ell_1]}{x_2 - x_1} & i = 1 \\ \frac{x_N \alpha[x_{N-1}, \ell_N] - \beta[x_{N-1}, \ell_N] - \delta_{+} \alpha[\ell_{+}, L_{+}]}{x_N - x_{N-1}} & i = N - 1 \\ 0 & i = N \\ \frac{x_{i+1} \alpha\left[x_i, x_{i+1}\right] - \beta\left[x_i, x_{i+1}\right]}{x_{i+1} - x_i} & \text{else} \end{cases} \tag{A.15}$$

It then follows that:

$$\int_{-\infty}^{\infty} f(x) \zeta(x) = \sum_{i=1}^{N} \left[J_i + K_i\right] y_i := \sum_{i}^{N} \nu_i y_i \tag{A.16}$$

And:

$$\sum_{i=1}^{N} \nu_i = \int_{L_{-}}^{L_{+}} \zeta(x) \mathrm{d}x \tag{A.17}$$

Hence, if $\zeta(x)$ is suitably normalised to 1 between these bounds, the integral weightings also sum to 1. Though more cumbersome to prove, it also follows that $\nu_i \geq 0$ as long as $\alpha, \beta \geq 0$, which follows from $\zeta(x)$ being a well defined PDF.

A.3 Bounding Yield Ratios

In §6.4, we devised a way to express the yield ratios in terms of arbitrary distribution functions across our gridded metallicity and rotation space:

$$\left[\frac{\mathrm{A}}{\mathrm{B}}\right] = \log_{10}\left(\frac{\sum_{i,j} \epsilon_i \mu_j \bar{y}_A(v_j, Z_i)}{\sum_{i,j} \epsilon_i \mu_j \bar{y}_B(v_j, Z_i)}\right) - \mathcal{R}_{\odot} \tag{A.18}$$

Where μ_j are the set of n_β rotation weightings and ϵ_i the set of n_Z metallicity weightings, $\mathcal{R}_{\odot}$ is the associated solar abundance value, and $\bar{y}_{A,i,j}$ are the IMF-integrated yield functions associated with a star of metallicity Z_j and ZAMS rotation v_i.

We wish to bound this function, with the only additional constraints being that $\{\epsilon_i\}$ and $\{\mu_j\}$ behave as if they were derived from a distribution function: namely that $\epsilon_i, \mu_i \geq 0$ and $\sum_i \epsilon_i = \sum_j \mu_j = 1$. We achieve these results by parameterising the vectors as follows:

$$\mu_i = \frac{\exp(x_i)}{\sum_j \exp(x_j)} \quad \text{and} \quad \epsilon_i = \frac{\exp(y_i)}{\sum_j \exp(y_j)} \tag{A.19}$$

$\mathbf{x}$ and $\mathbf{y}$ are then completely unconstrained vectors on $\mathbb{R}^{n_\beta}$ and $\mathbb{R}^{n_Z}$ respectively. Ignoring the solar offset, and utilising the monotonicity of the logarithm function, we therefore write the function to be extremised as:

$$\mathcal{I} = \frac{e^{\mathbf{y}} \cdot A e^{\mathbf{x}}}{e^{\mathbf{y}} \cdot B e^{\mathbf{x}}} := \frac{\alpha(\mathbf{x}, \mathbf{y})}{\beta(\mathbf{x}, \mathbf{y})} \tag{A.20}$$

Where A and B are the matrices obeying $X_{ij} = Y_X(v_i, Z_j)$ and $\mathbf{b} = e^{\mathbf{a}}$ is shorthand for $b_i = \exp(a_i)$. We note that although the normalising factor $\sum_i x_i$ factored out under division, it did so only under the condition that this factor was non-zero: solutions where $e^{\mathbf{x}} = \mathbf{0}$ are therefore disallowed. Similarly, the conditions on $\mathbf{x}, \mathbf{y}$ and Y_{Xij} mean that $\alpha, \beta > 0$.

Extremising this function gives:

$$\frac{\partial \mathcal{I}}{\partial \mathbf{y}} = \frac{e^{\mathbf{y}} \otimes \left[(A - \mathcal{I}B)\, e^{\mathbf{x}} \right]}{\beta} \tag{A.21}$$

$$\frac{\partial \mathcal{I}}{\partial \mathbf{x}} = \frac{e^{\mathbf{x}} \otimes \left[(A - \mathcal{I}B)^T e^{\mathbf{y}} \right]}{\beta} \tag{A.22}$$

Where $(\mathbf{a} \otimes \mathbf{b})_i = a_i b_i$ is the elementwise Hadamard Product. When the derivatives are zero, $(\mathbf{x}_e, \mathbf{y}_e)$ we have that $\mathcal{I} = \mathcal{I}_e$, the extremal ratio. The easiest way to find an extremal value is if $\mathbf{x}, \mathbf{y} \to -\infty$, however the condition that $e^{\mathbf{x}}, e^{\mathbf{y}} \neq \mathbf{0}$ forbids us from sending all elements to zero.

The solution must therefore have $e^{\mathbf{x}}$ with $0 < m_\beta \leq n_\beta$ non-zero elements, and similarly $e^{\mathbf{y}}$ must have $0 < m_Z \leq n_Z$ non-zero elements, chosen such that the second term in the numerators of Eq. (A.21) and Eq. (A.22) are zero.

The simplest way to achieve this is if there is only one non-zero element of each of $e^{\mathbf{x}}, e^{\mathbf{y}}$ - the p and q-th elements respectively. The above equations would then simplify to:

$$\alpha = A_{qp} e^{y_q + x_p} \tag{A.23}$$

$$\beta = B_{qp} e^{y_q + x_p} \tag{A.24}$$

$$\frac{\partial \mathcal{I}}{\partial \mathbf{y}} = \frac{\partial \mathcal{I}}{\partial y_q} \hat{e}_q = \beta^{-1} e^{y_q} \left(A_{qp} - \mathcal{I} B_{qp} \right) e^{x_p} \tag{A.25}$$

$$\frac{\partial \mathcal{I}}{\partial \mathbf{x}} = \frac{\partial \mathcal{I}}{\partial y_p} \hat{e}_p = \beta^{-1} e^{x_p} \left(A_{pq} - \mathcal{I} B_{pq} \right)^T e^{y_q} \tag{A.26}$$

The extremal value of $\mathcal{I}$ is then fixed by the requirement that the derivatives be zero and all other components are positive definite quantities:

$$\mathcal{I} = \frac{A_{qp}}{B_{qp}} \tag{A.27}$$

The null-derivative conditions are satisfied automatically, for any chosen value of q, p, y_q or x_p. We have shown the pointwise-ratios of the component matrices are extremal values of the function.

However: this is not the full list of extremal points, as we have not considered solutions where there is more than one non-zero element in our vectors (the sets $\{e^{q_i}\}$ and $\{e^{p_i}\}$ respectively), but which still satisfy the conditions

$$\sum_j \left(A_{p_i j} - \mathcal{I} B_{p_i j}\right) e^{z_j} = 0 \tag{A.28}$$

$$\sum_j \left(A_{j q_i} - \mathcal{I} B_{j q_i}\right) e^{v_{q_i}} = 0 \tag{A.29}$$

Rearranging these and dividing through by either $\prod_j B_{ij}$ or $\prod_i B_{ij}$ gives:

$$\mathcal{I} = \frac{\sum_j \left[\frac{A_{p_i j}}{B_{p_i j}} \times \frac{e^{z_j}}{\prod_{k\neq j} B_{p_i k}}\right]}{\sum_j \frac{e^{z_j}}{\prod_{k\neq j} B_{p_i k}}} = \frac{\sum_j r_{p_i j} Z_j}{\sum_j Z_j} \tag{A.30}$$

$$\mathcal{I} = \frac{\sum_j \left[\frac{A_{j q_i}}{B_{j q_i}} \times \frac{e^{v_j}}{\prod_{k\neq j} B_{k q_i}}\right]}{\sum_j \frac{e^{v_j}}{\prod_{k\neq j} B_{k q_i}}} = \frac{\sum_j r_{j q_i} V_j}{\sum_j V_j} \tag{A.31}$$

Where $r_{ij} = A_{ij}/B_{ij}$ is the pointwise yield ratio, and V_j and Z_j are shorthand for the terms appearing in the denominator: we only need to know that $V_j, Z_j > 0$. We now consider the general form of the above two equations:

$$\mathcal{T}_i = \frac{\sum_j r_i^j w_j}{\sum_j w_j} \tag{A.32}$$

Here r_i^j is either r_{ij} or r_{ji}, as needed, and $w_j > 0$ is an arbitrary vector. We now consider renumbering our indicies r_i^j such that $r_i^1 \geq r_i^2 \geq \ldots \geq r_i^N$, with $N = m_\beta$ or m_Z depending on the choice made for r_i^j.

Consider an arbitrary choice $\mathbf{w} = (w_1, ..., w_N)$. Subtracting r_i^1 from the resulting value of $\mathcal{T}_i$ gives:

$$\begin{aligned} \mathcal{T}_i - r_i^1 &= \frac{\sum_j w_j r_i^j}{\sum_j w_j} - \frac{\sum_j w_j r_j^1}{\sum_j w_j} \\ &= \frac{\sum_j w_j (r_i^j - r_i^1)}{\sum_j w_j} \end{aligned} \tag{A.33}$$

Since $r_i^1 \geq r_i^j$, the numerator is always zero or negative. Since $w_j \geq 0$, it must follow that $\mathcal{T}_i \leq r_i^1$: the maximum possible value of this function is r_i^1, which we defined as the maximum element of the pointwise-division of A and B. If we repeat this exercise, reordering such that $r_i^1 \leq r_i^j$, we would come to the opposite conclusion: $\mathcal{T}_i$ is bounded from below by the minimum value of r_{ij}.

If solutions with multiple non-zero elements of $e^{\mathbf{x}}$ and/or $e^{\mathbf{y}}$ exist, we have therefore shown that they produce extremal ratio values which are bounded by those one finds if you assume there to be only a single non-zero element. To find the value of the global maxima/minima (which is our aim here), it therefore suffices only to consider the single-element solutions.

Since the matrices are of finite size and perfectly known, it is a trivial exercise to list all the relevant extremal points and find the maximum and minimum values, which must therefore correspond to the global maxima and mimimum values.

$$\min\left(y_A \oslash y_B\right) \leq \frac{\sum_{i,j} \epsilon_i \mu_j y_{A,i,j}}{\sum_{i,j} \epsilon_i \mu_j y_{B,i,j}} \leq \max\left(y_A \oslash y_B\right) \tag{A.34}$$

Where $(A \oslash B)_{ij} = A_{ij}/B_{ij}$ is the pointwise division of the matrices and $\min(A)$ and $\max(A)$ are respectively the minimum and maximum values of the matrix elements.

B
Spin-Priors

In §6.5.1, we elected to choose a uniform prior on σ, despite the fact that we then proceeded to justify eliminating these models because of a relatively strong prior on σ, and correspondingly $\{\mu_i\}$. We now justify this decision.

Figure B.1 shows the effect of adding a Gaussian prior with $\mu = 0.5$ and varying standard deviations. We see that as the prior increased in strength (i.e. δ_σ decreased), there was a discontinuous jump in the location of the peak value of $\log(\mathcal{L})$, and that prior to this jump, there was little meaningful alteration in the peak value of σ. Figure B.2 shows that this was not the optimiser picking up on a as-yet-unseen local maximum in the distribution: the predicted curves for the optimal fit to $(\mu, \delta_\sigma) = (0.5, 0.1)$ are nor more incompatible with the observational evidence than those we found before: the [O/Fe] line is consistent up to high rotation speeds, but [Mg/Fe] is worse, and both [Si/Fe] and [Ca/Fe] remain too high for $\sigma > 0.04$.

It is clear that the prior has very little effect on the outcome, until it suddenly becomes the sole, dominant factor. As the position at which this happens was difficult to predict, we elected to use a uniform prior, as this made a numerically insignificant alteration to our outcome, and avoids the danger of the optimiser suddenly electing to prioritise the σ prior over the data.

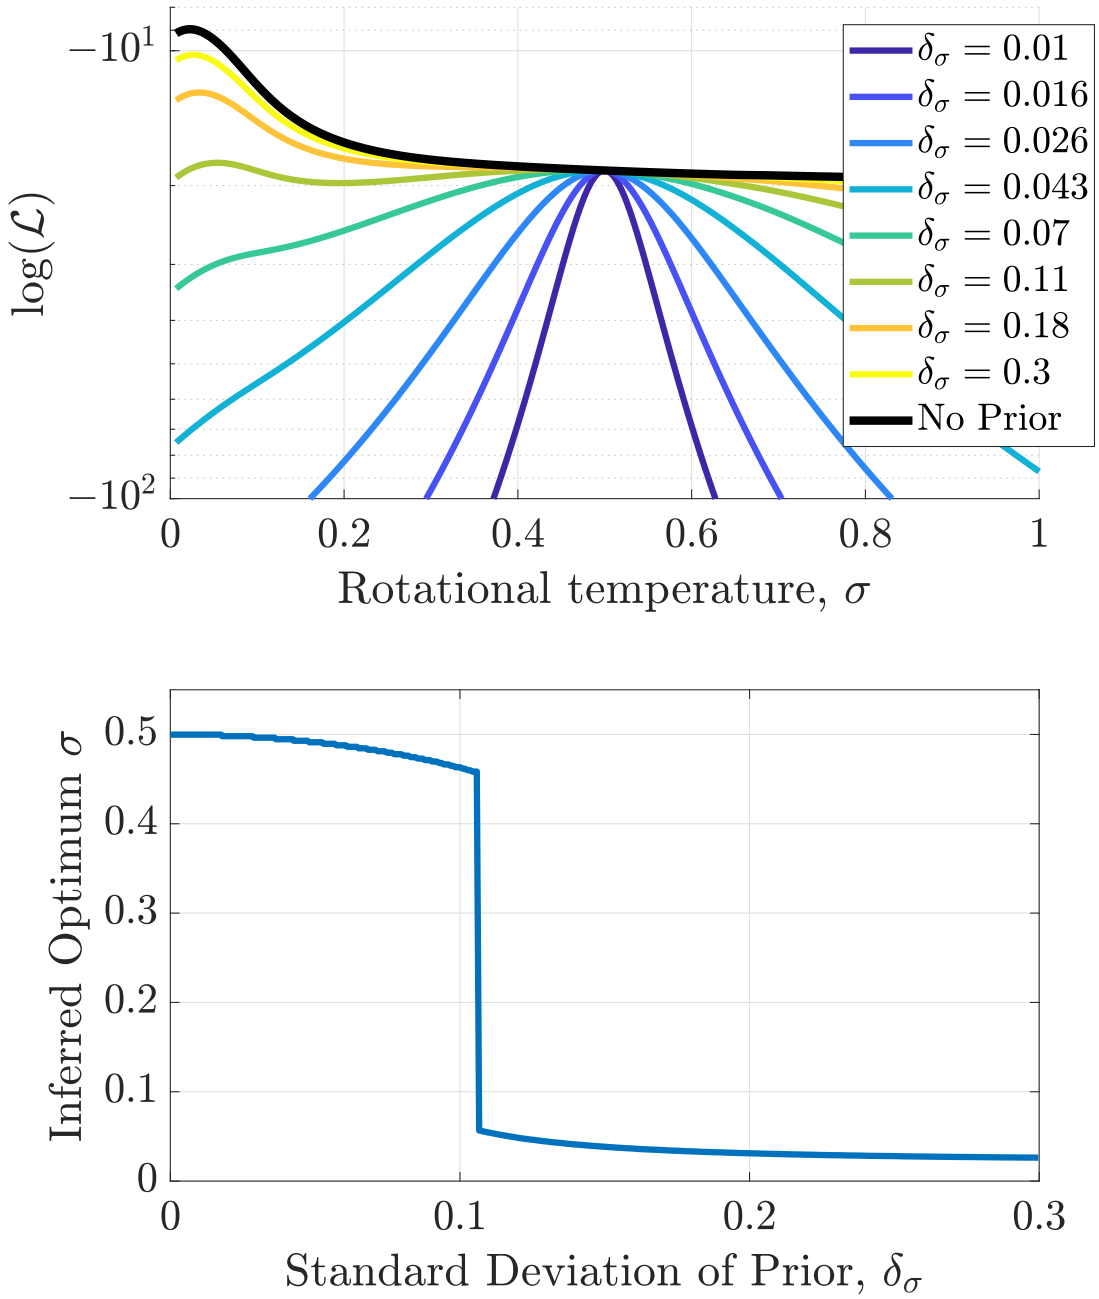

Figure B.1: The effect of adding a Gaussian prior with mean $\mu = 0.5$ to the Likelihood function and varying the standard deviation. All other parameters were fixed at those which maximised the Likelihood for $\delta_\sigma = 0.1$.

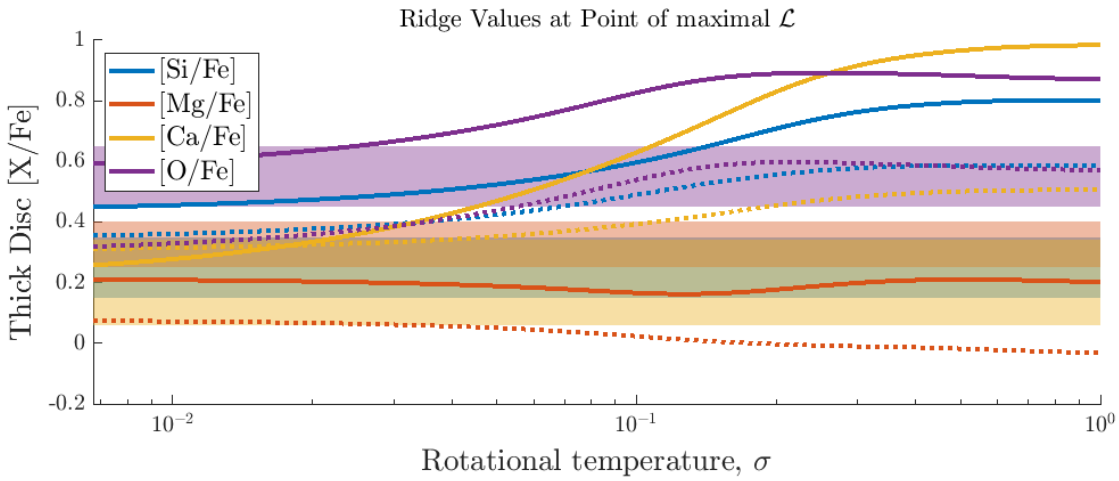

Figure B.2: Solid lines: the same as in Figure 6.6, **Dashed lines:** The effect the Gaussian prior with $(\mu, \delta_\sigma) = (0.5, 0.1)$ on the inferred values of the Ridgeline. All other parameters were fixed at those which maximised the Likelihood for $\delta_\sigma = 0.1$. Adding the prior shows a marked degredation in accuracy.

Bibliography

Abbott B. P., et al., 2017, Phys. Rev. Lett., 119, 161101

Adams S. M., Kochanek C. S., Gerke J. R., Stanek K. Z., Dai X., 2017, MNRAS, 468, 4968

Ahumada R., et al., 2020, APJS, 249, 3

Allen D. M., Ryan S. G., Rossi S., Beers T. C., Tsangarides S. A., 2012, A&A, 548, A34

Alpher R. A., Bethe H., Gamow G., 1948, Phys. Rev., 73, 803

Argast D., Samland M., Thielemann F. K., Qian Y. Z., 2004, A&A, 416, 997

Arnett W. D., Clayton D. D., 1970, Nature, 227, 780

Asplund M., Grevesse N., Sauval A. J., Scott P., 2009, Annual Review of Astronomy and Astrophysics, 47, 481

Asplund M., Amarsi A. M., Grevesse N., 2021, AAP, 653, A141

Aumer M., Binney J. J., 2009, Monthly Notices of the Royal Astronomical Society, 397, 1286

Bailey J. E., et al., 2015, Nature, 517, 56

Battistini C., Bensby T., 2016, A&A, 586, A49

Bensby T., Feltzing S., Oey M. S., 2014, A&A, 562, A71

Bergemann M., Collet R., Schönrich R., Andrae R., Kovalev M., Ruchti G., Hansen C. J., Magic Z., 2017, The Astrophysical Journal, 847, 16

Biemont E., Karner C., Meyer G., Traeger F., Zu Putlitz G., 1982, A&A, 107, 166

Bilitewski T., Schönrich R., 2012, Monthly Notices of the Royal Astronomical Society, 426, 2266

Binney J., Tremaine S., 2008, Galactic Dynamics: Second Edition

Binney J., Dehnen W., Bertelli G., 2000, Monthly Notices of the Royal Astronomical Society, 318, 658

Bloom J. S., Kulkarni S. R., Djorgovski S. G., 2002, AJ, 123, 1111

Bodenheimer P., Woosley S. E., 1983, ApJ, 269, 281

Bovy J., Rix H.-W., Hogg D. W., 2012, The Astrophysical Journal, 751, 131

Bressan A., Marigo P., Girardi L., Salasnich B., Dal Cero C., Rubele S., Nanni A., 2012, Monthly Notices of the Royal Astronomical Society, 427, 127

Buder S., et al., 2021, Monthly Notices of the Royal Astronomical Society, 506, 150

Burbidge E. M., Burbidge G. R., Fowler W. A., Hoyle F., 1957, Reviews of Modern Physics, 29, 547

Caffau E., Ludwig H. G., Steffen M., Freytag B., Bonifacio P., 2011, Solar Physics, 268, 255

Casagrande L., Schönrich R., Asplund M., Cassisi S., Ramírez I., Meléndez J., Bensby T., Feltzing S., 2011, A&A, 530, A138

Chabrier G., 2003, PASP, 115, 763

Chen Y., Bressan A., Girardi L., Marigo P., Kong X., Lanza A., 2015, Monthly Notices of the Royal Astronomical Society, 452, 1068

Chevalier R. A., Kirshner R. P., 1979, ApJ, 233, 154

Chiappini C., Matteucci F., Gratton R., 1997, The Astrophysical Journal, 477, 765

Chieffi A., Limongi M., 2004, ApJ, 608, 405

Chiosi C., 1980, A&A, 83, 206

Chomiuk L., Povich M. S., 2011, The Astronomical Journal, 142, 197

Chornock R., et al., 2017, The Astrophysical Journal, 848, L19

Clayton D., Fowler W., Hull T., Zimmerman B., 1961, Annals of Physics, 12, 331

Côté B., Denissenkov P., Herwig F., Ruiter A. J., Ritter C., Pignatari M., Belczynski K., 2018, The Astrophysical Journal, 854, 105

Cowan J. J., Rose W. K., 1977, ApJ, 212, 149

Demircan O., Kahraman G., 1991, Ap&SS, 181, 313

Denissenkov P. A., Herwig F., Battino U., Ritter C., Pignatari M., Jones S., Paxton B., 2017, The Astrophysical Journal, 834, L10

Ekström S., et al., 2012, AAP, 537, A146

El-Badry K., Ostriker E. C., Kim C.-G., Quataert E., Weisz D. R., 2019, MNRAS, 490, 1961

Fielding D. B., Ostriker E. C., Bryan G. L., Jermyn A. S., 2020, The Astrophysical Journal, 894, L24

Fischer T., Martínez-Pinedo G., Hempel M., Liebendörfer M., 2012, Phys. Rev.˜D, 85, 083003

Fischer T., Wu M.-R., Wehmeyer B., Bastian N.-U. F., Martínez-Pinedo G., Thielemann F.-K., 2020, ApJ, 894, 9

Foreman-Mackey D., 2016, The Journal of Open Source Software, 1, 24

Foreman-Mackey D., Hogg D. W., Lang D., Goodman J., 2013, Publications of the Astronomical Society of the Pacific, 125, 306

Fraser J., Schönrich R., 2022, MNRAS, 509, 6008

Freiburghaus C., Rosswog S., Thielemann F.-K., 1999, ApJl, 525, L121

Fujimoto S.-i., Kotake K., Yamada S., Hashimoto M.-a., Sato K., 2006, ApJ, 644, 1040

Fulbright J. P., 2000, The Astronomical Journal, 120, 1841

Fynbo J. P. U., et al., 2006, Nature, 444, 1047

Gao Y., Solomon P. M., 2004, ApJ, 606, 271

Gilmore G., Reid N., 1983, Monthly Notices of the Royal Astronomical Society, 202, 1025

Grevesse N., Sauval A., 1998, Space Science Reviews, 85, 161

Grevesse N., Asplund M., Sauval A. J., 2007, Space Science Reviews, 130, 105

Grisoni V., Spitoni E., Matteucci F., 2018, MNRAS, 481, 2570

Gronow S., Collins C. E., Sim S. A., Röpke F. K., 2021, AAP, 649, A155

Guerço R., Cunha K., Smith V. V., Hayes C. R., Abia C., Lambert D. L., Jönsson H., Ryde N., 2019, The Astrophysical Journal, 885, 139

Guiglion G., de Laverny P., Recio-Blanco A., Prantzos N., 2018, A&A, 619, A143

Hayes C. R., et al., 2018, The Astrophysical Journal, 852, 49

Haynes C. J., Kobayashi C., 2018, Monthly Notices of the Royal Astronomical Society, 483, 5123

Haywood M., 2008, Monthly Notices of the Royal Astronomical Society, 388, 1175

Healy B. F., McCullough P. R., Schlaufman K. C., 2021, arXiv e-prints, p. arXiv:2109.08692

Heger A., Langer N., Woosley S. E., 2000, The Astrophysical Journal, 528, 368

Heiderman A., Evans Neal J. I., Allen L. E., Huard T., Heyer M., 2010, ApJ, 723, 1019

Honda S., Aoki W., Ishimaru Y., Wanajo S., Ryan S. G., 2006, ApJ, 643, 1180

Hoyle F., 1946, MNRAS, 106, 343

Hoyle F., 1954, ApJs, 1, 121

Hoyle F., Fowler W. A., 1960, ApJ, 132, 565

Hwang U., Laming J. M., 2012, ApJ, 746, 130

ISO 2020, Technical report, International Standard ISO/IEC 14882:2020(E) – Programming Language C++., https://isocpp.org/std/the-standard. International Organization for Standardization, Geneva, CH, https://isocpp.org/std/the-standard

Iwamoto K., Brachwitz F., Nomoto K., Kishimoto N., Umeda H., Hix W. R., Thielemann F.-K., 1999, The Astrophysical Journal Supplement Series, 125, 439

Jones S., et al., 2019, AAP, 622, A74

Just A., Jahreiß H., 2007, arXiv preprint arXiv:0706.3850

Karakas A. I., Lugaro M., 2016, ApJ, 825, 26

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

Kennicutt R. C., Evans N. J., 2012, ARAA, 50, 531

Khoperskov S., Haywood M., Snaith O., Di Matteo P., Lehnert M., Vasiliev E., Naroenkov S., Berczik P., 2021, MNRAS, 501, 5176

Kobayashi C., Karakas A. I., Lugaro M., 2020, The Astrophysical Journal, 900, 179

Kraft R. P., 1967, The Astrophysical Journal, 150, 551

Krumholz M. R., McKee C. F., Tumlinson J., 2009, The Astrophysical Journal, 699, 850

Kulkarni S. R., et al., 1998, Nature, 395, 663

Langer N., Norman C. A., 2006, ApJl, 638, L63

Langer N., Heger A., Wellstein S., Herwig F., 1999, A&A, 346, L37

Lattimer J. M., Schramm D. N., 1974, ApJ, 192, L145

Leitherer C., Ekström S., Meynet G., Schaerer D., Agienko K. B., Levesque E. M., 2014, The Astrophysical Journal Supplement Series, 212, 14

Limongi M., Chieffi A., 2008, ORFEO database, http://orfeo.iaps.inaf.it/

Limongi M., Chieffi A., 2018, A&As, 237, 13

Loebman S. R., Debattista V. P., Ivezić Ž., Quinn T. R., Wadsley J., et al., 2011, The Astrophysical Journal, 737, 8

Luck R. E., Lambert D. L., 2011, AJ, 142, 136

MacFadyen A. I., Woosley S. E., 1999, The Astrophysical Journal, 524, 262

Maeder A., 1992, A&A, 264, 105

Magg E., et al., 2022, arXiv e-prints, p. arXiv:2203.02255

Marigo P., 2001, A&A, 370, 194

Marigo P., et al., 2017, The Astrophysical Journal, 835, 77

Matt S. P., Brun A. S., Baraffe I., Bouvier J., Chabrier G., 2015, ApJL, 799, L23

Matteucci F., Romano D., Arcones A., Korobkin O., Rosswog S., 2014, MNRAS, 438, 2177

Mazzali P. A., et al., 2003, ApJl, 599, L95

McCarthy I. G., Font A. S., Crain R. A., Deason A. J., Schaye J., Theuns T., 2012, MNRAS, 420, 2245

McKee C. F., Parravano A., Hollenbach D. J., 2015, The Astrophysical Journal, 814, 13

Meynet G., Maeder A., 2000, arXiv preprint astro-ph/0006404

Mishenina T. V., Pignatari M., Korotin S. A., Soubiran C., Charbonnel C., Thielemann F. K., Gorbaneva T. I., Basak N. Y., 2013, A&A, 552, A128

Miyaji S., Nomoto K., Yokoi K., Sugimoto D., 1980, PASJ, 32, 303

Mor R., Robin A. C., Figueras F., 2019, in Montesinos B., Asensio Ramos A., Buitrago F., Schödel R., Villaver E., Pérez-Hoyos S., Ordóñez-Etxeberria I., eds, Highlights on Spanish Astrophysics X. pp 424–424

Nishimura N., Sawai H., Takiwaki T., Yamada S., Thielemann F. K., 2017, ApJl, 836, L21

Nordström B., et al., 2004, AAP, 418, 989

Norman C. A., Ikeuchi S., 1989, ApJ, 345, 372

Ojima T., Ishimaru Y., Wanajo S., Prantzos N., François P., 2018, The Astrophysical Journal, 865, 87

Palla M., Matteucci F., Spitoni E., Vincenzo F., Grisoni V., 2020, MNRAS, 498, 1710

Pannarale F., Ohme F., 2014, ApJl, 791, L7

Parker E. N., 1958, Reviews of Modern Physics, 30, 955

Pasquini L., et al., 2002, The Messenger, 110, 1

Pastorelli G., et al., 2020, Monthly Notices of the Royal Astronomical Society, 498, 3283

Perley D. A., et al., 2016, The Astrophysical Journal, 817, 8

Pettini M., Madau P., Bolte M., Prochaska J. X., Ellison S. L., Fan X., 2003, ApJ, 594, 695

Podsiadlowski P., Ivanova N., Justham S., Rappaport S., 2010, MNRAS, 406, 840

Prantzos N., Abia C., Limongi M., Chieffi A., Cristallo S., 2018, MNRAS, 476, 3432

Ramírez-Agudelo O. H., et al., 2013, AAP, 560, A29

Reddy B. E., Lambert D. L., Prieto C. A., 2006, Monthly Notices of the Royal Astronomical Society, 367, 1329

Reeves H., Fowler W., Hoyle F., 1970, Nature, 226, 727

Robitaille T. P., Whitney B. A., 2010, ApJl, 710, L11

Rosswog S., Liebendörfer M., Thielemann F. K., Davies M. B., Benz W., Piran T., 1999, A&A, 341, 499

Rosswog S., Sollerman J., Feindt U., Goobar A., Korobkin O., Wollaeger R., Fremling C., Kasliwal M. M., 2018, A&A, 615, A132

Sakari C. M., et al., 2018, ApJ, 868, 110

Salpeter E. E., 1955, ApJ, 121, 161

Schmidt M., 1959, The Astrophysical Journal, 129, 243

Schneider F. R. N., et al., 2018, A&A, 618, A73

Schönrich R., Binney J., 2009a, MNRAS, 396, 203

Schönrich R., Binney J., 2009b, MNRAS, 399, 1145

Schönrich R., McMillan P. J., 2017, Monthly Notices of the Royal Astronomical Society, 467, 1154

Schönrich R. A., Weinberg D. H., 2019, MNRAS, 487, 580

Shapiro P. R., Field G. B., 1976, The Astrophysical Journal, 205, 762

Sharma S., et al., 2018, Monthly Notices of the Royal Astronomical Society, 473, 2004

Siegel D. M., Barnes J., Metzger B. D., 2019, Nature, 569, 241

Sneden C., McWilliam A., Preston G. W., Cowan J. J., Burris D. L., Armosky B. J., 1996, ApJ, 467, 819

Sneden C., Cowan J. J., Gallino R., 2008, ARA&A, 46, 241

Sollerman J., et al., 2006, A&A, 454, 503

Spitoni E., Verma K., Silva Aguirre V., Calura F., 2020, AAP, 635, A58

Spitoni E., et al., 2021, AAP, 647, A73

Suess H. E., Urey H. C., 1956, Rev. Mod. Phys., 28, 53

Sweet P., 1950, Monthly Notices of the Royal Astronomical Society, 110, 548

Symbalisty E., Schramm D. N., 1982, Astrophys. Lett., 22, 143

Tang J., Bressan A., Rosenfield P., Slemer A., Marigo P., Girardi L., Bianchi L., 2014, Monthly Notices of the Royal Astronomical Society, 445, 4287

Tinsley B. M., 1979, ApJ, 229, 1046

Travaglio C., Hillebrandt W., Reinecke M., 2005, Astronomy & Astrophysics, 443, 1007

Tsujimoto T., Baba J., 2019, The Astrophysical Journal, 878, 125

Umeda H., Nomoto K., 2002, ApJ, 565, 385

Vartanyan D., Burrows A., Radice D., Skinner M. A., Dolence J., 2018, Monthly Notices of the Royal Astronomical Society, 482, 351

Villante F. L., Serenelli A., 2020, arXiv e-prints, p. arXiv:2004.06365

Wanajo S., Ishimaru Y., 2006, Nuclear Phys.˜A, 777, 676

Wanajo S., Kajino T., Mathews G. J., Otsuki K., 2001, ApJ, 554, 578

Wanajo S., Janka H.-T., Müller B., 2011, ApJl, 726, L15

Weinberg D. H., Andrews B. H., Freudenburg J., 2017, The Astrophysical Journal, 837, 183

Wheeler J. C., Hansen C. J., 1971, APSS, 11, 373

Whelan J., Iben Icko J., 1973, ApJ, 186, 1007

Wolf C., Podsiadlowski P., 2007, Monthly Notices of the Royal Astronomical Society, 375, 1049

Woosley S. E., 1993, in American Astronomical Society Meeting Abstracts #182. p. 55.05

Woosley S. E., Heger A., 2006, The Astrophysical Journal, 637, 914

Woosley S. E., Weaver T. A., 1986, The Physics of Supernovae. p. 91, doi:10.1007/3-540-16764-1_7

Woosley S. E., Weaver T. A., 1995, ApJS, 101, 181

Woosley S. E., Wilson J. R., Mathews G. J., Hoffman R. D., Meyer B. S., 1994, ApJ, 433, 229

Yoon S. C., Langer N., 2005, A&A, 443, 643

Yoon S. C., Langer N., Norman C., 2006, Nuovo Cimento B Serie, 121, 1631

Zhao G., et al., 2016, ApJ, 833, 225

Zorec J., Royer F., 2012, A&A, 537, A120