

UNIVERSITÀ
DI PAVIA

Dipartimento di Fisica

Corso di laurea magistrale in Scienze fisiche

Vincoli sulla struttura termica dei dischi protoplanetari dalla cinematica del gas: un test su osservazioni sintetiche

Constraining the thermal structure of protoplanetary discs with kinematics: a test on synthetic observations

Tesi per la laurea di:
Matteo Raffaelli

Relatore:
Prof. Giuseppe Lodato

Correlatore interno:
Prof. Andrea Giuliani

Anno accademico 2025 – 2026

Abstract

Planets form in protoplanetary discs, structures of gas and dust which surround young stars, and how much mass such a disc contains, how that mass is distributed and how warm its gas is decide what can form in it and where. None of these quantities are observed directly, so they have to be inferred from the light the disc emits. Developing methods to measure such quantities from the emitted light is essential in understanding how planetary systems form and evolve and in answering questions like "Why is our Solar system so different than other planetary systems we have observed through our telescopes?".

The gas in a disc rotates at a speed which is close to Keplerian, and the small departure from Keplerian rotation is what carries the information. It is set by the mass of the disc itself through its self-gravity, and by the radial gradient of the gas pressure, which depends in turn on the thermal structure and density of the disc. Measuring the rotation curve of a molecular line therefore gives the mass of the central star, the mass of the disc and the radius over which the gas is distributed. The thermal structure, which is what sets the pressure gradient together with the density, is the one ingredient which the method is not currently used to measure. It is taken instead from the brightness temperature of optically thick lines, a procedure which does not reach the midplane and is dependent on chemical models.

This thesis asks whether the rotation curve can measure the thermal structure as well, that is whether the thermal parameters of a disc can be extracted from the rotation curves of two molecular tracers which emit at different heights above the midplane. We answer the question on synthetic observations, for which the underlying structure is known exactly.

We first study the model with Fisher's method. The thermal parameters are the hardest to constrain, because they are degenerate with the stellar mass and with each other, but a second tracer emitting from a different height reduces those degeneracies, and when the only source of error is thermal noise on the curve every parameter is recovered within its errors. We then build a realistic pipeline, which passes a synthetic disc through the radiative transfer code MCFOST and extracts its rotation curve with discminer exactly as it would be extracted from an ALMA observation, and we apply it to a sample of forty discs with two tracers each. The extraction carries a systematic error of ~ 2 per cent of the velocity, always of the same sign, smooth in radius and three times the uncertainty quoted on each point. Fisher's analysis converts a change of the thermal parameters into a change of the curve, and against that scale the systematic

error is twice as large. With an extraction of this accuracy the thermal parameters cannot yet be measured.

Sommario

I pianeti si formano nei dischi protoplanetari, strutture di gas e polvere che circondano le stelle giovani, e quanta massa un disco contenga, come quella massa sia distribuita e quanto sia caldo il suo gas decidono che cosa possa formarsi al suo interno e dove. Nessuna di queste grandezze è osservata direttamente, quindi vanno dedotte dalla luce che il disco emette. Sviluppare metodi per misurarle a partire da quella luce è essenziale per capire come i sistemi planetari si formino ed evolvano, e per rispondere a domande come «perché il nostro Sistema solare è così diverso dagli altri sistemi planetari che abbiamo osservato con i nostri telescopi?».

Il gas di un disco ruota a una velocità vicina a quella kepleriana, e il piccolo scostamento dalla rotazione kepleriana è ciò che porta l'informazione. È fissato dalla massa del disco stesso, attraverso la sua autogravità, e dal gradiente radiale della pressione del gas, che dipende a sua volta dalla struttura termica e dalla densità del disco. Misurare la curva di rotazione di una riga molecolare fornisce quindi la massa della stella centrale, la massa del disco e il raggio su cui il gas è distribuito. La struttura termica, che insieme alla densità è ciò che fissa il gradiente di pressione, è l'unico ingrediente che il metodo non viene attualmente usato per misurare. Viene invece ricavata dalla temperatura di brillanza di righe otticamente spesse, una procedura che non arriva al piano mediano e che dipende da modelli chimici.

Questa tesi si chiede se la curva di rotazione possa misurare anche la struttura termica, cioè se i parametri termici di un disco possano essere estratti dalle curve di rotazione di due traccianti molecolari che emettono ad altezze diverse sopra il piano mediano. Rispondiamo alla domanda su osservazioni sintetiche, per le quali la struttura sottostante è nota esattamente.

Studiamo prima il modello con il metodo di Fisher. I parametri termici sono i più difficili da vincolare, perché sono degeneri con la massa stellare e fra loro, ma un secondo tracciante che emette da un'altezza diversa riduce quelle degenerazioni, e quando l'unica fonte di errore è il rumore termico sulla curva ogni parametro viene recuperato entro i suoi errori. Costruiamo poi una pipeline realistica, che fa passare un disco sintetico attraverso il codice di trasporto radiativo MCFOST e ne estrae la curva di rotazione con discminer esattamente come la si estrarrebbe da un'osservazione ALMA, e la applichiamo a un campione di quaranta dischi con due traccianti ciascuno. L'estrazione porta con sé un errore sistematico pari a circa il 2 per cento della velocità, sempre dello stesso segno, regolare in raggio e tre volte l'incertezza dichiarata su ogni punto. L'analisi di Fisher converte una variazione dei parametri termici in una

variazione della curva, e su quella scala l'errore sistematico è grande il doppio. Con un'estrazione di questa accuratezza i parametri termici non sono ancora misurabili.

Contents

Abstract **iii**

Sommario **v**

1 Introduction **1**
1.1 Context . . . 3
1.2 State of the art . . . 6
1.3 Aims and outline . . . 8

2 The theory **11**
2.1 Gas dynamics in protoplanetary discs . . . 11
2.1.1 Disc energetics . . . 11
2.1.2 Assumptions . . . 12
2.1.3 The fundamental equations . . . 13
2.1.4 The viscous stress tensor . . . 14
2.1.5 Azimuthal component . . . 16
2.1.5.1 Solving the Diffusion Equation . . . 18
2.1.6 Vertical component . . . 21
2.1.7 Radial component . . . 22
2.2 The model . . . 23
2.2.1 Emission surfaces . . . 27
2.2.2 The model in summary . . . 28

3 Testing pipeline **31**
3.1 MCFOST . . . 31
3.1.1 The grid . . . 32
3.1.2 The synthetic disc . . . 32
3.1.3 Radiative transfer and the cube . . . 34
3.2 Discminer . . . 36
3.2.1 Brightness temperature . . . 39
3.2.2 Rotation curve extraction . . . 43
3.2.3 Beam smearing correction . . . 43

4 Parameter space **47**
4.1 Fisher's method . . . 47
4.2 Isothermal case . . . 49

4.2.1 Degeneracies 50
4.2.2 Number of tracers 52
4.2.3 Radial range 54
4.3 Stratified case 55
4.3.1 Degeneracies 56
4.3.2 Number of tracers 58
4.3.3 Radial range 58
4.4 A simple test 58

5 Results **67**
5.1 The extraction bias 67
5.2 Retrieval of the parameters 69
5.2.1 Without any correction 69
5.2.2 With the beam smearing correction 71
5.2.3 What the rotation curve measures 72

6 Conclusions **75**
6.1 Summary 75
6.2 Future work 76

A SPH simulation **79**
A.1 What is an SPH simulation 79
A.2 The code 79
A.3 The time loop 81
A.4 The spreading of a ring 82
A.5 The self-similar disc 82
A.6 Known errors and how we handle them 84

Bibliography **87**

Acknowledgements **91**

Temp'era dal principio del mattino,
e 'l sol montava 'n sù con quelle stelle
ch'eran con lui quando l'amor divino
mosse di prima quelle cose belle;

Dante Alighieri, *Inferno*, I, 37–40

Chapter 1

Introduction

Astronomical objects have always fuelled our creativity and guided us throughout our history. However, only after Galileo have we really started to pull back the covers of the great machine that is our universe. Ever since 2011 the ALMA interferometer, a descendant of the first telescope used by the father of modern science, has allowed us to zoom in on new worlds, giving us unprecedented insights on their birth. Figure 1.1 shows an image of the not so small[1] cradle of planets just like the ones in our own system. These giant cradles are aptly called protoplanetary discs and are structures made up of dust and gas found orbiting a central star in a fluid-like manner.

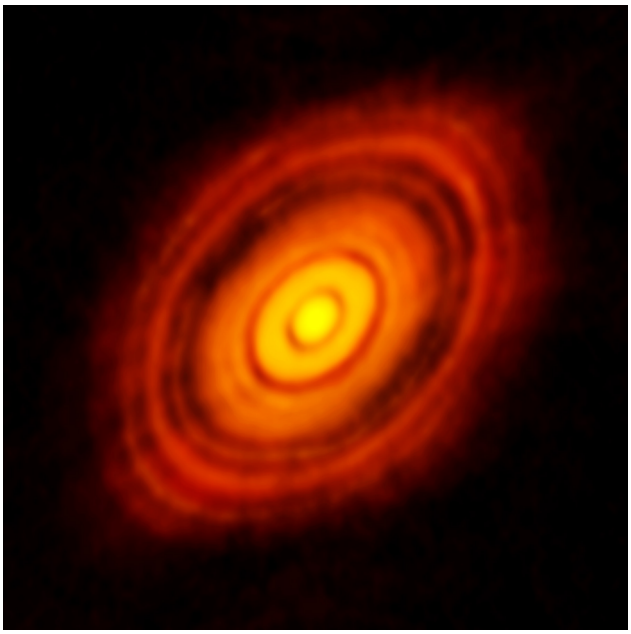


Figure 1.1: ALMA image of the protoplanetary disc around HL Tauri. The ring-shaped dark gaps are likely the result of the gravitational influence of newly formed planets. Image credit: ALMA (ESO/NAOJ/NRAO) [1].

[1] Of the order of hundreds of au (astronomical unit). An astronomical unit is a unit of length equal to the average distance from the Earth to the Sun, measuring precisely 149,597,870,700 metres.

The science behind these structures will be thoroughly developed in the next chapters, therefore I will use this introduction to explain why protoplanetary discs are so interesting and which mysteries could be solved by studying them.

As of August 2026, 6333 exoplanets have been discovered [2], and their distribution in mass and semi-major axis, seen in Figure 1.2, is full of things we cannot explain. Why are there planets as massive as Jupiter orbiting their star in less than ten days[2], when we believe that giant planets can only form at greater orbits? Why is the most common kind of planet around Sun-like stars, bigger than the Earth and smaller than Neptune [3][3], the one our own system does not have? And are Earth-like planets really rare, or are our instruments simply blind to a small planet on a wide orbit? These and many other questions still need to be answered not only because of curiosity but also because of their deep connection with the problem of the origin of life. It is widely known that this is a chemistry problem, however chemical reactions are regulated by temperature, the presence of an atmosphere and its composition, gravity and many other physical factors that are direct consequences of the nature of a planet. To solve the origin of life problem it is therefore essential to understand how planets form and we can do that by looking at them being formed[4].

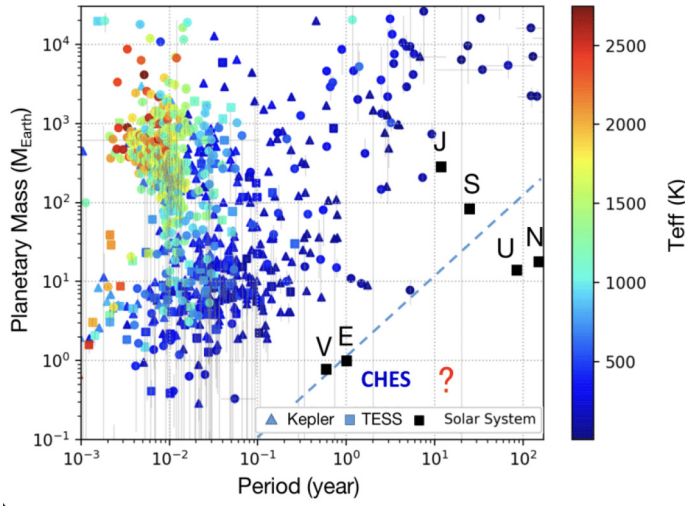


Figure 1.2: Mass and orbital period distribution of exoplanets. Reproduced from Ji et al. [4], Fig. 1.

Another, more personal, reason for studying protoplanetary discs is the elegant way in which they can be described. These astronomical objects are basically giant spinning discs of fluid that follow two very famous equations: the conservation of mass and Navier-Stokes. The model tested in this thesis will be derived directly from these equations with elementary but long calculations, and its final form will seem almost obvious in its clarity. During the calculations other famous mathematical objects will come up like the diffusion equation or the Gaussian distribution. I am deeply fas-

[2]These planets are called Hot-Jupiters.
[3]These planets are called super-Earths or sub-Neptunes
[4]No pun intended.

cinated by the way totally unrelated phenomena are described so well by the same mathematical objects. I also find that this work perfectly sums up what I have learned during my studies in theoretical physics.

1.1 Context

This thesis aims to make its contribution to the field of dynamical measurements of protoplanetary discs, an area of astrophysics which concerns itself with obtaining important information by looking at the way the gaseous fluid orbits the central star at different distances. The reasoning behind it can be rapidly summed up.

1. Starting by essential equations and reasonable assumptions it is possible to derive a model for the azimuthal velocity[5] at different radii and heights which will depend on physical features like the mass of the star, the mass of the orbiting fluid and the temperature.
2. By comparing the model's predicted rotation curve to the one measured on real data it is possible to estimate the value and uncertainty of the physical features.

The real data used for this comparison comes from the sub-millimetre, the spectral region where the rotational transitions of the gas fall, and which can be explored at high resolution by the Atacama Large Millimeter Array, ALMA for short, the largest radio telescope in the world. Its outstanding zoom capabilities are due to a series of factors. Firstly, ALMA was built in the Atacama desert, in Chile, one of the driest places on Earth which means no clouds. The Atacama desert is also the highest in the world, meaning less atmosphere can interfere with the observations. The clarity of the night sky in Atacama can be seen in Figure 1.3. ALMA also takes advantage of the great spaces of the desert by using interferometry techniques to combine the signal of a wide array of antennas which is equivalent to using a single antenna with a 15 km dish [5]. The results of this combination of factors are spectacular images like the one in Figure 1.1, obtained by observing the dust continuum thermal emission [1]. On the other hand, observing gas line emission produces the images on which this work is based, like Figure 1.4.

A typical dynamical measurement is based on observations of the disc in a specific spectral window centred on the rest frequency of a transition of a certain isotopologue[6] in order to separate its emission from the one of other sources. This is done because each isotopologue has its own chemically determined abundance and optical depth, which means that its emission starts at a specific height in the disc and therefore probes a different region, as explained in section 2.2.1. The reason why these observations are not taken at the exact transition frequency but rather over a spectral window is the doppler effect. Like every other object in the universe the disc has a systemic velocity with respect to Earth, which is the same everywhere and which we subtract, and on top of it the rotation of the gas shifts the frequency of the emission by an amount

[5] The graph of the velocity of the orbiting fluid against the radii will be called from now on the "rotation curve".

[6] A molecule that has the same chemical formula and structure of another, but differs in its isotopic composition because at least one atom contains a different number of neutrons.

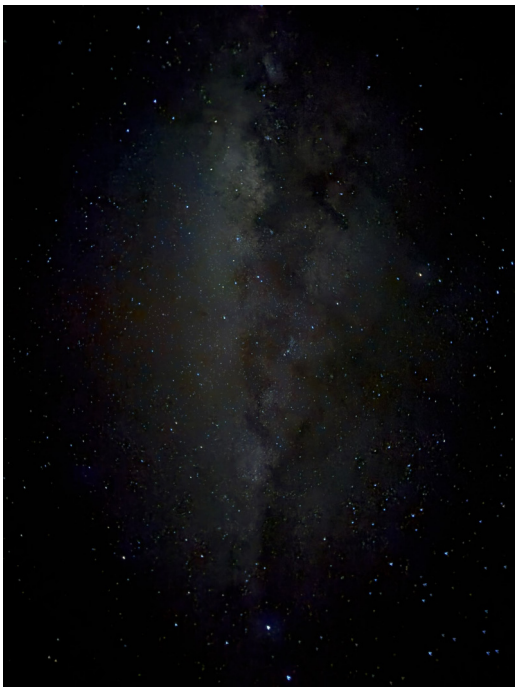


Figure 1.3: Photo of the Atacama night sky, taken with a simple mobile phone.

which changes from point to point, since what enters is the projection of the velocity along our line of sight. The window has to contain all of these shifts, from the fast gas of the inner regions to the slow gas of the outer ones, and this is what makes it several kilometres per second wide. The line emitted by a single parcel of gas is not infinitely narrow either, because the molecules move with their thermal agitation and with the turbulent motions which we do not resolve, and inside the beam of the telescope there are parcels which rotate at slightly different speeds. By dividing the spectral window into smaller intervals called channels, we get a series of images, each showing the region of the disc whose gas moves along our line of sight at the velocity of that channel, a band rather than a line because of the widths we have just described. From these regions we can then extract the rotation curve, with the procedure described in Section 3.2.2. The entire observational data is called a data cube because it has three dimensions: two spatial and one spectral, as shown in Figure 1.4. Read along the spectral direction at a fixed position on the sky, the same cube gives the spectrum of that point, the intensity as a function of velocity, and the shape it takes is the line profile, whose centre is set by the projected velocity of the gas there and whose width by the motions we have just described.

Generally the work in the dynamical measurements field is then twofold: a theoretical part followed by an experimental one based on observations to test the retrieval capability of the method.

Other methods used in the literature to measure the mass of protoplanetary discs are based on the dust continuum thermal emission, on chemical gas tracers and on gravitational instabilities. The first two, the continuum and the chemical tracers, are chemical in nature and measure the physical features of the disc indirectly, since they infer the amount of gas from the emission of something else and therefore depend on

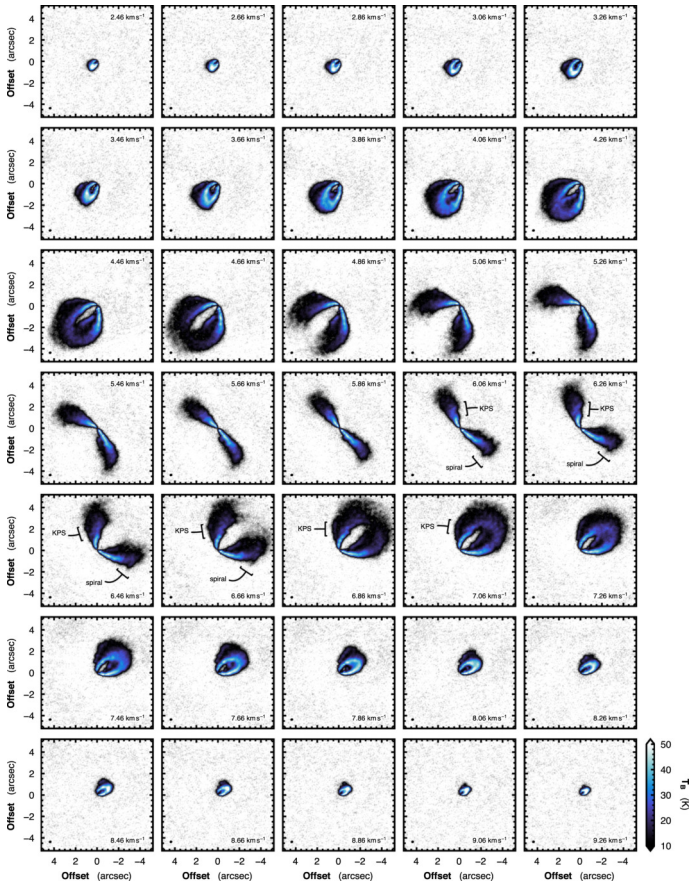


Figure 1.4: The slices of a datacube, showing the isovelocity regions of a disc. The "butterfly" pattern is the result of the fact that in a rotating disc, two different regions have the same velocity along our line of sight. These two regions are symmetric with respect to the major axis of the ellipse resulting from the inclination of the disc in the plane of the sky. This symmetry is exact only for a geometrically thin disc. When the emission comes from an elevated surface the axis of symmetry is displaced by $z \sin i$, where z is the height of the emitting layer and i is the inclination, and measuring this displacement allows us to reconstruct the emission surface [6]. Channel maps reproduced from Teague et al. [7].

uncertain conversion factors [8]. The third, which relies on gravitational instabilities, is indirect as well because it weighs the disc through the spiral structure that develops when the disc is self-gravitating, and it does not give us information on features like the temperature. The dynamical method is a direct measurement that bypasses chemical assumptions [9] and could, in theory, allow us to obtain additional information like the temperature [10]. Its main drawback, which is also encountered in this work, is that the extraction of the rotation curve needs to reach a level of precision which is not commonly found in astrophysics, a difficulty that we will quantify in Chapter 5.

1.2 State of the art

The dynamical measurements approach developed in this work naturally follows from a series of papers all aimed at measuring the mass of the central star, the mass of the disc and the length of the scale radius[7] by comparing a theoretical rotation curve to real measured data. The model for the curve is derived directly from the conservation of mass and Navier-Stokes equations, as shown in Chapter 2. These calculations are thoroughly developed in the paper by Lodato et al. [11], which refines the method used for the first time by Veronesi et al. [12] on the disc of Elias 2-27, and are based on sound physical assumptions. However, for simplicity, the assumed vertical temperature structure is isothermal and not stratified, contrary to what discs have historically shown. The dynamical method is applied to two sources[8] showing inconsistent results between the fits of the two isotopologues considered[9], specifically a factor of four between the two estimated values for the disc mass for the GM Aur source. The two isotopologues of IM Lup agree on a disc mass of $0.1\,M_{\odot}$, with a systematic uncertainty of a factor of two which comes from the method used to extract the rotation curve. This inconsistency is what prompts the successive paper by Martire et al. [10], where the calculations are reformulated under the more realistic hypothesis of a stratified thermal vertical structure, as shown in Figure 1.5.

The analytical solution is then tested against Smoothed-Particle-Hydrodynamics (SPH) numerical simulations of the disc, of which a simplified example can be seen in Appendix A. The test verifies that the density field and the pressure gradient predicted by the stratified model reproduce the simulated ones while the isothermal model does not. The model is then applied to real observations of the MAPS sample[10], and the results show that adding the stratified temperature significantly improves the fit to the rotation curves, as seen in Figure 1.6, and reconciles most of the inconsistencies between the two isotopologues, with the only exception of the outer part of IM Lup where neither of the two models explains the difference. Part of this improvement could come from the extraction and not from the model, since the rotation curves used by Martire et al. [10] have been extracted again and are different from the ones of Lodato

[7] Intuitively this is the radius of the disc but its role will be made more precise in the next chapter.

[8] GM Aur and IM Lup, two young T Tauri stars surrounded by large and bright discs, both at roughly 160 parsecs (pc) from us. A parsec is the distance at which one astronomical unit subtends an angle of one arcsecond, and it is equal to about 3.086×10^{16} metres, that is roughly 206,265 au.

[9] The ^{12}CO and the ^{13}CO.

[10] Molecules with ALMA at Planet-forming Scales, an ALMA Large Program which observed five nearby protoplanetary discs in about twenty molecular lines at high angular and spectral resolution.

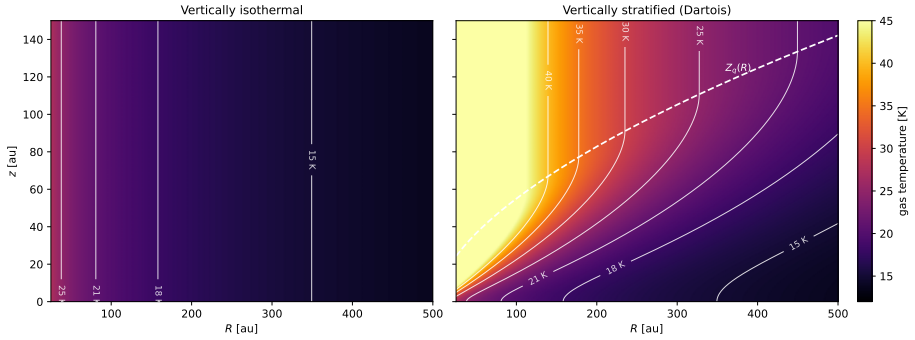


Figure 1.5: Comparison of isothermal and stratified vertical structure. Imagining the disc as a cake, these two graphs are showing the side of a slice. The two profiles are computed with the model of Chapter 2 for the parameters of LkCa 15, and are shown only to illustrate the difference between the two prescriptions.

et al. [11]. However, these are real sources and their true parameters are unknown, therefore the improvement can only be judged from the quality of the fit and from the agreement between the tracers. A comparison with a known truth needs synthetic observations. This was done by Veronesi et al. [13] for isothermal self-gravitating discs and by Andrews et al. [14], who built synthetic observations from models with a stratified temperature and recovered the disc mass with a bias below 20% for a disc to stellar mass ratio of 0.1. In both of these tests the thermal structure is an input of the fit, and at this stage it is still obtained through means different than the dynamical one, in this case by measuring the brightness temperature [15].

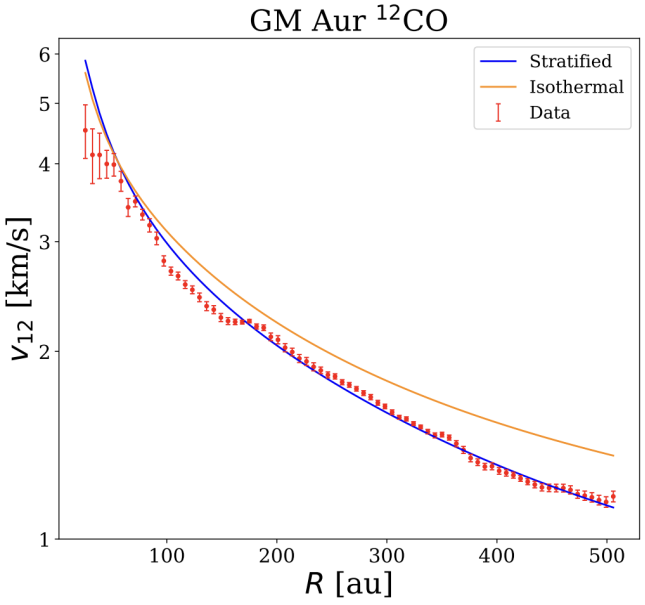


Figure 1.6: Rotation curve of ^{12}CO in the GM Aur disc compared with the isothermal and stratified models. Reproduced from the left panel of Fig. 6 of Martire et al. [10].

In the article by Pezzotta et al. [9] the method is applied to the single disc of HD 163296 in the same way as Martire's, with an identical fitting procedure, but instead

of only two tracers, seven are used simultaneously. This extension shows an improvement in the fit since it reduces the chi squared to about 60% of its initial value, and it makes the method more robust by probing the vertical structure down to the midplane. The errors on the fitted parameters are also ten times smaller than Martire's, but the authors attribute this to the tool used to extract the curves, eddy [16] in place of Discminer[11] [17], since its error bars only contain the statistical part, and not to a real gain in precision. As a consequence the formal precision of these measurements is not a good indicator of their accuracy.

The step from single sources to a sample is taken by Longarini et al. [18], who fit the rotation curves of the ten discs of the exoALMA programme with the same dynamical model and obtain the mass of the star, the mass of the disc and the scale radius of each of them. Two of their conclusions matter here. The first is that the disc cannot in general be treated as a test particle around its star, since fitting a purely Keplerian curve and a full dynamical one return appreciably different stellar masses on the more massive discs. The second, and the one this thesis starts from, is that the thermal structure is what dominates the systematic uncertainty of the whole method, because it is the quantity which sets the pressure gradient.

1.3 Aims and outline

The papers described in the previous section show that it is possible to directly derive the masses of the star and the disc together with the scale radius from a dynamical analysis. They show that assuming a stratified thermal structure improves the physical fit, but its parametric form has been until now derived from an independent measurement of a different nature [15], based on the brightness temperature of optically thick lines [19]. This procedure has limitations of its own which we will discuss in Chapter 3. Above all it never reaches the midplane and therefore the midplane temperature is not measured but extrapolated. Therefore obtaining the temperature parameters with a dynamical measurement is not only more coherent with the rest of the method but also more direct since it rests on more general assumptions. The aim of this work is exactly to try and prove the capability of this method to retrieve the thermal parameters. This will be verified by reiterating the following steps under different conditions:

1. Using the physical model developed by Martire to generate a spatial three-dimensional grid over which values of density, velocity and temperature are known.

2. Applying the radiative transfer software MCFOST to this grid, convolving with a beam size similar to ALMA's[12] and adding thermal noise to reproduce real data.

3. Using Discminer software to extract the rotation curve from the synthetic data.

[11] The software we use in this work as well, described in Chapter 3.

[12] This simply means "blurring" in the same way as the real telescope would, due to its finite resolution.

4. Fitting the model to the synthetic data to try and retrieve the known underlying physical structure.

The steps above are repeated changing the temperature of the disc, the molecular tracers and the quality of the synthetic data. The test is passed when the parameters recovered by the fit agree with the ones used to generate the data, and when they do not agree the difference measures the systematic error introduced by the chain that goes from the model to the extracted curve.

The thesis is organised as follows. In Chapter 2 the model is derived from the conservation of mass and the Navier-Stokes equations, from the diffusion of the surface density up to the rotation curve of a thermally stratified disc. Chapter 3 describes the pipeline used to produce the synthetic observations and to extract the rotation curve from them. Chapter 4 studies with Fisher's method which parameters the rotation curve can constrain, and validates it with fits to synthetic curves with Gaussian noise. Chapter 5 presents the results, the retrieval of the thermal parameters, the bias introduced by the extraction and the attempt at correcting it. Chapter 6 summarises the work and suggests how to continue it.

Chapter 2

The theory

2.1 Gas dynamics in protoplanetary discs

Planets form in protoplanetary discs, doughnut-like structures that are found around newborn stars. The geometry of these discs is a consequence of the conservation of angular momentum along the process of star formation. It all starts with the collapse of a molecular cloud that shrinks down due to gravity and therefore increases its rotational velocity and forms a protostar. The balance between rotation and gravity sets a precise radius inside which material falls onto the central object, while it does not do so outside. This radius, usually found at around 300 au, is therefore the starting point of the disc while the end point is usually found at around 700 au. Systems composed of an accreting young star and its exterior envelope are called Young Stellar Objects and are classified according to their spectral energy distribution [20], a graph of the energy as a function of the frequency. This graph is the result of the emission of the central object, reprocessed by the surrounding envelope. The envelope, being the result of molecular cloud collapse, is mostly made up of gas and dust and the ratio between their masses is usually assumed to be 100. Therefore the disc dynamics are controlled by the gas since it is 100 times more abundant than the dust. However the gas is mostly made up of molecular hydrogen, a symmetric molecule which does not emit through roto-vibrational transitions and is therefore invisible to observations. In order to study the gas, other tracers are then used, like the CO isotopologues [21], whose abundances relative to the molecular hydrogen range from 10^{-7} for $C^{18}O$ to 10^{-4} for ^{12}CO. These tracers emit in the sub-millimetre for their rotational transitions while the dust has a continuum emission.

2.1.1 Disc energetics

Protoplanetary discs are divided into active and passive, depending on the mechanism which produces their luminosity. If a disc is active, material is currently accreting onto the central star and the luminosity is given by [22]:

$$L_{act} = \frac{GM_*\dot{M}}{2R_*} \tag{2.1}$$

intuitively, the falling material heats up through friction, and the thermal energy it emits is related to the gravitational potential energy it possessed at the start of the process. This potential energy is given by the gravitational constant G times the stellar mass M_* times the mass of the falling material M divided by the distance (radius) from the star R. The virial theorem can be applied to a single element of mass giving a total energy of: $-GM_*M/2R$ where the minus is due to the fact that the potential energy is negative since we are dealing with a bound system. This energy, equal in modulus to half of the potential energy, gets converted to luminosity. To get this luminosity we simply have to differentiate the total energy with respect to time getting the accretion rate $\dot{M} = dM/dt$. This shows that half of the potential energy gets emitted while the other half is converted to kinetic energy, which makes sense because a falling object usually accelerates.

If the disc is passive, its luminosity is the result of the reprocessing of the stellar emission and is given by [23, 24]:

$$L_{pas} = \frac{L_*}{4} \tag{2.2}$$

This formula can be intuitively understood via the following reasoning. Imagine sitting on the disc and looking towards the centre; only one fourth of the star is visible since in the sky we can only see half of the sphere and the bottom part is covered by the disc. Therefore only a fourth of the light reaches you.

We can then compare the formulas for the two disc regimes to get an accretion rate threshold to distinguish them:

$$\dot{M} \geq \frac{R_* L_*}{2GM_*} \tag{2.3}$$

Inputting the Sun's values we get $\dot{M} = 10^{-8}\, M_\odot\, \mathrm{yr}^{-11}$ while typical values vary from 10^{-6} to 10^{-8} [25, 26].

2.1.2 Assumptions

In order to describe the accretion process and the disc in general we treat it as a single continuous fluid. The description as a continuum is justified by the average molecule number density, which is around $10^8\,\mathrm{cm}^{-3}$. Treating the gas and the dust as a single fluid is possible due to the gas-to-dust ratio of 100 which makes the dust dynamically irrelevant [27]. We then assume a gas of constant mean molecular weight $\mu = 2.37$, measured in units of the hydrogen mass. We said before that the gas is mainly made up of molecular hydrogen [28], however other gaseous species like helium account for the additional 0.37.

The obvious choice of coordinates to describe this system is the cylindrical one. z is used for the vertical direction, ϕ for the azimuthal angle and R for the radius, as shown in Figure 2.1.

It is reasonable to assume axisymmetry to describe our disc, which implies that all ∂_ϕ are equal to zero where ϕ is the azimuthal angle in the coordinate system of the

[1] In Astronomy, when we want to say that a certain quantity is the one of the sun, we indicate it with the subscript $\odot$. If we mean a generic star we use $*$.

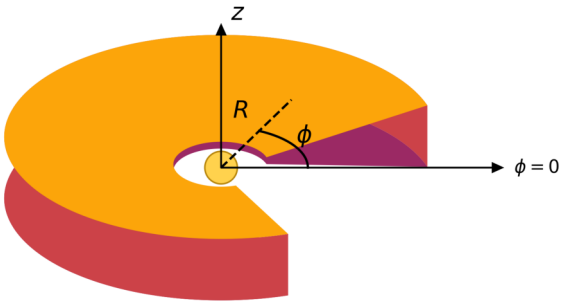


Figure 2.1: Visualisation of the cylindrical coordinates superimposed on a stylised image of a disc.

disc, because there is no preferred direction. Similarly we assume that the observed data is under the condition of stationarity because the viscous time, which will be explained in Section 2.1.5, is of the order of millions of years while the observations only last a couple of hours. This means that the time derivative of the velocity field is equal to zero, because the rotation law adjusts itself on the orbital time which is much shorter than the viscous one. However the surface density does change on the viscous time and its evolution is the subject of Section 2.1.5, which is why its time derivative is the one we keep. The same large time scale, through the slowness of the viscous process which has the effect of dredging mass towards the centre, implies that the radial and vertical velocities of the fluid, which will be written respectively as v_R and v_z, are negligible with respect to the orbital one. The radial velocity is negligible in the balance of forces but not in the conservation of angular momentum. There it is the only term that carries mass across a given radius, and it has nothing to compete with.

2.1.3 The fundamental equations

Having stated our assumptions we then lay the fundamental bricks of our theoretical construction. The first of these bricks is a continuity equation and describes the behaviour of mass through its surface density Σ [29]:

$$\frac{\partial \Sigma}{\partial t} + \frac{1}{R}\frac{\partial}{\partial R}(R\Sigma v_R) = 0 \tag{2.4}$$

This equation can be derived by considering an annulus of the disc between R and $R + dR$ of mass $dM = 2\pi\Sigma R dR$. Differentiating in time we get the rate of mass $r_M = 2\pi R\Sigma v_R$ that crosses the circle of radius R, entering (or exiting)[2] the annulus at the internal radius and exiting (or entering) at the external one. The rate of change of the mass of the annulus $\partial_t(2\pi R\Sigma dR)$ is therefore equal to the difference between $r_M(R)$ and $r_M(R + dR)$ which is $-\partial_R r_M dR$. Dividing by $2\pi R dR$ we get the wanted result.

[2]It depends on the sign of the velocity.

The second fundamental brick is the Navier-Stokes equation which describes the conservation of linear momentum, and whose azimuthal component describes the transport of angular momentum:

$$\rho\left[\frac{\partial \mathbf{v}}{\partial t}+(\mathbf{v}\cdot\nabla)\mathbf{v}\right]=-\nabla P-\rho\nabla\Phi+\nabla\cdot\boldsymbol{\sigma} \tag{2.5}$$

This equation is basically Newton's second law for a fluid element, divided by the volume. On the left-hand side we can see the volumetric density ρ times the total derivative of the velocity, or in other words the acceleration. On the right-hand side we can see a sum of the external forces acting on our fluid element. The first force is simply given by the opposite of the gradient of the pressure P. The second force is obtained by multiplying the volumetric density by the opposite of the gradient of the gravitational potential Φ. The total gravitational potential results from the sum of two terms. The first is simply the gravitational potential of the star, $\Phi_* = -GM_*/r$, where $r = \sqrt{R^2+z^2}$ is the spherical radius and not the cylindrical one. The second is the self-gravitating term due to the mass of the disc itself. This term has a form which is less familiar than the stellar one, and the reason is geometrical. For a spherical body only the mass inside the radius counts, and it pulls as if it were concentrated in the centre, while a disc is flat and this property is lost: every ring of gas attracts our fluid element, the ones inside it and the ones outside it as well. The self-gravitating term is therefore an integral over the whole disc, where the ring of radius r enters with its surface density $\Sigma(r)$ times the radial force it exerts at the position of the element. This force is written with the complete elliptic integrals K and E, which come from going around the ring, since the distance between the element and the ring changes with the azimuth. The rings closest to the element are the ones that weigh the most, while the disc as a whole adds its mass to the one of the star: in the outer part, where almost all of the disc is inside the radius, the rotation curve tends to the keplerian one around $M_* + M_d$. The radial derivative of this potential, which is the quantity that enters the rotation curve, was computed by Bertin and Lodato [30]:

$$R\frac{\partial \Phi_d}{\partial R}(R,z) = G\int_0^\infty \left[K(k)-\frac{1}{4}\left(\frac{k^2}{1-k^2}\right)\left(\frac{r}{R}-\frac{R}{r}+\frac{z^2}{Rr}\right)E(k)\right]\sqrt{\frac{r}{R}}\,k\,\Sigma(r)\,dr \tag{2.6}$$

where $k^2 = 4Rr/[(R+r)^2+z^2]$ measures how close the element is to the ring. Since Σ scales with the mass of the disc, this term is of order M_d/M_* with respect to the keplerian one.

The last force is a bit more complex and is obtained by taking the divergence of the viscous stress tensor $\boldsymbol{\sigma}$, which we describe in the next section.

2.1.4 The viscous stress tensor

The viscous stress tensor is generally determined by the following equation [31] in cartesian coordinates x_1, x_2 and x_3:

$$\sigma_{ij} = \rho\nu\left(\frac{\partial v_i}{\partial x_j} + \frac{\partial v_j}{\partial x_i} - \frac{2}{3}\delta_{ij}\nabla\cdot\mathbf{v}\right) + \rho\zeta\delta_{ij}\nabla\cdot\mathbf{v} \tag{2.7}$$

Each element σ_{ij} gives the force per unit area along i on a surface whose normal points in the j direction. The first term of the definition starts with the product of the volumetric density ρ with the kinematic viscosity ν, which is measured in square centimetres per second and tells us how difficult it is to deform our element of fluid. In fact the velocity derivatives kept inside the bracket of the second factor describe a change in shape without a change in volume. The first derivative tells us the difference in velocity between adjacent layers of fluid, which if not negligible turns cubes of fluid into oblique parallelepipeds, as shown in Figure 2.2.

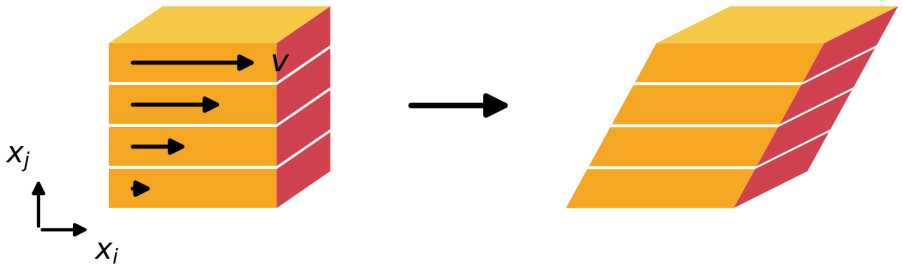


Figure 2.2: Effect of a difference in velocity between adjacent layers of fluid. The layers of the cube on the left move along x_i with a velocity which grows along x_j, and after some time they have slid one over the other, deforming the cube into an oblique parallelepiped.

The second derivative exists because the tensor needs to be symmetrical in order to conserve angular momentum, so this term is an expression of Newton's third law. The third derivative serves to cancel the volume deformations, the ones where the velocity is differentiated with respect to its own direction. The factor 2 is due to the presence of the first two terms that count the diagonal elements twice. The dividing factor 3 is due to the fact that we are describing a three-dimensional system and therefore the tensor has three diagonal terms. The divergence of the velocity gives all of the diagonal terms because a stretch along one direction is significant only in comparison with the stretches along the other directions. For example if the stretch along all directions is equal, the result of the bracket is zero because the shape did not change, only the volume did. Any change in volume is captured by the last term of the definition which starts with a coefficient ζ and is usually negligible for molecular gases. We then take the divergence of the viscous stress tensor because what matters for an element of fluid is the difference between two different application points of the viscous force, much like for the pressure. However complex this definition might look, it simply boils down to one single non-vanishing component. This is due to the fact that only the angular velocity derivative along the radial direction, in cylindrical coordinates, is non-zero. Imagine a keplerian disc where the fluid's angular velocity is determined

by the following centripetal motion equation:

$$\Omega^2 = \frac{GM_*}{R^3} \Longrightarrow \Omega \propto R^{-\frac{3}{2}} \tag{2.8}$$

Then the inner rings of the disc rotate faster than the outer ones and their interaction is captured through the viscosity. The unique non-vanishing component in cylindrical coordinates is:

$$\sigma_{R\phi} = \rho\, \nu R \frac{d\Omega}{dR} \tag{2.9}$$

Having stated our assumptions and having clarified the starting equations, we now proceed to solve them in order to obtain a verifiable description of reality.

2.1.5 Azimuthal component

We start from the azimuthal component of the Navier-Stokes equation, where the viscous term is non-vanishing and plays an important role. Because of axisymmetry the pressure and gravitational potential terms give zero contribution. Here it is important to remember that we are using cylindrical coordinates and therefore the gradients are in cylindrical form. For example the gradient of the pressure gives:

$$\nabla P = \frac{\partial P}{\partial R}\hat{\mathbf{R}} + \frac{1}{R}\frac{\partial P}{\partial \phi}\hat{\boldsymbol{\phi}} + \frac{\partial P}{\partial z}\hat{\mathbf{z}} \tag{2.10}$$

Here the azimuthal component vanishes because of axisymmetry, as stated within the assumptions. The same reasoning applies to the gravitational potential. The advective term of the total derivative of the velocity in cylindrical coordinates reads:

$$(\mathbf{v}\cdot\nabla)\mathbf{v} = \left(v_R\frac{\partial}{\partial R} + \frac{v_\phi}{R}\frac{\partial}{\partial \phi} + v_z\frac{\partial}{\partial z}\right)(v_R\hat{\mathbf{R}} + v_\phi\hat{\boldsymbol{\phi}} + v_z\hat{\mathbf{z}}) \tag{2.11}$$

From which it is easy to select the wanted component. The Navier-Stokes equation for the azimuthal component thus becomes:

$$\rho\left[\frac{\partial v_\phi}{\partial t} + v_R\left(\frac{v_\phi}{R} + \frac{\partial v_\phi}{\partial R}\right)\right] = \nabla\cdot\boldsymbol{\sigma}|_\phi \tag{2.12}$$

Integrating vertically we obtain:

$$\Sigma\left[\frac{\partial v_\phi}{\partial t} + v_R\left(\frac{v_\phi}{R} + \frac{\partial v_\phi}{\partial R}\right)\right] = \nabla\cdot\mathbf{T}|_\phi \tag{2.13}$$

where $\mathbf{T}$ is the vertically integrated stress tensor whose only non-vanishing component is $T_{R\phi} = \Sigma\nu R\Omega'$. Here Ω' is the derivative of the angular velocity with respect to the radius. Axisymmetry and vertical hydrodynamical equilibrium make the angular velocity a function of R alone, therefore this derivative is uniquely defined. In order to obtain a conservation law for the angular momentum we first multiply the whole

equation by R. The radius is a coordinate and does not depend on time, so $R\partial_t v_\phi$ is the time derivative of Rv_ϕ. Similarly the two terms inside the round bracket combine into $\partial_R(Rv_\phi)$. The equation then reads:

$$\Sigma\left[\frac{\partial}{\partial t}(Rv_\phi)+v_R\frac{\partial}{\partial R}(Rv_\phi)\right]=R\nabla\cdot\mathbf{T}|_\phi \tag{2.14}$$

where Rv_ϕ is the angular momentum per unit mass and the bracket is simply its total derivative. We then multiply the continuity equation by Rv_ϕ and add it to the one above. The time derivatives sum to $\partial_t(\Sigma Rv_\phi)$ and the radial ones to $\frac{1}{R}\partial_R(Rv_R\Sigma Rv_\phi)$, due to the fact that in both cases we are adding the two halves of the derivative of a product. In this way we shift from a description that follows the fluid to one written in terms of fluxes.

On the right-hand side we can still see the azimuthal component of the divergence of a tensor. In cylindrical coordinates and for our single component this is given by:

$$\nabla\cdot\mathbf{T}|_\phi=\frac{1}{R}\frac{\partial}{\partial R}(RT_{R\phi})+\frac{T_{R\phi}}{R}=\frac{1}{R^2}\frac{\partial}{\partial R}(R^2T_{R\phi}) \tag{2.15}$$

The additional term $T_{R\phi}/R$ is the result of the curvature of the coordinates and turns $RT_{R\phi}$ into $R^2T_{R\phi}$. The second radius is the lever arm of the viscous stress, and $R^2T_{R\phi}$ is therefore the torque acting on the annulus. Multiplying by the R of the first step and substituting $T_{R\phi}=\Sigma\nu R\Omega'$ we get the explicit angular momentum per unit area ΣRv_ϕ conservation law:

$$\frac{\partial}{\partial t}(\Sigma Rv_\phi)+\frac{1}{R}\frac{\partial}{\partial R}(Rv_R\Sigma Rv_\phi)=\frac{1}{R}\frac{\partial}{\partial R}(\nu\Sigma R^3\Omega') \tag{2.16}$$

Notice that the right-hand side of this equation comes from a divergence and therefore, because of the divergence theorem, is a surface term. The process we are mathematically describing is a transport and not a dissipation. We can define the angular momentum per unit mass $h=Rv_\phi=R^2\Omega$ to rewrite Equation 2.16. Here the temporal derivative $\partial_t(\Sigma h)$ is equal to $\partial_t(\Sigma)h+\Sigma\partial_t(h)$. However h is independent of time, as stated in Section 2.1.2, therefore the second term is equal to zero. The radial derivative $\partial_R(R\Sigma v_R h)$ is equal to $\partial_R(R\Sigma v_R)h+R\Sigma v_R\partial_R h$. Dividing this last sum by R and adding it to the temporal derivative we get, for the left-hand side of Equation 2.16, $h[\partial_t\Sigma+(1/R)\partial_R(R\Sigma v_R)]+\Sigma v_R\partial_R h$. Inside the square brackets we recognise the continuity equation and therefore substitute it with zero. We are left with:

$$\Sigma v_R\frac{dh}{dR}=\frac{1}{R}\frac{\partial}{\partial R}(\nu\Sigma R^3\Omega') \tag{2.17}$$

Isolating v_R and writing h explicitly we get an expression for the radial velocity:

$$v_R=\frac{1}{R\Sigma(R^2\Omega)'}\frac{\partial}{\partial R}(\nu\Sigma R^3\Omega') \tag{2.18}$$

Assuming mass and angular momentum conservation and considering the transport caused by viscosity, we have obtained a formula for the radial velocity that tells us where mass moves inside the disc. To get this formula we injected information from the conservation of mass into the law for the conservation of angular momentum, but we can also go the other way. If we multiply the radial velocity formula by $R\Sigma$ and do some simplification we get:

$$R\Sigma v_R = \frac{1}{(R^2\Omega)'}\frac{\partial}{\partial R}(\nu\Sigma R^3\Omega') \tag{2.19}$$

substituting inside the continuity equation:

$$\frac{\partial \Sigma}{\partial t} = -\frac{1}{R}\frac{\partial}{\partial R}\left[\frac{1}{(R^2\Omega)'}\frac{\partial}{\partial R}\left(\nu\Sigma R^3\Omega'\right)\right] \tag{2.20}$$

This last equation is a diffusion one: one temporal derivative equal to a second derivative with respect to a spatial coordinate, plus some factors deriving from the choice of a coordinate system. The reader might have encountered this equation while studying thermodynamics under the name of heat equation. However in this case the equation is non-linear since ν, the kinematic viscosity coefficient, can depend on the surface density Σ. We will solve this equation in the next subsection because the result will be essential in building our model.

2.1.5.1 Solving the Diffusion Equation

We now want to solve Equation 2.20 for the surface density and we will do it by following the approach by Lynden-Bell and Pringle [32]. All the subsequent theory will use this result.

We start with a variable change by using the angular momentum per unit mass h, the actual transported quantity, and the viscous torque $g = -2\pi\Sigma\nu R^3\Omega'$. Here we treat ν as independent of time and this will become exact once we write it as a power law of R. The resulting equation is:

$$-\frac{1}{2\pi\nu R^3\Omega'}\frac{\partial g}{\partial t} = \frac{1}{2\pi R}\frac{\partial}{\partial R}\left[\left(\frac{\partial h}{\partial R}\right)^{-1}\frac{\partial g}{\partial R}\right] \tag{2.21}$$

which can be tidied up by noting that $\frac{\partial}{\partial R} = \frac{\partial h}{\partial R}\frac{\partial}{\partial h}$:

$$\frac{\partial^2 g}{\partial h^2} = -\left(\nu R^2\Omega'\frac{\partial h}{\partial R}\right)^{-1}\frac{\partial g}{\partial t} \tag{2.22}$$

At this point we make two assumptions with the objective of getting a closed form solution for the equation. The first assumption is that the viscosity coefficient ν can be written as $\nu_0\left(\frac{R}{R_0}\right)^b$ where $b < 2$. The second assumption is that the disc is keplerian which means that the gravitational potential of the disc itself is negligible compared to the stellar one and that the orbital velocity is simply $v_\phi = \sqrt{\frac{GM_*}{R}}$, so that $h = \sqrt{GM_*R}$.

To keep the notation light we define l as $\frac{1}{2(2-b)}$ and k as $\frac{2}{2-b}\left[\frac{R_0^b}{3\nu_0}(GM_*)^{b-2}\right]^{1/2}$. The result is:

$$\frac{\partial^2 g}{\partial h^2} = \frac{1}{4}\left(\frac{k}{l}\right)^2 h^{1/l-2}\frac{\partial g}{\partial t} \tag{2.23}$$

This equation is solvable for $g(h,t)$ through the method of the superposition of modes by supposing that the temporal dependence of each mode is $g \propto e^{-st}$ with $s > 0$. The result of this hypothesis is:

$$\frac{\partial^2 g}{\partial h^2} + \frac{1}{4}\left(\frac{K}{l}\right)^2 h^{1/l-2} g = 0 \tag{2.24}$$

where $K^2 = sk^2$. We then apply the following chain of variable changes:

$$x = h^{1/(2l)} \tag{2.25}$$

$$g_1 = \frac{g}{x^l} \tag{2.26}$$

$$z = Kx \tag{2.27}$$

And finally we get Bessel's equation:

$$z^2\frac{\partial^2 g_1}{\partial z^2} + z\frac{\partial g_1}{\partial z} + (z^2 - l^2)g_1 = 0 \tag{2.28}$$

This is a known differential equation that frequently arises in problems regarding wave motion and heat transport. It was systematically studied by astronomer Friedrich Bessel around 1824. For this reason the following functions, used to solve the equation, are named after him and their expression is:

$$J_l(Kx) = \left(\frac{Kx}{2}\right)^l \sum_{n=0}^{+\infty} \frac{(-1)^n (Kx/2)^{2n}}{n!\,\Gamma(l+n+1)} \tag{2.29}$$

where l does not need to be an integer. The general solution $g(x,t)$ is given by the superposition of all modes which are the product of the temporal part and a spatial part. This second part is given by the linear combination of a Bessel function with parameter l and its symmetric counterpart with parameter $-l$. The coefficients of the linear combination A and B depend on K.

$$g(x,t) = \int_0^{+\infty} \exp\left(-\frac{K^2}{k^2}t\right)(Kx)^l\left[A(K)J_l(Kx) + B(K)J_{-l}(Kx)\right]dK \tag{2.30}$$

The coefficients of the linear combination can be found by setting boundary conditions. We use the physical fact that the star rotates much more slowly than the gas just outside its surface, which moves at the keplerian velocity. Therefore the angular velocity must reach a maximum between the two where its derivative Ω' vanishes [29], and the torque g vanishes there as well since it is proportional to Ω'. The stellar radius is negligible compared to the size of the disc and we can set $g(h,t)$ equal to zero in the origin. From this we get that $B(K)$ must be equal to zero, because near the origin J_{-l} diverges as x^{-l} and once multiplied by $(Kx)^l$ it tends to a non-zero constant, while the

J_l term naturally goes to zero. Using the Fourier-Bessel theorem [33] we can derive the following expression for the coefficient $A(K)$:

$$A(K') = \int_0^{+\infty} g(x,0) J_l(K'x)(K'x)^{1-l} dx \tag{2.31}$$

By choosing a starting condition $g(x,0)$, using it in Equation 2.31, inserting the result back into Equation 2.30 with $B(K)$ set to zero and following the chain of variable changes backwards we are able to compute $\Sigma(R,t)$. Let's apply this method for the following starting condition:

$$g(h,0) = Ch \exp\left(-ah^{1/l}\right) \tag{2.32}$$

where C is a normalisation constant and $a = (GM_* R_0)^{-1/(2l)}$. This condition goes to zero for $h = 0$ as required by the boundary condition, and it is chosen in order to solve the following integrals analytically. However this choice is less arbitrary than it seems, since a diffusion process gradually forgets the details of its initial condition. With the variable x the starting condition becomes $g(x,0) = Cx^{2l} \exp(-ax^2)$, and inserting it in Equation 2.31 we get a tabulated integral [33]:

$$A(K) = CK(2a)^{-(l+1)} \exp\left(-\frac{K^2}{4a}\right) \tag{2.33}$$

We then insert this coefficient back into Equation 2.30. The new integral in K has the same form as the previous one and can be solved in the same way. Going back to the variable h the result is:

$$g(h,t) = ChT^{-(l+1)} \exp\left(-\frac{ah^{1/l}}{T}\right) \tag{2.34}$$

where we defined the dimensionless time $T = 1 + 4at/k^2$. In the keplerian case the definition of the torque gives $g = 3\pi\nu\Sigma h$. Finally, substituting the power law for the viscosity we arrive at the self-similar solution:

$$\Sigma(R,t) = \frac{C}{3\pi\nu_0}\left(\frac{R}{R_0}\right)^{-b} T^{-\eta} \exp\left(-\frac{(R/R_0)^{2-b}}{T}\right) \tag{2.35}$$

where $\eta = (5/2 - b)/(2 - b)$. The quantity T is dimensionless, therefore the coefficient of t inside it must be the inverse of a time. This is the viscous time we used in Section 2.1.2 to justify the condition of stationarity:

$$t_\nu = \frac{k^2}{4a} = \frac{R_0^2}{3(2-b)^2\nu_0} \tag{2.36}$$

The solution is called self-similar because R and t only appear through the combination $(R/R_0)^{2-b}/T$, which means that the profile keeps the same shape as time goes on: it

spreads outwards while its amplitude decreases. For $t = 0$ we have $T = 1$ and the profile becomes $\Sigma \propto (R/R_0)^{-b} \exp[-(R/R_0)^{2-b}]$. This is the surface density we will adopt in our model, where R_0 plays the role of the scale radius.

A practical example of the concepts described in this section can be found in Appendix A, where we describe a small SPH simulation of a viscous disc.

2.1.6 Vertical component

We now solve the vertical component of the Navier-Stokes equation. The time a sound wave needs to cross the thickness of the disc is of the order of the orbital period, which is much shorter than the viscous time and therefore we can assume hydrostatic equilibrium in this direction. Intuitively, the pressure disturbances along the vertical direction are redistributed by sound waves much faster than the disc evolves. For this reason we neglect the vertical component of the velocity and the left-hand side of the equation is zero. Moreover the viscosity's vertical components are zero and therefore we obtain the following equation:

$$\frac{1}{\rho}\frac{\partial P}{\partial z} = -\frac{\partial \Phi}{\partial z} \tag{2.37}$$

This equation exactly describes a vertical hydrostatic equilibrium. Using the ideal gas law, which relates the pressure to the density through the square of the sound speed, $P = c_s^2 \rho$, we can rewrite the left-hand side as a function of the density:

$$\frac{1}{\rho}\frac{\partial (c_s^2 \rho)}{\partial z} = -\frac{\partial \Phi}{\partial z} \tag{2.38}$$

We now solve the equation under strict assumptions in order to give an intuition of how the vertical hydrostatic equilibrium works. For the final model we will relax these assumptions to describe a more general picture. However for now let us assume that the disc is vertically isothermal and that therefore the sound speed does not depend on z. The sound speed is related to the temperature by the ideal gas law:

$$c_s^2 = \frac{k_B T}{\mu m_p} \tag{2.39}$$

where k_B is the Boltzmann constant and m_p is the proton mass. Furthermore we assume that the disc is vertically thin $z \ll R$ and we can therefore approximate the spherical radius r with the cylindrical one R. We then make the final assumption that the disc mass is negligible with respect to the stellar one. The vertical equilibrium becomes:

$$\frac{c_s^2}{\rho}\frac{\partial \rho}{\partial z} = -\frac{GM_*}{R^3} z = -\Omega_k^2 z \tag{2.40}$$

This is a basic differential equation whose solution is a Gaussian. Let us define $H = c_s/\Omega_k$ as the scale height. Notice that c_s is a velocity while Ω_k is an angular velocity and therefore their ratio has the dimension of length. The scale height H is a quantity

that will often come up from now on so it is useful to develop an intuition for it. It basically tells us that the vertical scale of the system we are describing is determined by how fast the pressure disturbances caused by Ω_k can be smoothed out. The explicit solution is:

$$\rho(R,z) = \rho_{\mathrm{mid}}(R)\exp\left[-\frac{z^2}{2H^2}\right] \tag{2.41}$$

where $\rho_{\mathrm{mid}}(R)$ is the density at the midplane. Since H is the standard deviation of the Gaussian we know that 99.7% of the mass of the disc is vertically contained within a distance $3H$ from the midplane. This kind of reasoning applies even with more relaxed starting assumptions.

2.1.7 Radial component

Let us finally consider the radial component of the Navier-Stokes equation. We have established that the accretion process sits on a longer timescale than all the other ones and therefore we neglect the radial component of velocity. We also know that the viscous stress has no radial components. The equation then reads:

$$v_\phi^2 = R\frac{\partial\Phi}{\partial R} + \frac{R}{\rho}\frac{\partial P}{\partial R} \tag{2.42}$$

This equation tells us that the azimuthal velocity is the result of the sum of two different terms. The first one obviously being the gravitational influence of the star and of the disc itself when non negligible. The second term is a pressure gradient determined by density and temperature that typically slows down the rotation and pushes molecules from regions of higher pressure to regions of lower pressure. Let us give an intuition for this equation by solving it under strict assumptions, as for the vertical component case. Let's assume that the disc mass is negligible and that the disc is isothermal. The first term on the right-hand side can be written as:

$$R\frac{\partial\Phi}{\partial R} = \frac{GM_*R^2}{(R^2+z^2)^{3/2}} = v_k^2\left(1+\frac{z^2}{R^2}\right)^{-3/2} \tag{2.43}$$

And the second one:

$$\frac{R}{\rho}\frac{\partial P}{\partial R} = c_s^2\frac{\partial\log\rho}{\partial\log R} = v_k^2\left(\frac{H}{R}\right)^2\frac{\partial\log\rho}{\partial\log R} \tag{2.44}$$

Therefore the azimuthal velocity is:

$$v_\phi^2 = v_k^2\left[\left(1+\frac{z^2}{R^2}\right)^{-3/2} + \left(\frac{H}{R}\right)^2\frac{\partial\log\rho}{\partial\log R}\right] \tag{2.45}$$

By making a further assumption on the functional form of the density, that is $\rho \propto R^{-n}$, we can get a clearer insight into this first simple solution in the midplane:

$$v_\phi = v_k\left[1-n\left(\frac{H}{R}\right)^2\right]^{1/2} \tag{2.46}$$

Here we can see that the pressure gradient adds a negative term scaling as $(H/R)^2$ to the keplerian velocity thus slowing down the rotation.

2.2 The model

Up until now we have stated our assumptions and developed a rigorous theoretical framework for the hydrodynamical analysis of protoplanetary discs. We have solved the corresponding equations along their three components deriving the full solution for the azimuthal one and solving the other two under strict assumptions in order to give an intuition for them. We will now present the full model clearly stating the starting points, the mathematical steps and the final results. We have four equations in total but five unknowns. In the continuity equation we have the surface density Σ and the radial velocity v_R. In the Navier-Stokes equations the unknowns are the three velocity components, the volumetric density ρ and the pressure P. Since the surface density can be obtained by vertically integrating the volumetric one, we have exactly five unknowns corresponding to one vector field and two scalar ones. We are therefore missing a relation to solve the system, which is the energy balance. Solving it would require the radiative transfer of the stellar light through the dust, which we do not do. Radiative transfer is expensive inside a fit which evaluates the model hundreds of thousands of times, and it depends on uncertain quantities like the dust opacities. It would also make the temperature an input of the model, while the aim of this work is to measure it from the kinematics. Radiative transfer is still used in this work, but only to produce the synthetic observations, as explained in Chapter 3. In order to avoid this complexity we cleverly choose a parametrisation for the temperature based on observations, whose parameters can then be fitted. We know from the literature that protoplanetary discs are thermally stratified and not vertically isothermal [15, 19], therefore we set:

$$T(R,z) = T_{\text{mid}}(R) f(R,z) \tag{2.47}$$

where f is a function that describes the dependency of temperature on the height. This relation also implies the following one:

$$c_s^2(R,z) = c_{s,\text{mid}}^2(R) f(R,z) \tag{2.48}$$

For the surface density we choose the self-similar profile because it is the asymptote of every starting condition for the diffusion equation. This can also be seen in Appendix A. Intuitively the self-similar solution makes sense because it is the product between a power law, a general functional form that can describe well the mass content, and an exponential term that smoothly puts an external end to the disc. The position of this outer edge is set by the evolution itself and is measured by the scale radius. The same form is used by the works that this thesis follows [10, 11]. We report the self-similar density solution for the convenience of the reader:

$$\Sigma(R) = \frac{(2-\gamma)M_d}{2\pi R_c^2}\left(\frac{R}{R_c}\right)^{-\gamma}\exp\left[-\left(\frac{R}{R_c}\right)^{2-\gamma}\right] \tag{2.49}$$

where M_d is the mass of the disc, R_c is the scale radius and γ plays the role of the exponent b of Equation 2.35, while the normalisation ensures that the total mass of the disc is M_d. From here on we fix $\gamma = 1$ as in the works we follow [10, 18], since the rotation curve constrains it poorly and its choice mainly affects the estimate of the scale radius [14].

We note that the self-similar solution was derived assuming that the disc rotates at the keplerian velocity, while the discs we are trying to model do not because the temperature introduces an additional term to the rotation curve. We have shown that in general the angular velocity differs from the keplerian one by quantities of order $(H/R)^2$ and M_d/M_*, and the evolution equation depends on the rotation only through the specific angular momentum, therefore the shape of the profile changes by the same amount. Since we use it as a description of the disc at the time of the observation, with M_d and R_c fitted to the data, a deviation of this size is absorbed by the parameters themselves. The disc density at the midplane is thus:

$$\rho_{\text{mid}}(R) = \frac{\Sigma(R)}{\sqrt{2\pi} H_{\text{mid}}(R)} \tag{2.50}$$

where $H_{\text{mid}} = c_{s,\text{mid}}/\Omega_k$ is the scale height computed with the midplane temperature. This relation follows from integrating the Gaussian profile of Section 2.1.6 along the vertical direction, since $\Sigma = \int \rho \, dz = \sqrt{2\pi} H \rho_{\text{mid}}$. Next we write the general volumetric density as:

$$\rho(R,z) = \rho_{\text{mid}}(R) g(R,z) \tag{2.51}$$

where g describes how the density changes with the height and should therefore be equal to one at the midplane. Unlike f, whose form we can choose, g is not free since it follows from the vertical hydrostatic equilibrium, as we will show. Using the ideal gas law $P = c_s^2 \rho$ we can finally write the pressure as:

$$P(R,z) = P_{\text{mid}}(R) f\, g(R,z) = c_{s,\text{mid}}^2(R) \rho_{\text{mid}}(R) f\, g(R,z) \tag{2.52}$$

where in the last equality we have used the ideal gas law at the midplane. The weight of being the unsolvable unknown is now moved from the pressure to the function f, for which there exist various parametrisations in the literature [19, 34] that avoid solving the energy balance.

Now we will solve the vertical hydrostatic equilibrium to get $g(R,z)$. As in Martire et al. [10], we take into account the potential of the star alone along the vertical direction while keeping the self-gravity of the disc in the radial component. This is because in the vertical direction the self-gravity is of the order of $1/Q$ with respect to the stellar gravity, where Q is the Toomre parameter [35]. To gain an intuition for this reasoning, imagine a mass element of fluid at a certain height above the midplane. Two forces pull it back down, the vertical component of the gravity of the star and the gravity of the gas found between the mass element and the midplane. They both are functions of the height of the same order since the gravitational law is the same and therefore their ratio along this direction is a real number. As seen in Section 2.1.6 when the star's gravity dominates the scale height is: $H = c_s/\Omega_k$. If the self-gravity term is the dominant one we have $H_{sg} = c_s^2/(\pi G \Sigma)$ [36]. Thus, the ratio between the two is:

$$\frac{H_{sg}}{H} = \frac{c_s \Omega_k}{\pi G \Sigma} = Q \tag{2.53}$$

The two gravitational terms are comparable when Q approaches one and it is said that in this regime the disc is unstable. For the discs considered in this work Q is around 5 in the inner region and therefore our assumption is justified.

We now solve the vertical hydrostatic equilibrium. The Navier-Stokes equation along the vertical direction is:

$$\frac{1}{\rho}\frac{\partial P}{\partial z} = -\frac{\partial \Phi_*}{\partial z} \tag{2.54}$$

where, using the assumption of no self-gravity along this direction, the gravitational potential is:

$$\Phi_*(R,z) = -\frac{GM_*}{\sqrt{R^2+z^2}} \Longrightarrow \frac{\partial \Phi_*}{\partial z} = \Omega_k^2 z\left[1+\left(\frac{z}{R}\right)^2\right]^{-3/2} \tag{2.55}$$

And based on Equation 2.52 the left-hand side is:

$$\frac{1}{\rho}\frac{\partial P}{\partial z} = c_{s,\mathrm{mid}}^2 f \frac{\partial \log(f\,g)}{\partial z} \tag{2.56}$$

We can solve for $\log(f\,g)$ obtaining:

$$\log(f\,g) = -\frac{1}{H_{\mathrm{mid}}^2}\int_0^z \frac{z'}{f(R,z')}\left[1+\left(\frac{z'}{R}\right)^2\right]^{-3/2} dz' \tag{2.57}$$

From which we can get an expression for the density only dependent on f and not g:

$$\rho(R,z) = \frac{\rho_{\mathrm{mid}}(R)}{f(R,z)}\exp\left[-\frac{1}{H_{\mathrm{mid}}^2}\int_0^z \frac{z'}{f(R,z')}\left[1+\left(\frac{z'}{R}\right)^2\right]^{-3/2} dz'\right] \tag{2.58}$$

Now choosing a parametric form for f determines both unknowns of pressure and density.

Solving the radial component of Navier-Stokes's equation as in Section 2.1.7 tells us how to compute the orbital velocity:

$$v_\phi^2(R,z) = \frac{R}{\rho}\frac{\partial P}{\partial R}(R,z) + R\frac{\partial \Phi_*}{\partial R}(R,z) \tag{2.59}$$

The first term, the partial derivative of the pressure with respect to the radius, can be written as[3]:

$$\frac{R}{\rho}\frac{\partial P}{\partial R} = c_{s,\mathrm{mid}}^2 f\left[\frac{\partial \log P_{\mathrm{mid}}}{\partial \log R} + \frac{\partial \log(f\,g)}{\partial \log R}\right] \tag{2.60}$$

And the second term:

$$R\frac{\partial \Phi_*}{\partial R} = v_k^2\left[1+\left(\frac{z}{R}\right)^2\right]^{-3/2} \tag{2.61}$$

Finally the rotation curve in the self-similar case is:

$$\begin{aligned} v_\phi^2(R,z) = v_k^2\Bigg\{&\left[1+\left(\frac{z}{R}\right)^2\right]^{-3/2} \\ &-\left[\gamma' + (2-\gamma)\left(\frac{R}{R_c}\right)^{2-\gamma} - \frac{\partial \log(f\,g)}{\partial \log R}\right]\left(\frac{H_{\mathrm{mid}}}{R}\right)^2 f(R,z)\Bigg\} \\ &+ R\frac{\partial \Phi_d}{\partial R} \end{aligned} \tag{2.62}$$

[3] More details can be found in Appendix A of Martire et al. [10].

where $\gamma' = \gamma + (3 - q_{\rm mid})/2$ collects the radial slopes of the surface density, of the temperature and of the scale height. The two terms inside the curly brackets have a different physical origin. The first term corrects the keplerian velocity for the height above the midplane. The second term originates from the pressure and is determined by density and temperature through the function f and the hydrostatic equilibrium. The term outside the brackets is the self-gravitating one.

In this thesis we are going to use two temperature prescriptions: the isothermal one and Dartois'. The first case is the simpler one where $f = 1$ and the rotational velocity reduces to:

$$v_\phi^2(R,z) = v_k^2 \Bigg\{ 1 - \left[\gamma' + (2-\gamma) \left(\frac{R}{R_c} \right)^{2-\gamma} \right] \left(\frac{H_{\rm mid}}{R} \right)^2 + q_{\rm mid} \left[1 - \left(1 + \left(\frac{z}{R} \right)^2 \right)^{-1/2} \right] \Bigg\} + R \frac{\partial \Phi_d}{\partial R} \tag{2.63}$$

while the density becomes:

$$\rho(R,z) = \rho_{\rm mid}(R) \exp \left\{ \frac{R^2}{H_{\rm mid}^2} \left[\left(1 + \left(\frac{z}{R} \right)^2 \right)^{-1/2} - 1 \right] \right\} \tag{2.64}$$

Martire et al. [10] parametrise the stratified thermal structure with the prescription of Dullemond et al. [19]. In this work we adopt the one of Dartois, Dutrey, and Guilloteau [34], as in Longarini et al. [18], which smoothly connects the midplane, where $z = 0$, with the rest of the disc and in whose form the thermal structures of the exoALMA discs have been measured [37]. The f is:

$$f(R,z) = \begin{cases} \dfrac{T_{\rm atm}(R) + [T_{\rm mid}(R) - T_{\rm atm}(R)] \cos^2 \left(\dfrac{\pi z}{2 z_q} \right)}{T_{\rm mid}(R)} & z < z_q \\ \dfrac{T_{\rm atm}(R)}{T_{\rm mid}(R)} & z \geq z_q \end{cases} \tag{2.65}$$

$$z_q(R) = Z_0 \left(\frac{R}{100\,{\rm au}} \right)^\beta$$

The radial dependence of the two temperatures is described by two power laws [18]:

$$T_{\rm mid}(R) = T_{\rm mid,100} \left(\frac{R}{100\,{\rm au}} \right)^{q_{\rm mid}}, \qquad T_{\rm atm}(R) = T_{\rm atm,100} \left(\frac{R}{100\,{\rm au}} \right)^{q_{\rm atm}} \tag{2.66}$$

where $T_{\rm mid,100}$ and $T_{\rm atm,100}$ are the temperatures at 100 au, $q_{\rm mid}$ and $q_{\rm atm}$ are their slopes, which are negative because the disc cools outwards, while Z_0 and β of Equation 2.65 give the height of the transition between the two and its radial dependence. The rotation curve keeps the form of Equation 2.62, since the only term which changes with f is the logarithmic one. Since f has two branches, the integral of Equation 2.57

has to be split at z_q. Below the transition it becomes:

$$\log(f\,g) = -\frac{1}{H_{\rm mid}^2}\int_0^z \frac{z'\,T_{\rm mid}}{T_{\rm atm}+(T_{\rm mid}-T_{\rm atm})\cos^2\left(\frac{\pi z'}{2z_q}\right)}\left[1+\left(\frac{z'}{R}\right)^2\right]^{-3/2} dz', \qquad z < z_q \tag{2.67}$$

This integral has no closed form, therefore we compute it numerically together with its logarithmic derivative. Above the transition f is constant and the remaining part of the integral can be solved analytically:

$$\log(f\,g) = \log(f\,g)\big|_{z_q} - \frac{T_{\rm mid}}{T_{\rm atm}}\frac{R^2}{H_{\rm mid}^2}\left\{\left[1+\left(\frac{z_q}{R}\right)^2\right]^{-1/2} - \left[1+\left(\frac{z}{R}\right)^2\right]^{-1/2}\right\}, \qquad z \geq z_q \tag{2.68}$$

The density is then given by Equation 2.58 with this f. Equations 2.49, 2.39, 2.6, 2.57, 2.58, 2.62, 2.63, 2.64, 2.65, 2.66, 2.67 and 2.68 constitute our model on which the rest of the thesis is built. Section 2.2.2 collects them in one place, together with the free parameters of the two cases, five for the isothermal one and nine for the stratified one. The objective of this work is to verify, through a synthetic pipeline, when and if different combinations of the thermal parameters are actually recoverable from real data.

2.2.1 Emission surfaces

As explained in the introduction of this thesis, ALMA's observations are datacubes of specific isotopologues. Each isotopologue emits from a different height inside the disc and therefore the extracted rotation curve we see is actually the $v_\phi(R,z)$ at a specific $z(R)$ called emission surface profile. Information about this profile is needed for two reasons. Firstly, as we will see in Section 3.2.2, this information is used to actually extract the rotation curve data from the datacubes by the software Discminer. Secondly, we need to know the specific height of the velocity field in order to correctly fit it with our model.

In this section we will briefly describe the chemical mechanism behind the emission surface profile and talk about the uncertainties when measuring it. In the gas phase, molecules like ^{12}CO are present only in an intermediate layer of the disc. Different mechanisms are at hand: the dense midplane is shielded from the stellar radiation and is cold enough for the molecules to freeze onto the dust grains, while in the rarefied upper layers the ultraviolet radiation is not shielded and photodissociates them. Inside this layer, much like for a star, there exists a surface under which the gas is simply too thick to let photons pass and therefore we mostly see the gas near the height at which light becomes able to escape. Mathematically, this surface is defined as the vertical height $z(R)$ for which the optical depth τ_ν, depending on the frequency ν, is equal to one. The formula is:

$$\tau_\nu(s) = \int_{s_i}^{s} k_\nu(R(s'),z(s'))\rho(R(s'),z(s'))ds' = 1 \tag{2.69}$$

where k_ν is the isotopologue specific opacity and s is the spatial coordinate along the line of sight of the observer. The integral goes from s_i, the location of the observer, to

s, a point inside the disc. The functions $R(s)$ and $z(s)$ depend on the inclination and position angle of the disc.

The emission surface $z(R)$, solution of Equation 2.69, is usually parametrised in the literature [15, 17] as:

$$z(R) = z_0 \left(\frac{R}{100\,\mathrm{au}} \right)^{\psi} \exp\left[-\left(\frac{R}{R_t} \right)^{q_t} \right] \tag{2.70}$$

The four parameters are the height of the surface at 100 au z_0, its slope ψ, which sets the flaring, the taper radius R_t, beyond which the surface bends back towards the midplane, and the taper exponent q_t, which sets how steeply it does so. In this work we are going to use the software Discminer to extract the parameters of the emission surface to use them in our pipeline. The height of the emission surface is uncertain, and for the faintest lines the relative uncertainty can reach 50% [9]. We note that this extraction process is subject to known methodological difficulties which will be explained in Section 3.2.

2.2.2 The model in summary

The pieces of the model are scattered across the derivations of this chapter, and the chapters which follow use them constantly. We therefore collect them here, first the part which is common to the two thermal prescriptions, then the isothermal case and then the stratified one. The equations are repeated without a number, and each of them carries the reference to the point where it was obtained.

The surface density is the self-similar profile, with $\gamma = 1$ throughout this work,

$$\Sigma(R) = \frac{(2-\gamma) M_d}{2\pi R_c^2} \left(\frac{R}{R_c} \right)^{-\gamma} \exp\left[-\left(\frac{R}{R_c} \right)^{2-\gamma} \right] \qquad \text{(Equation 2.49)}$$

the midplane temperature is a power law,

$$T_{\mathrm{mid}}(R) = T_{\mathrm{mid},100} \left(\frac{R}{100\,\mathrm{au}} \right)^{q_{\mathrm{mid}}} \qquad \text{(Equation 2.66)}$$

and it fixes the sound speed and with it the scale height of the midplane, $H_{\mathrm{mid}} = c_{s,\mathrm{mid}}/\Omega_k$,

$$c_s^2 = \frac{k_B T}{\mu m_p} \qquad \text{(Equation 2.39)}$$

The vertical structure follows from the hydrostatic equilibrium and depends on the thermal prescription only through the function $f = T(R,z)/T_{\mathrm{mid}}(R)$,

$$\rho(R,z) = \frac{\rho_{\mathrm{mid}}(R)}{f(R,z)} \exp\left[-\frac{1}{H_{\mathrm{mid}}^2} \int_0^z \frac{z'}{f(R,z')} \left[1 + \left(\frac{z'}{R} \right)^2 \right]^{-3/2} dz' \right] \qquad \text{(Equation 2.58)}$$

with $\rho_{\rm mid} = \Sigma/(\sqrt{2\pi}H_{\rm mid})$. The rotation curve is

$$v_\phi^2(R,z) = v_k^2\left\{\left[1+\left(\frac{z}{R}\right)^2\right]^{-3/2} - \left[\gamma' + (2-\gamma)\left(\frac{R}{R_c}\right)^{2-\gamma} - \frac{\partial\log(fg)}{\partial\log R}\right]\left(\frac{H_{\rm mid}}{R}\right)^2 f(R,z)\right\} + R\frac{\partial\Phi_d}{\partial R} \quad \text{(Equation 2.62)}$$

where $\gamma' = \gamma + (3-q_{\rm mid})/2$ and the last term, the self-gravity of the disc, is the integral of Equation 2.6. Finally, the curve is not observed at the midplane but on the emission surface of each tracer,

$$z(R) = z_0\left(\frac{R}{100\,\text{au}}\right)^\psi \exp\left[-\left(\frac{R}{R_t}\right)^{q_t}\right] \quad \text{(Equation 2.70)}$$

which is an input of the model and not a parameter we fit.

In the isothermal case $f = 1$ and both the curve and the density take a closed form,

$$v_\phi^2(R,z) = v_k^2\left\{1 - \left[\gamma' + (2-\gamma)\left(\frac{R}{R_c}\right)^{2-\gamma}\right]\left(\frac{H_{\rm mid}}{R}\right)^2 + q_{\rm mid}\left[1-\left(1+\left(\frac{z}{R}\right)^2\right)^{-1/2}\right]\right\} + R\frac{\partial\Phi_d}{\partial R} \quad \text{(Equation 2.63)}$$

$$\rho(R,z) = \rho_{\rm mid}(R)\exp\left\{\frac{R^2}{H_{\rm mid}^2}\left[\left(1+\left(\frac{z}{R}\right)^2\right)^{-1/2} - 1\right]\right\} \quad \text{(Equation 2.64)}$$

and the free parameters are five, the mass of the star M_*, the mass of the disc M_d, the scale radius R_c, the temperature at 100 au $T_{\rm mid,100}$ and its slope $q_{\rm mid}$.

In the stratified case f is the prescription of Dartois, Dutrey, and Guilloteau [34],

$$f(R,z) = \begin{cases} \dfrac{T_{\rm atm}(R) + [T_{\rm mid}(R) - T_{\rm atm}(R)]\cos^2\left(\dfrac{\pi z}{2z_q}\right)}{T_{\rm mid}(R)} & z < z_q \\ \dfrac{T_{\rm atm}(R)}{T_{\rm mid}(R)} & z \geq z_q \end{cases} \quad \text{(Equation 2.65)}$$

$$z_q(R) = Z_0\left(\frac{R}{100\,\text{au}}\right)^\beta$$

with the temperature of the atmosphere given by its own power law,

$$T_{\rm atm}(R) = T_{\rm atm,100}\left(\frac{R}{100\,\text{au}}\right)^{q_{\rm atm}} \quad \text{(Equation 2.66)}$$

The rotation curve keeps the form written above, but the logarithmic term has to be split at the transition height. Below it the integral has no closed form and we compute

it numerically,

$$\log(f\,g) = -\frac{1}{H_{\rm mid}^2}\int_0^z \frac{z'\,T_{\rm mid}}{T_{\rm atm} + (T_{\rm mid} - T_{\rm atm})\cos^2\left(\frac{\pi z'}{2z_q}\right)}\left[1 + \left(\frac{z'}{R}\right)^2\right]^{-3/2} dz'$$

(Equation 2.67)

while above it f is constant and the integral is analytic,

$$\log(f\,g) = \log(f\,g)\big|_{z_q} - \frac{T_{\rm mid}}{T_{\rm atm}}\frac{R^2}{H_{\rm mid}^2}\left\{\left[1 + \left(\frac{z_q}{R}\right)^2\right]^{-1/2} - \left[1 + \left(\frac{z}{R}\right)^2\right]^{-1/2}\right\}$$

(Equation 2.68)

The free parameters are now nine, the five of the isothermal case plus the temperature of the atmosphere at 100 au $T_{\rm atm,100}$, its slope $q_{\rm atm}$, the transition height at 100 au Z_0 and its slope β. Whether and when these nine can be recovered from a rotation curve is the question of Chapter 4.

Chapter 3

Testing pipeline

In this chapter we are going to illustrate the pipeline used for producing realistic synthetic data. As explained in the state of the art section, in order to validate theoretical models we test them against synthetic data which was generated from chosen parameters. We start by generating a three-dimensional grid with MCFOST over which we compute at each point the density, temperature and velocity with our model. We then use interpolation to fill in the gaps of this grid. MCFOST then transforms it into an ideal datacube. We add thermal noise and convolve the datacube with the beam, and through this process we get our synthetic data. We extract the rotational curve using Discminer which is the state of the art for kinematic data retrieval. We will illustrate this process by using the disc of LkCa 15 as an example, with the parameters measured by Longarini et al. [18]. The example uses the vertically isothermal prescription. It is the simpler of the two, it is the one the sample of Chapter 5 is built on, and the stratified prescription enters the pipeline in exactly the same way, so nothing of what follows depends on the choice.

3.1 MCFOST

MCFOST is a three-dimensional radiative transfer code based on the Monte Carlo method [38]. The radiation emitted by the star is simulated as a large number of photon packets which are scattered, absorbed and re-emitted by the dust as they travel through the disc, and the radiation field is reconstructed from the statistics of their paths. The transfer of molecular lines is treated separately: the code computes the populations of the rotational levels of the molecule, here in local thermodynamic equilibrium, and integrates the line and continuum opacities along the line of sight to produce a spectral cube. In this work the density, the temperature and the velocity field are the ones of our model, therefore the code solves neither the thermal balance nor the vertical structure of the disc.

3.1.1 The grid

MCFOST provides us with an empty grid that we can fill with our model. The grid is cylindrical with 100 radial cells spaced logarithmically, 50 vertical cells and a single azimuthal cell because of axisymmetry. We can set the bounds and for our example with LkCa 15 we choose 10 to 700 au. The logarithmic spacing increases the resolution in the inner region where the rotation curve varies faster. This same grid is shared by all of our experiments, since it depends only on the domain and on the resolution.

Evaluating our model over the grid provided is computationally expensive because the self-gravity and the pressure term have no closed form in the stratified case, as can be seen in Equations 2.6 and 2.67. We also want to explore many combinations of the parameters during the fits, and computing these two term would take too much time. We therefore tabulate them on their own grids, which cover the whole region the fits can explore, and we interpolate. The self-gravity table depends only on the radius, the height and the scale radius, so it is the same for every disc and is built once. The pressure table instead needs an axis for every thermal parameter which is left free, and it is therefore needed only in the stratified case: the disc of this chapter is vertically isothermal, its vertical integral has the closed form of Equation 2.64, and we evaluate it directly without tabulating anything. The self-gravity grid has 50 radial points, a vertical spacing of 2.5 au and 20 points in the scale radius, which for our example means about 260 thousand points, each one requiring the integral over the whole disc. We compute it with unit masses, since the term is the product of the ratio between the disc mass and the stellar mass with a kernel which depends only on the radius, the height and the scale radius, and we restore the masses afterwards. The pressure grid, when it is needed, has 100 radial points, a vertical spacing of 1.25 au and 9 points for every thermal parameter left free, which in the stratified case are the temperature at 100 au and its slope. We compute it for one solar mass, since the stellar mass enters only through the scale height, and we rescale it afterwards.

The axes of the two grids for our example are reported in Table 3.1. They are used both to fill the cells for MCFOST and to evaluate the model during the fits. The interpolation is not allowed to extrapolate, so that a point which falls outside the covered domain is rejected instead of being invented, and the extent of the thermal axes is also the interval that the fits can explore.

3.1.2 The synthetic disc

Once every grid is built we evaluate our model at the centre of every cell and we write three quantities. The first one is the volumetric density from the vertical hydrostatic equilibrium, normalised so that its vertical integral gives the self-similar surface density. The second one is the temperature of the prescription we are testing. The third one is the azimuthal velocity of Equation 2.62, which already contains the pressure gradient and the self-gravity. MCFOST reads these three quantities instead of building the structure of the disc, therefore it solves neither the vertical equilibrium nor the thermal balance.

For the example of this chapter we use the parameters measured for LkCa 15 [18]: a

Axis	Interval	Self gravity	Pressure
Radius [au]	10 – 700	50	100
Height [au]	0 – 650	260	520
Scale radius [au]	10 – 400	20	–
$T_{\mathrm{mid},100}$ [K]	7 – 65	–	9
q_{mid}	−0.78 – 0.42	–	9

Table 3.1: Axes of the two tabulated grids and number of points along each of them. The first three axes are those of the self-gravity table, which does not depend on the parameters of the disc and is therefore identical for every model we simulate. The last two belong to the pressure table and exist only in the stratified case, where they are centred on the temperature of the disc we simulate, from 0.2 to 1.9 times its value, and are wider than the physical intervals of Table 4.1 because they also have to contain the region explored by the fits.

stellar mass of 1.118 solar masses, a disc mass of 0.108, a scale radius of 150 au and a temperature which falls as $R^{-0.23}$. For its value at 100 au we do not take the 20 K measured at the midplane but 34 K, which is the average between that value and the 48 K measured in the atmosphere [37]. The reason is that at this stage our model is vertically isothermal and therefore has a single temperature, which has to play two roles at once. It sets the pressure gradient of the rotation curve, where the value of the midplane is the right one, and it sets the brightness of the lines, which in a real disc comes from the warm atmosphere. At 20 K the synthetic cubes are half as bright as the observed ones and at the sensitivity of exoALMA the channel maps are dominated by the noise. The choice is a compromise, because the disc we simulate is thicker than the real one, and what this costs us on the extraction of the rotation curve is measured in Chapter 5.

For every tracer we also write its abundance, which we keep constant through the disc for simplicity. The height of the emitting layer is set by this value alone, since the line is optically thick and we observe it where the optical depth reaches unity, as described in Section 2.2.1, and that height depends on how many molecules lie above. We therefore tune the abundance of each tracer so that its surface reproduces the one measured for the exoALMA discs [18, 37]. The values we adopt are listed in Table 3.2, relative to molecular hydrogen. A constant abundance is a simplification, because in a real disc the molecule is destroyed by the ultraviolet radiation in the upper layers and freezes onto the dust grains close to the midplane, so that it survives only in an intermediate layer. Our models therefore lack chemical and thermal realism, and computing the abundances and the temperature of the gas self consistently is one of the natural developments of this work.

For the example of this chapter we use two molecules, the ^{12}CO and the ^{13}CO, whose abundances are chosen so that their surfaces sit at z/R of 0.29 and 0.16 at 100 au, the heights measured for the discs of exoALMA [18, 37]. The choice is verified afterwards, because MCFOST returns the surface at which the optical depth of the line reaches unity: the surface of the ^{12}CO comes out at 29 au at 100 au with a flaring exponent of 1.08, against the 29.0 au and the 1.06 tabulated for LkCa 15 [18]. We use more than

one tracer because each of them emits from a different height, and it is the vertical structure which carries the thermal information. We follow in this the approach of Pezzotta et al. [9], who fits seven molecules of the same disc simultaneously, and the gain is quantified in Section 4.2.2. Figure 3.1 shows the structure of the disc in the R, z plane with the two surfaces drawn on top.

Parameter	Symbol	Value
Stellar mass	M_* [$M_\odot$]	1.118
Disc mass	M_d [$M_\odot$]	0.108
Scale radius	R_c [au]	150
Midplane temperature at 100 au	$T_{\mathrm{mid},100}$ [K]	20
Midplane temperature slope	q_{mid}	−0.23
Atmospheric temperature at 100 au	$T_{\mathrm{atm},100}$ [K]	48
Atmospheric temperature slope	q_{atm}	−0.55
Transition height at 100 au	Z_0 [au]	55
Transition height slope	β	0.59
^{12}CO surface	z_0 [au], ψ, R_t [au], q_t	29.3, 1.08, 639, 6.82
^{13}CO surface	z_0 [au], ψ, R_t [au], q_t	15.8, 1.04, 306, 4.74
^{12}CO abundance	X_{12}	2.9×10^{-7}
^{13}CO abundance	X_{13}	2.3×10^{-8}

Table 3.2: Parameters of LkCa 15 used throughout this thesis. The masses and the scale radius are the ones measured by Longarini et al. [18], the thermal structure the one measured by Galloway-Sprietsma et al. [37], and the emission surfaces are the ones we impose when we generate the model, chosen so that each tracer emits from the height measured for the exoALMA discs. The vertically isothermal model of this chapter uses a single temperature of 34 K, the average between the midplane and the atmosphere, for the reason explained above.

3.1.3 Radiative transfer and the cube

MCFOST integrates the equation of radiative transfer along the line of sight for the two isotopologues and returns a spectral cube, that is one image of the sky for every velocity channel. We compute the rotational transition $J = 3{-}2$ of both isotopologues, at 345.796 GHz for the ^{12}CO and at 330.588 GHz for the ^{13}CO, which are the lines observed by exoALMA [39]. The disc also contains dust, with a mass equal to one hundredth of the mass of the gas, and the dust emits a continuum which the code computes together with the line: we subtract it channel by channel, as is done on the observations, so that what is left is the line alone. As explained in Chapter 1, the line which each parcel of gas emits is not infinitely narrow. Part of its width is thermal, because the molecules move, and at the temperature of our disc it amounts to about 140 m/s for the CO, and part comes from the turbulent motions which we do not resolve, which the code adds in quadrature as a parameter. We set the latter to 20 m/s, a few per cent of the sound speed, which is the upper limit measured in real discs [40]. The wider the line, the wider the bands of Figure 3.2. The disc is placed at 157 pc, inclined by 50.4 degrees and rotated by a position angle of 61.9 degrees, the values

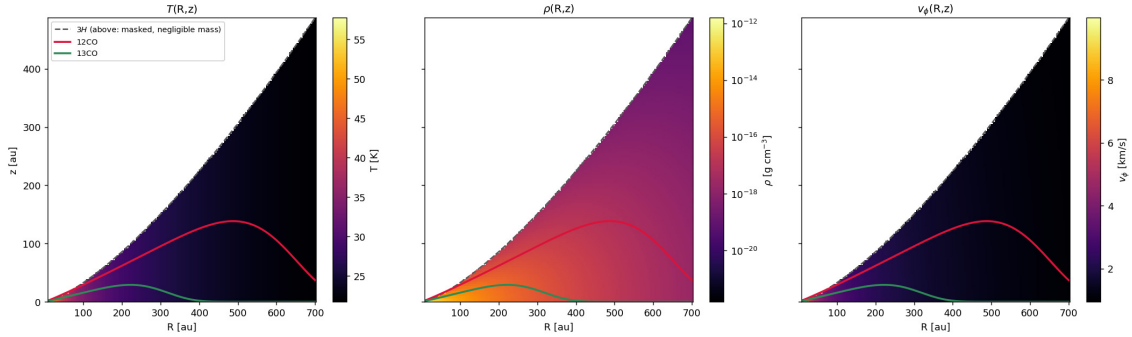


Figure 3.1: Vertical structure of the synthetic disc used as an example in this chapter, at the parameters with which it is generated. Temperature, density and azimuthal velocity in the R, z plane, with the emission surfaces of the two tracers drawn on top. The region above three scale heights is masked, since it contains a negligible fraction of the mass of the column.

measured for LkCa 15 [18]. The geometry of the cube is chosen on the observations we are imitating, the ones of the exoALMA programme [39]: 80 channels of 0.1 km/s, which cover the whole line, and 512 pixels of 0.025 arcsec, six of which fit inside the beam we will use. The result is an ideal observation, without any resolution limit and without noise. Figure 3.2 shows one of its channels, 0.4 km/s away from the systemic velocity. The emission traces the band along which the velocity projected on the line of sight equals the one of the channel, and it leaves dark the regions which rotate too fast or too slowly to contribute to it. The two tracers give the same picture, the ^{13}CO being the smaller of the two because it emits from a lower and more compact layer.

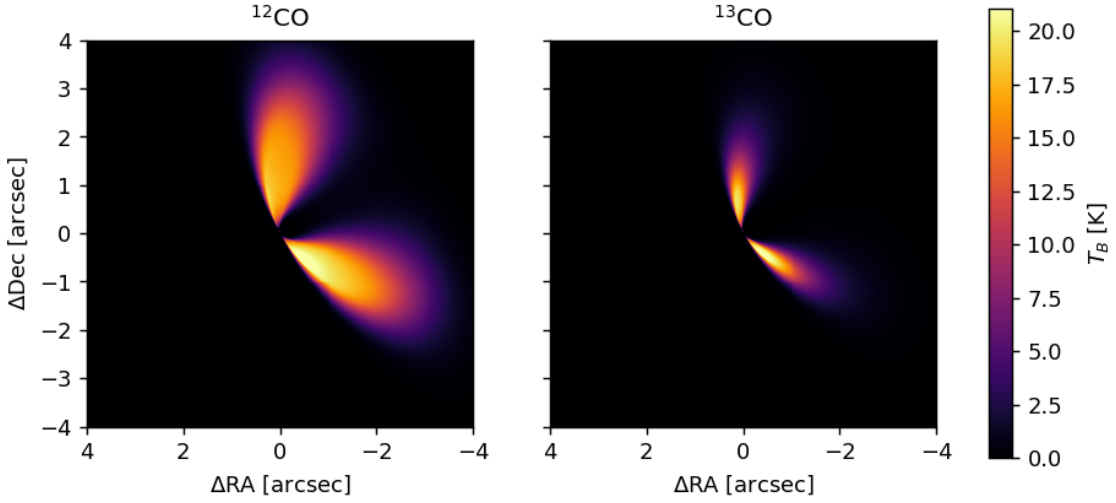


Figure 3.2: A channel 0.4 km/s away from the systemic velocity, of the ideal cubes returned by MCFOST for the two tracers, in brightness temperature. The colour scale is the same in Figures 3.3 and 3.4, so that the three steps can be compared.

The first thing we do to this ideal cube is to add the noise. In a real observation the noise is produced by the receivers, so it enters the visibilities and goes through the imaging together with the signal: we therefore add it before the convolution,

drawing it independently in every pixel of every channel. Its amplitude is chosen so that the map we end up with has a root mean square of 3 mJy per beam, which at these frequencies is a brightness temperature of 1.4 K, the sensitivity of about 1.5 K per channel reached by the exoALMA observations [39]. Before the beam the noise is much larger than that, about 12 K in a single pixel, and Figure 3.3 shows that at this stage the disc barely stands out of it.

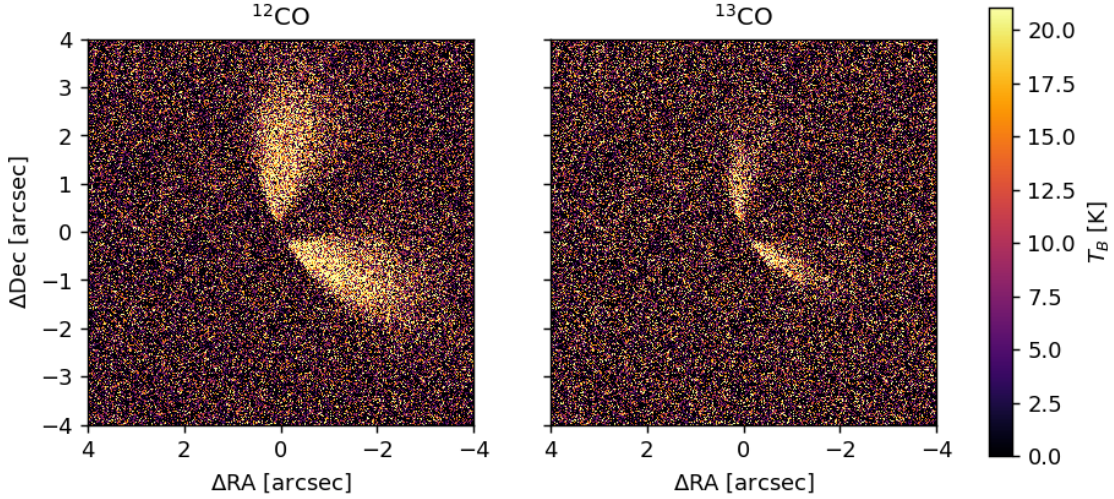


Figure 3.3: The same channel after the noise has been added to the cube, before the convolution. In a single pixel the noise is about 12 K, comparable with the emission, and the disc is barely visible.

The cube still has the resolution of the model, so we convolve every channel with a circular beam of 0.15 arcsec, the resolution reached by exoALMA [39], which at this distance is 24 au. The beam covers some forty pixels, and averaging them lowers the noise by a factor of nine while it leaves the signal, which is spread over much larger scales, almost untouched. This is how the disc of Figure 3.4 emerges from the grain of Figure 3.3. The emission itself looks very similar, what is lost is the sharp structure. The brightness is affected instead in the channels at the two ends of the line, which come from the innermost region where the rotation is fast and the emission covers three or four beams. Here the peak falls by eighteen per cent at 3 km/s from the systemic velocity and by one half at the very ends. This loss is the beam smearing, and it is the reason for the correction described in Section 3.2.3.

3.2 Discminer

Discminer is the software we use to extract the rotation curve from the cubes [17], and it is the same one applied to the discs of MAPS and of exoALMA [41, 42]. Its idea is to fit the whole datacube with a model for it. It builds a synthetic cube from an analytic prescription of the disc, whose parameters are the ones of Table 3.3, and it compares it with the data channel by channel, exploring them with a Markov chain. The com-

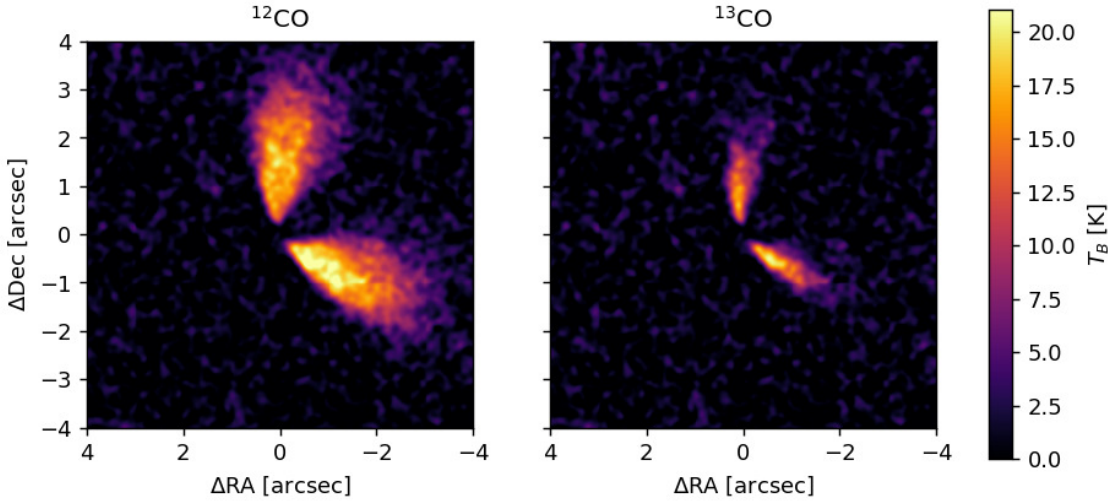


Figure 3.4: The same channel after the convolution with the circular beam of 0.15 arcsec, which at the distance of LkCa 15 is 24 au. The beam averages the noise down to 1.4 K and the disc emerges. This is the synthetic observation from which the rotation curve is extracted.

parison is made on the cube binned by seven pixels[1], which keeps every evaluation of the likelihood fast enough to run a chain of a few thousand steps. The prescription contains the geometry of the disc as well, the inclination and the position angle, but we keep those fixed at the values of the model. Leaving them free adds two parameters which the lines constrain poorly and which are degenerate with the stellar mass, so that the curve extracted afterwards comes out worse, and the geometry of a disc is in any case measured more accurately on the dust continuum, which is brighter and more compact than the lines. The velocity field of this model is keplerian, and this is what makes the comparison useful. Everything which is not keplerian, the gradient of the pressure, the self-gravity of the disc and the kinks produced by an embedded planet, is absent from the model by construction and therefore survives in the residuals between the data and the best fit.

Three of these parameters describe the brightness, and they do so with a power law broken in two, R^{p_0} inside $R_{\rm break}$ and R^{p_1} outside it, with the two branches joined without a jump at the break, which is also the radius where I_0 is measured. A single slope would have to average two different regimes, the inner one where the line is optically thick and the brightness follows the temperature, falling slowly, and the outer one where the gas runs out, the line becomes thin and the brightness collapses.

Four of these parameters describe the profile of the emission line, the intensity of a

[1] Binning by seven means averaging blocks of seven by seven pixels into a single one, which takes the map from 0.025 to 0.175 arcsec per pixel, just above the beam. The cube is oversampled, with six pixels across the beam, and since both the signal and the noise are correlated on that scale the neighbouring pixels carry almost the same information, so that binning them costs almost nothing and makes every evaluation of the likelihood fifty times cheaper.

Parameter	Symbol	Meaning
Stellar mass	M_*	Mass of the star
Peak intensity	I_0	brightness of the emission line at the break radius
Inner slope	p_0	radial slope of the brightness inside the break
Outer slope	p_1	radial slope of the brightness outside the break
Break radius	R_{break}	radius at which the two slopes meet
Line width	L_0	width of the profile at 100 au
Width slope	p_L	radial slope of the width
Line slope	L_s	steepness of the wings of the profile
Slope gradient	p_s	radial slope of L_s
Surface height	z_0	height of the emitting layer at 100 au
Flaring	ψ	how quickly the layer rises with the radius
Taper radius	R_t	radius beyond which the layer bends back
Taper exponent	q_t	slope of the bend

Table 3.3: Free parameters of the Discminer fit. The last four describe one emission surface and are fitted twice, once for the upper surface and once for the lower one, so that the parameters are seventeen in total. The outer radius of the emission and the vertical dependence of brightness and width are kept fixed.

pixel as a function of velocity. Discminer does not use a gaussian but a bell,

$$I(v) = \frac{I_{\text{peak}}}{1 + \left| \frac{v - v_0}{L} \right|^{2L_s}} \tag{3.1}$$

where the width is set by L and the shape is set by L_s. The larger L_s, the flatter the top and the steeper the wings, while $L_s = 1$ gives a lorentzian. This form reproduces the lines of an optically thick disc better than a gaussian, which is why it is the one adopted by exoALMA [42]. The three numbers which describe the bell at a given pixel are its centre, its width and its height. The centre is the velocity at which the profile sits, what is called the centroid of the line, and it is the number we take as the velocity of the gas on that line of sight. Figure 3.6 shows these profiles measured on our cube, with the fit drawn on top of them.

A disc seen at an inclination shows two emitting surfaces. The molecular layer wraps the midplane from above and from below, so a line of sight which enters the disc meets first the surface of the near side, and then the one of the far side. Figure 3.5 shows the geometry. The two surfaces lie at the same height but on opposite sides of the midplane, therefore at the same position on the sky the gas rotates with velocities whose projections along the line of sight are different, and the spectrum of a single pixel carries two components instead of one, as panels 6 and 11 of Figure 3.6 show. For this reason we model the line of every pixel with two bells, one for each surface, the kernel which Discminer calls doublebell [42]. A single bell would be forced to place its centroid somewhere between the two peaks, and the velocity it returns would belong to neither surface. The two components can be combined in two ways. They can be added, which is the right thing to do when the line is thin enough that both

surfaces reach us, or one can keep at every velocity the brighter of the two, which is what happens when the line is so thick that the near side absorbs the far one. We use the second for both tracers, because it is the one which fits our cubes better, and the difference is largest for the ^{12}CO, which is the thicker of the two.

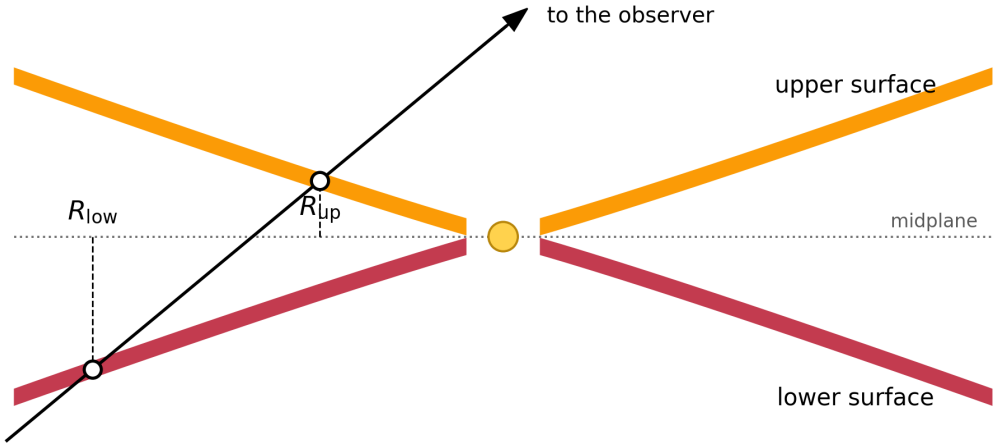


Figure 3.5: The two emission surfaces of an inclined disc

The best fit of this model is not a measurement for us. Its velocity field is keplerian, therefore its stellar mass absorbs everything which the model does not contain, and taking it as an estimate of the mass of the star would mean attributing to gravity what belongs to the pressure. From this fit we take two things only, the emission surface, which is needed to deproject the velocities, and the maps of the line centroid, from which the rotation curve is extracted. The parameters of our model are estimated afterwards, with a fit to the extracted curve.

The parameters of the best fit show this well. The stellar mass comes out at 1.02 solar masses against the 1.118 of the model, and the disagreement is the one we have just described, since the keplerian mass has to account for the pressure gradient as well and it does so by coming out smaller. The emission surface sits lower than the one from which the line really emerges in our disc, about half of it at 100 au and two thirds of it between 200 and 400 au. Surfaces extracted from the data are known to be biased downwards in the outer disc [43], and in a parametric fit like this one the surface also absorbs part of what the rest of the model cannot reproduce, so it is better read as an effective surface than as a measurement of the layer where the optical depth reaches unity.

The quality of the fit is judged on the maps and on the channels themselves. Figure 3.7 compares the map of the centroid velocity of our synthetic data with the one of the best fit, and Figure 3.8 does the same on the channel maps.

3.2.1 Brightness temperature

Radio observations express the intensity of a line as a brightness temperature, that is the temperature at which a black body would emit the intensity we measure. The

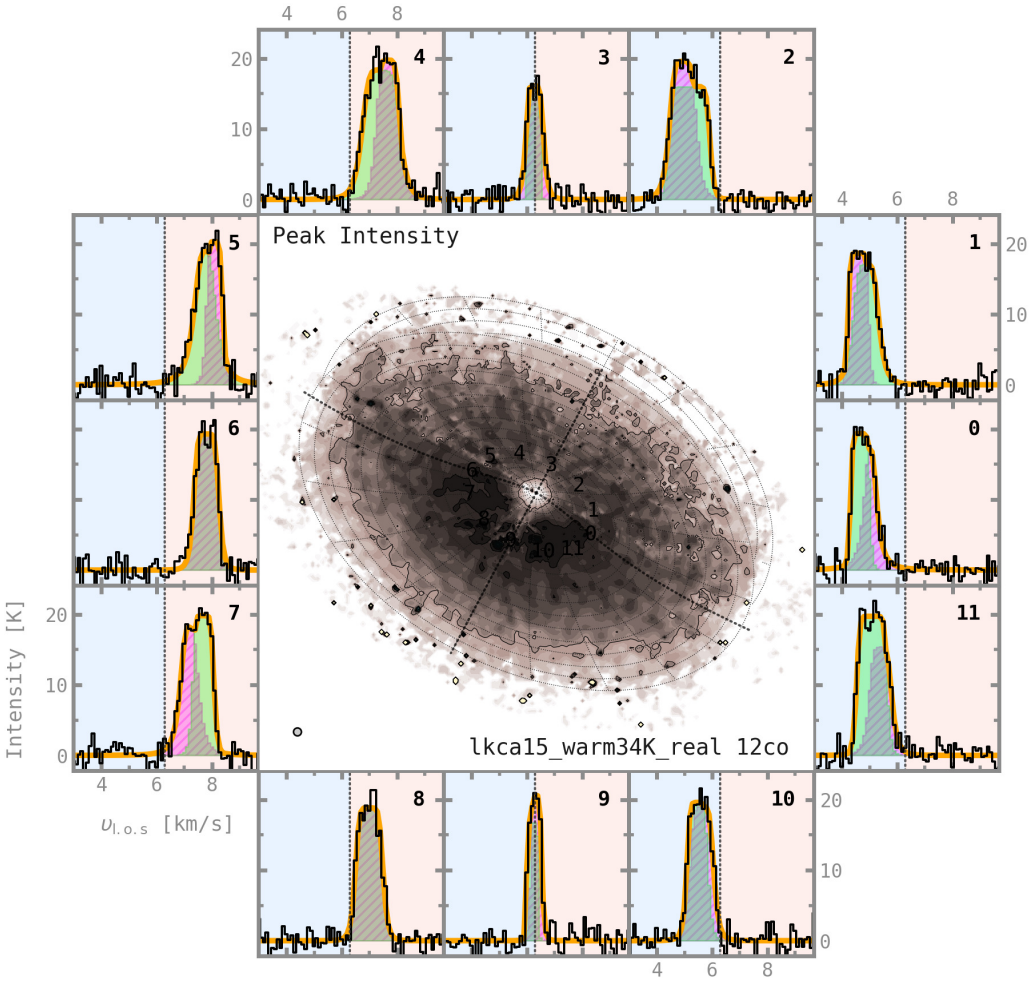


Figure 3.6: Line profiles of the ^{12}CO measured on our synthetic cube along the annulus at 200 au. The central panel is the map of the peak intensity with the twelve azimuthal positions marked on it, and the panels around it are the spectra of those pixels, the observed one in black and the fit of the doublebell in orange. The two shaded components are the two emission surfaces. Close to the minor axis, panels 3 and 9, the two surfaces are seen at almost the same projected velocity and the line has a single peak, while away from it the two components separate and the profile becomes double, as in panels 6 and 11. Produced with the `spectra` tool of Discminer [17].

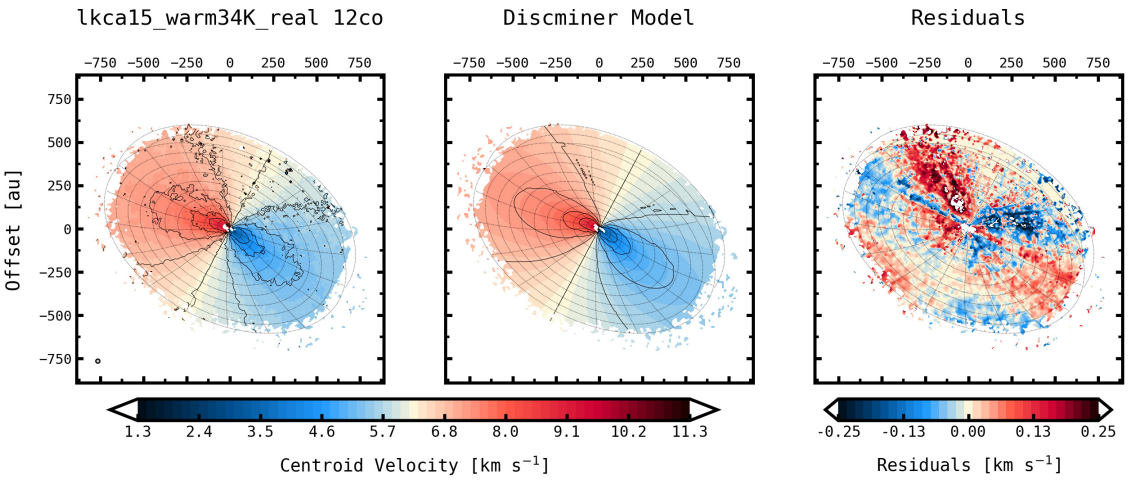


Figure 3.7: Map of the centroid velocity of the ^{12}CO, for our synthetic data on the left, for the best fit of Discminer in the middle and for their difference on the right. The dotted lines are the rings and the axes of the disc, drawn on the fitted surface. Produced with the `moment+residuals` tool of Discminer [17].

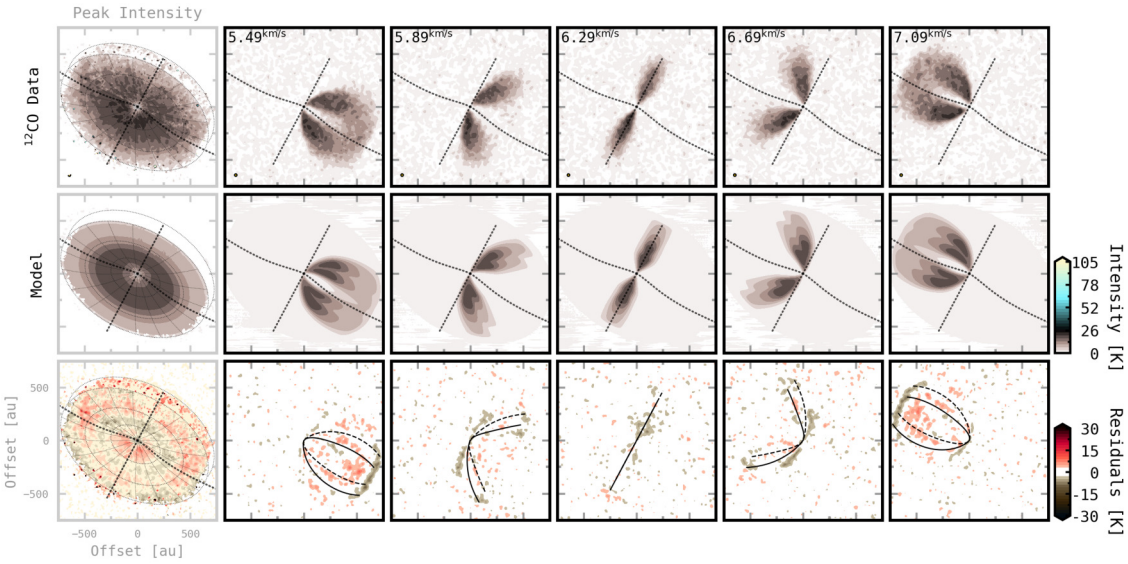


Figure 3.8: Data, model and residuals of the ^{12}CO, in brightness temperature. The first column is the map of the peak intensity and the others are five channels around the systemic velocity, the fourth of which is the one of Figures 3.2 to 3.4. The rows are, from top to bottom, our synthetic data, the best fit of Discminer and their difference. Produced with the `channels+peakint` tool of Discminer [17].

reason for this choice is that our lines are optically thick, so that the radiation which reaches us leaves the disc where the optical depth reaches unity, and in local thermodynamic equilibrium it carries the temperature of the gas at that point. The peak of the line at every pixel is therefore a measurement of the temperature of the emitting layer.

We measure the temperature following the procedure used on the observations [37, 44]. We take the peak intensity of the line, that is the maximum of the spectrum of every pixel, we convert it into a temperature with the full Planck function, and we average it azimuthally along the emission surface, on the same annuli and with the same deprojection with which the rotation curve is extracted. The two profiles are then fitted with a power law between 60 and 400 au.

Figure 3.9 shows the result. The temperature at 100 au is recovered within fifteen per cent, 29.0 K for the ^{12}CO and 29.5 K for the ^{13}CO against the 34 K of the model, which is the accuracy quoted for this measurement [45]. The slope is not because the power laws fitted to the two profiles give -0.057 and -0.382 against the -0.230 of the model. All of this, moreover, concerns the temperature at the emission surface, which for our tracers lies tens of au above the midplane. The temperature of the midplane does not enter these numbers and can only be reached indirectly. However, in our model this makes no difference, since the temperature does not change with height at this stage and the layer and the midplane are at the same temperature. In a real disc the two differ, and the brightness of the line gives the atmosphere while the pressure gradient of Equation 2.62 is set by the midplane. It is precisely because this measurement cannot reach the midplane that we try to reach it with the rotation curve, and the section which follows describes how we extract it.

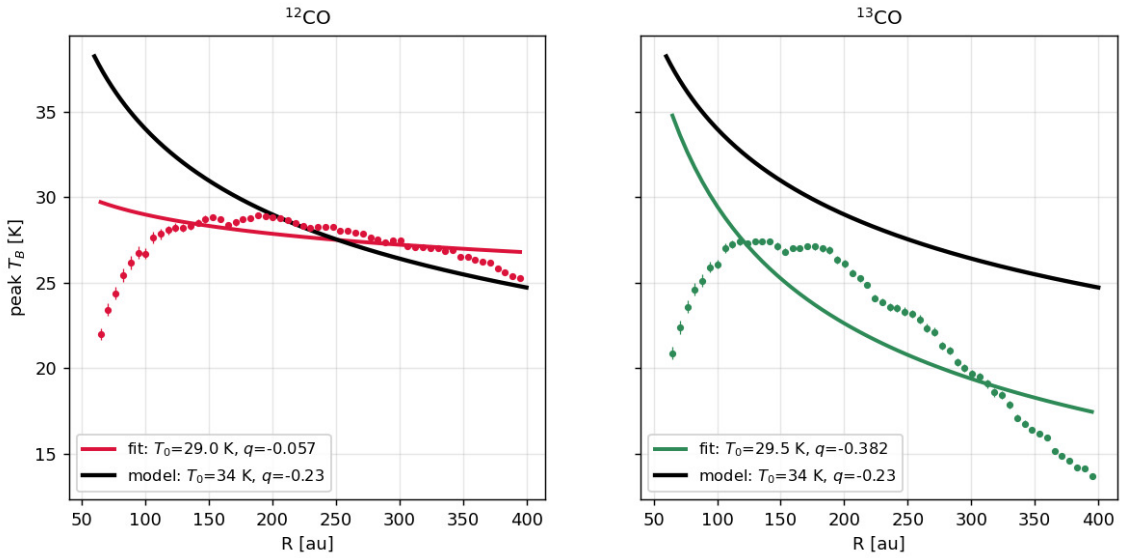


Figure 3.9: Peak brightness temperature of the two tracers as a function of the radius. The points are the azimuthal averages, the coloured lines the power laws fitted to them and the black line the real temperature of the model. Everything is calculated at the emission surface fitted by Discminer.

3.2.2 Rotation curve extraction

The fit gives, at every pixel, the centre of the profile of each of the two surfaces, and we keep the one of the upper surface. That velocity is the projection along the line of sight of the motion of the gas at the point where the line of sight meets the surface. In order to turn it into a rotation velocity we need to know which point of the disc that is, and this is what the fitted surface is for. Discminer inverts the projection and assigns to every pixel the radius and the azimuth of the point it is seeing, and the rotation velocity follows from the geometry,

$$v_\phi = \frac{v_{\rm los} - v_{\rm sys}}{\sin i \, \cos \phi} \tag{3.2}$$

where ϕ is measured from the major axis of the disc.

The points are then collected in annuli of constant radius drawn on the surface, spaced by a quarter of a beam, and in each annulus we take the median of the deprojected velocities. The thirty degrees around the minor axis are discarded. There $\cos \phi$ goes to zero and the deprojection of Equation 3.2 diverges, so that a small error on the velocity along the line of sight becomes an arbitrarily large error on v_ϕ. The extraction is made on the cube binned by two pixels instead of seven, because here we want the finest sampling we can afford and not the speed of the likelihood.

The error bar of each point of the curve is the dispersion of the values along the annulus divided by the square root of the number of beams which fit along it. This is not thermal noise because it contains everything which makes the velocity change with the azimuth at a given radius, from the noise of the cube to the imperfections of the deprojection, and it is the error which Chapter 4 uses as the noise of the curve.

Figure 3.10 shows the curves extracted from the two tracers, together with the curve of our model evaluated on the same surface, and the residual between the two. The residual is not noise, and it has the same shape for both tracers. Inside 60 au it plunges, down to −30 per cent for the ^{12}CO and −43 for the ^{13}CO at 30 au, where the disc is barely wider than a beam. Between 60 and 400 au it settles on a plateau which is slightly negative, −6.3 per cent on average for the ^{12}CO and −2.7 for the ^{13}CO, with a dispersion of 1.7 and 2.0 per cent. Further out it changes sign, at about 320 au for the ^{13}CO and only beyond 450 au for the ^{12}CO, and it grows outwards until it reaches five and seven per cent at 600 au and eleven and sixteen at 700 au. A residual of this kind, smooth and shared by the two tracers, cannot be averaged away, and Chapter 5 is about what it does to the parameters we are after.

3.2.3 Beam smearing correction

The collapse of the residual inside 60 au is not a property of the disc but of the telescope. Section 3.1.3 showed that the beam takes half of the peak away in the channels at the two ends of the line, which are the ones emitted by the fast gas of the inner disc. What the beam does to the brightness it does to the velocity as well, because where the disc is only a few beams wide a single resolution element collects gas whose projected velocities differ by more than the width of the line itself. The profile which comes out

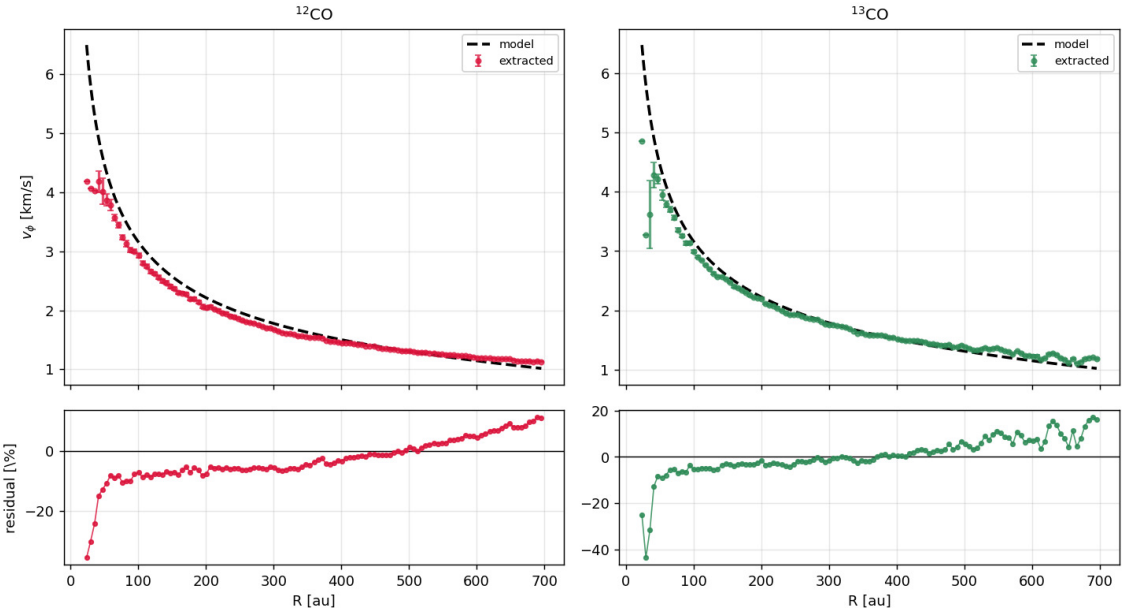


Figure 3.10: Rotation curves extracted from the two tracers of our example, with the curve of the model evaluated on the same emission surface drawn as a dashed line, and the residual between the two underneath. The error bars are the ones described in the text and are smaller than the points over most of the disc.

of that mixture is broader than the profile of any of the gas elements inside it, and its centre lies closer to the systemic velocity than the fastest of them, so that the velocity we extract is systematically too small. This is the beam smearing, and it is the reason why analyses of real discs usually discard the radii within two beams of the centre.

We correct those radii instead of discarding them, and we can do it thanks to Discminer. The best fit of Discminer is a disc in exact keplerian rotation, so its rotation curve is known analytically from its stellar mass and its emission surface. We therefore pass the model cube through the same beam and the same extraction as the data, and we compare the curve which comes out with the analytic one. Their ratio,

$$\xi(R) = \frac{v_{\rm Kep}(R)}{v^{\rm ext}_{\rm Kep}(R)} \tag{3.3}$$

where $v^{\rm ext}_{\rm Kep}$ is the extracted one, is the factor by which our own procedure slows down a curve whose truth we know, and we multiply the data by it, error bars included.[2]

Figure 3.11 shows what it does. The factor is 1.23 for the ^{12}CO and 1.17 for the ^{13}CO at the innermost radius we keep, it falls to a few per cent above unity by 60 au and it stays there, which is the behaviour expected of something produced by the resolution. Between 25 and 60 au the residual goes from −16.8 to −8.7 per cent for the ^{12}CO and from −18.9 to −7.6 for the ^{13}CO, and between 60 and 400 au it goes from −6.3 to −3.6 and from −2.7 to −1.4. Beyond 400 au the factor is within three per cent of unity

[2]The correction is due to Pezzotta and collaborators, who kindly shared the recipe with us ahead of its publication.

and the residual which grows outwards survives untouched. The correction gives us back the inner disc, which is the region which weighs most on the stellar mass, and leaves the outer residual where it was, which is already an indication that the latter does not come from the resolution.

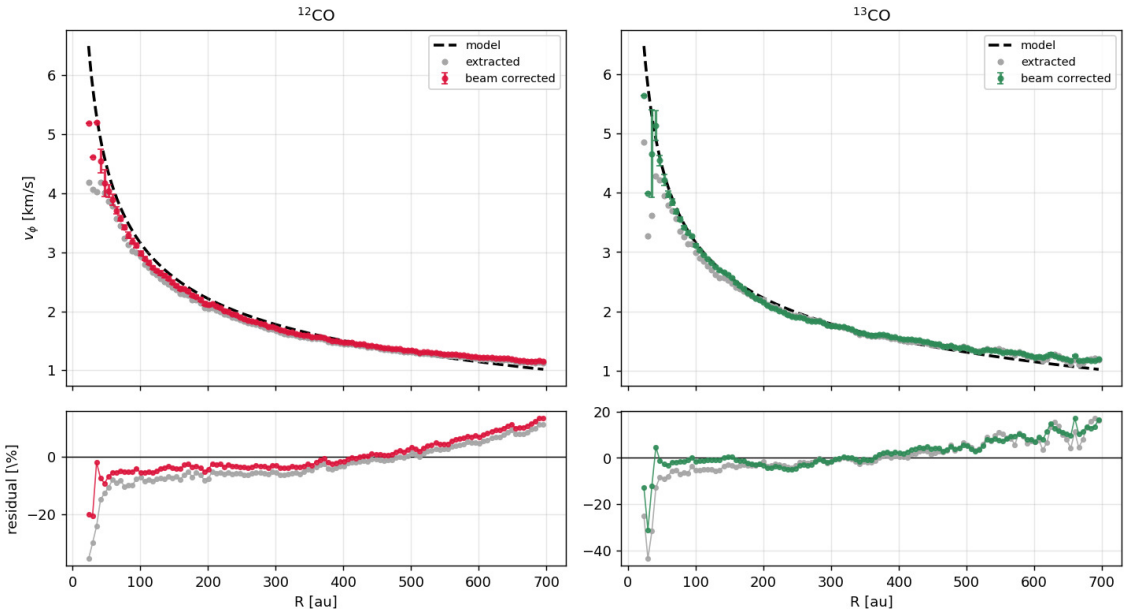


Figure 3.11: Rotation curves of the two tracers after the beam smearing correction, in colour, with the extracted curves of Figure 3.10 underneath them in grey and the curve of the model as a dashed line. The residuals with respect to the model are shown below, for the corrected curve and for the raw one.

Chapter 4

Parameter space

This chapter was not part of the original idea for this thesis but is a consequence of the trial and error process. The thermal parameters, as will be shown, are particularly difficult to recover and therefore a more thorough analytical study was needed. In this chapter we apply Fisher's method first to the isothermal model and then to the stratified one. For each of them we characterise the degeneracies which belong to the model itself and we study how they change with the number of tracers and with the radial range, which guides the choices adopted in Chapter 5. Finally we validate the method with fits to synthetic curves with Gaussian noise.

4.1 Fisher's method

The problem we want to solve is characterising how the rotation curve reacts to changes in parameters before we fit it with our model. This will allow us to concentrate our efforts in the right direction. We introduce in this section an analytical method to study a priori how well each parameter of our model can be retrieved. The connection with the uncertainties on the parameters comes from the likelihood of the data given our model, which for Gaussian and independent errors is written through the chi squared [46] as:

$$\mathscr{L}(\theta) \propto \exp\left[-\frac{1}{2}\chi^2(\theta)\right], \qquad \chi^2(\theta) = \sum_{n=1}^{N} \frac{\left[v_n - v_\phi(R_n;\theta)\right]^2}{\sigma_n^2} \tag{4.1}$$

where v_n is the measured velocity at the n-th point and $v_\phi(R_n;\theta)$ is the one of our model. Let us expand the chi squared around its minimum, which is the maximum of the likelihood, up until the second order:

$$\chi^2(\theta) \simeq \chi^2(\hat{\theta}) + \frac{1}{2}\sum_{ij} \frac{\partial^2\chi^2}{\partial\theta_i\partial\theta_j}\Delta\theta_i\Delta\theta_j, \qquad \Delta\theta_i = \theta_i - \hat{\theta}_i \tag{4.2}$$

The linear term at a minimum is equal to zero while the quadratic one survives. The

second-order derivatives of the chi squared explicitly are:

$$\frac{\partial^2 \chi^2}{\partial \theta_i \partial \theta_j} = 2\sum_{n=1}^{N} \frac{1}{\sigma_n^2} \frac{\partial v_\phi(R_n)}{\partial \theta_i} \frac{\partial v_\phi(R_n)}{\partial \theta_j} - 2\sum_{n=1}^{N} \frac{v_n - v_\phi(R_n;\theta)}{\sigma_n^2} \frac{\partial^2 v_\phi(R_n)}{\partial \theta_i \partial \theta_j} \tag{4.3}$$

If we average this expression over all of the possible noise realisations the second term amounts to zero because it contains the difference between the data and the model. The first term of the right-hand side of Equation 4.3, averaged over the realisations of the noise, is exactly twice Fisher's matrix. This holds if the model is correct, because otherwise the residuals keep a systematic component and their average is not zero. To read this expression consider our rotation curve made up of N points at regular radial intervals, a generic parameter θ, and a vector D_θ of dimension N. Inside this vector the n-th element is the orbital velocity of our model derived with respect to θ, evaluated at the n-th point and divided by the noise at that point σ_n, so that each element measures the shift of the curve in units of the error of that point. The first term of expression 4.3 is then twice the scalar product between two of these vectors. The matrix whose elements are defined as $F_{ij} = D_i \cdot D_j$ for two generic parameters i and j is called Fisher's matrix. The length of D_θ tells us how much the curve reacts to changes in the θ parameter, while the cosine between the vectors of two different parameters tells us how much their effects on the curve overlap. The correlation between the errors of two parameters is read from the inverse matrix and, in the case of two parameters, it is equal to minus this cosine. Fisher's matrix is therefore the curvature of the chi squared around its minimum. From this reasoning we see that the chi squared can be approximated, around its minimum, as the sum of a constant and a quadratic form. The likelihood is then approximated by a Gaussian whose exponent contains Fisher's matrix, which is therefore the inverse of the covariance of the parameters. We can then extract the marginalised errors[1] directly from the diagonal elements of the covariance matrix [47]:

$$\sigma(\theta_i) = \sqrt{(F^{-1})_{ii}} \tag{4.4}$$

The Cramér-Rao theorem [48] tells us that the marginal error obtained through this method constitutes an upper limit on the precision obtainable by any unbiased estimator, including the MCMC fits we are going to use. We can also extract the conditional errors[2]:

$$\sigma_{\text{cond}}(\theta_i) = \frac{1}{\sqrt{F_{ii}}} \tag{4.5}$$

The ratio of these two uncertainties measures how well the other parameters can reproduce the effect of θ_i on the curve:

$$\frac{\sigma(\theta_i)}{\sigma_{\text{cond}}(\theta_i)} = \frac{1}{\sqrt{1-r_i^2}} \tag{4.6}$$

where r_i is the multiple correlation coefficient between θ_i and all the other parameters, which reduces to the correlation coefficient r_{ij} between the two parameters when

[1] The uncertainty of a parameter obtained by fitting on data while the other parameters are unknown and free to move.

[2] The uncertainty of a parameter obtained by fitting on data while all of the other parameters are fixed at the right value.

these are only two [47]. And finally we can infer a degeneracy ellipse between two parameters, which is the set of points at constant $\Delta\chi^2$ around the minimum:

$$\Delta\theta^T C^{-1} \Delta\theta = \Delta\chi^2, \qquad C = F^{-1} \tag{4.7}$$

This can be done by computing the eigenvalues and the eigenvectors of the block of the covariance matrix corresponding to the two parameters. Each axis is oriented along the direction of the corresponding eigenvector, and its length is the square root of the corresponding eigenvalue times the square root of $\Delta\chi^2$, which for two parameters is 2.30 at the 68% level [47]. Since the eigenvectors live in the space of the parameters, their directions are linear combinations of the two, and they are meaningful once the two axes are normalised, usually by taking the logarithms of the parameters or their relative variations. Intuitively, moving along the major axis changes the model by less than the noise, so the two parameters are indistinguishable, while moving along the minor axis changes it in a way the data can detect.

4.2 Isothermal case

We now use the method developed in Section 4.1 to characterise the degeneracies of our model, starting from the isothermal case, which is the simplest one.

We keep the emission surface fixed, because what we want here are the degeneracies which belong to the model. This means that its parameters are not free in Fisher's matrix, while their values span the intervals measured for the ^{12}CO and the ^{13}CO of the MAPS discs [10]. This is an optimistic choice, since the height of the emitting layer, described in Section 2.2.1, is an input of the model and carries an uncertainty of its own, which for the faintest lines can reach half of the height itself [9] and which enters the rotation curve through the term of order $(z/R)^2$. Its effect on the recovery of the parameters has been studied with synthetic observations by Andrews et al. [14] and we come back to it in Chapter 5.

We start by choosing intervals for our parameters in order to thoroughly study the whole space:

We will sample the space by taking three values in the interval of every parameter of Table 4.1 and of the four parameters of the emission surface, the two ends and the centre, which gives 3^9 combinations. Then we consider an ideal rotation curve of 24 independent points, going from 50 au to 600 au with a beam of 0.15″ at a distance of 160 pc[3]. At this stage we will be using only one tracer, because in this way we isolate the degeneracies which belong to the model itself. We will study the effects of adding tracers in Section 4.2.2. We take as errors on every point the ones we obtain by running the experimental pipeline of Section 3.2.2, which contain only the statistical component and therefore make our estimate an optimistic one. We then apply the method under these conditions by following these steps:

1. For every one of the 3^9 combinations we compute v_ϕ on the whole radial grid.

[3] A typical distance for a disc observed by ALMA.

Parameter	Symbol	Interval
Stellar mass	M_* [$M_\odot$]	0.47 – 2.03
Disc mass	M_d [$M_\odot$]	0.04 – 0.16
Scale radius	R_c [au]	90 – 370
Temperature at 100 au	$T_{\mathrm{mid},100}$ [K]	13 – 32
Temperature slope	q_{mid}	−0.37 – −0.08

Table 4.1: Intervals of the parameters of the isothermal model. The masses and the scale radius cover the discs of the exoALMA [18] and MAPS [10] samples bracketing the values they measure. The two masses are sampled independently, so that part of the space falls below the threshold $M_d/M_* \approx 0.05$ under which the disc mass cannot be measured [13, 14], which Section 5.2.3 uses as a test, while the midplane temperature and its slope cover the ten exoALMA discs, whose thermal structures are measured by Galloway-Sprietsma et al. [37].

2. We compute the derivative of v_ϕ at every R with respect to every parameter θ with the central difference method, moving one parameter at a time by a small fraction of its interval.
3. We divide each derivative by the error at that point, obtaining the vector D_θ.
4. We build F with the scalar products of the D vectors, we check its condition number and we invert it.
5. We extract the marginalised errors from the diagonal of the inverse, the conditional ones from the diagonal of F, and the correlations by normalising the off diagonal elements of the inverse [47].
6. We repeat the procedure on every combination and we report the median and the percentiles over the whole sample, since the result is a distribution over the parameter space and not a single number.
7. For the pairs of parameters with the strongest correlation we draw the ellipses of Equation 4.7, using the corresponding block of the inverse matrix and $\Delta\chi^2 = 2.30$.

4.2.1 Degeneracies

Thanks to this procedure we were able to identify the strongest intrinsic degeneracies of the model. These are presented in Figure 4.1.

In the first panel of Figure 4.1 we can see the stellar mass and temperature degeneracy. Intuitively this degeneracy comes from the fact that increasing the temperature, through the pressure gradient, tends to slow down the rotation while increasing the stellar mass has the opposite effect. The link between the thermal structure and the stellar mass has been found also by Martire et al. [10], whose isothermal fits underestimate the stellar mass of the MAPS discs, and Longarini et al. [18] identify the thermal structure as the dominant systematic uncertainty on the dynamical masses, since it sets the pressure gradient. In the second panel we can see the degeneracy

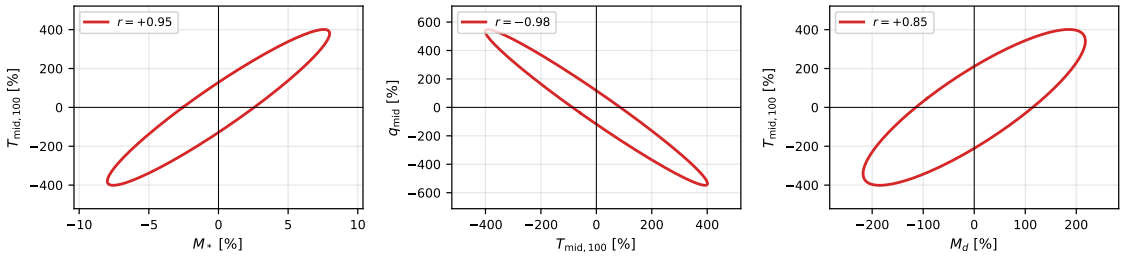


Figure 4.1: Confidence ellipses at the 68% level for the three pairs of parameters with the strongest correlation, drawn in relative variations of the parameters. Each ellipse is built with the median of the two marginal errors and the median of their correlation over the whole sampled space, therefore it summarises the degeneracy of the model and it does not describe a single disc. The two axes of each panel have different scales.

between the temperature at 100 au and the slope which is due to the power-law form of the temperature function. The third panel shows the disc mass and temperature degeneracy which is the result of the fact that both parameters act on the external part of the rotation curve in the same way. The same degeneracy between pressure support and self-gravity is described by Andrews et al. [14], who find it stronger for discs with a gradual outer taper, and Martire et al. [10] observe that an isothermal fit compensates for an inadequate thermal structure by increasing the disc mass. Other relevant degeneracies are:

- The one between the stellar mass and the temperature slope, with $r = -0.95$. It is the degeneracy of the first panel seen through the second one, since the slope changes the temperature over the observed range and therefore the pressure support, which the stellar mass compensates.
- The one between the disc mass and the temperature slope, with $r = -0.72$, which works in the same way in the outer part of the disc, where the self-gravity acts.
- The one between the stellar mass and the disc mass. Their direct link, measured by the partial correlation[4], is an anticorrelation with $r = -0.97$, since in the outer part of the disc the curve measures the mass enclosed within the radius and less stellar mass can be compensated by more disc mass [14]. With one tracer the marginal correlation is +0.42, because both masses are dragged by their degeneracy with the temperature, as we will see in Section 4.2.2.
- The one between the disc mass and the scale radius, with $r = +0.58$. Taken alone the two parameters compensate each other, since a heavier disc rotates faster because of its self-gravity and a larger scale radius rotates faster because the pressure term becomes smaller, and their partial correlation is −0.80. The positive marginal correlation comes from the degeneracy that both parameters share with the temperature. According to Andrews et al. [14] the degeneracies

[4]The correlation computed from Fisher's matrix itself as $-F_{ij}/\sqrt{F_{ii}F_{jj}}$, which describes the link between two parameters when all the other ones are fixed.

between the parameters of the surface density profile are also the main reason why the disc mass cannot be measured below $M_d/M_* \approx 5\%$.

We also report our results for the conditional and marginal errors in Table 4.2. In this table we can see that the conditional errors are all of the order of a few per cent or smaller which means that if the exact values for all of the other parameters were known, the parameter would be recoverable with an error of only a few per cent. On the other hand the marginal errors show that if the other parameters were free the uncertainties explode. This is the consequence of the degeneracies.

Parameter	Marginal error [%]	Conditional error [%]	Ratio
M_*	5.3 [1.7, 16]	0.11 [0.09, 0.18]	47 [15, 129]
M_d	144 [10, 1100]	1.5 [0.8, 3.5]	85 [9, 483]
R_c	119 [9, 664]	2.4 [0.8, 5.7]	59 [6, 235]
$T_{\mathrm{mid},100}$	264 [53, 1120]	1.0 [0.5, 1.9]	317 [52, 902]
q_{mid}	361 [92, 2250]	3.9 [1.7, 9.0]	95 [32, 475]

Table 4.2: Relative marginal and conditional errors of the parameters of the isothermal model and their ratio, for one tracer. For each quantity we report the median over the 3^9 combinations of the sampled space and, in brackets, the 16th and 84th percentiles. The ratio is the median of the ratios and not the ratio of the two medians.

4.2.2 Number of tracers

In this section we quantify the improvement obtained by increasing the number of tracers used in the analysis. Assuming that the errors on different tracers are independent, the F matrix is simply the sum of the F matrices of the single tracers. This is because the probability of observing two curves is the product of the probabilities of the single curves and, taking the logarithm of this product, the chi squared of the two curves becomes the sum of their chi squared. Therefore, by using n tracers with a similar sensitivity to the parameters, the conditional errors are reduced by a factor of about $\sqrt{n}$. Moreover each tracer emits from a different height and adds information about the rotation at that height, which can weaken some of the degeneracies of Section 4.2.1.

In order to measure this second effect we repeat the analysis of Section 4.2.1 with up to five tracers. The first one has the emission surface sampled as before, while the other ones are copies of the same surface lowered by factors evenly spaced between one and 0.6, which is the median ratio between the heights at 100 au of the ^{13}CO and ^{12}CO surfaces of the MAPS discs [10]. The results are reported in Table 4.3.

The conditional errors follow the statistical expectation, for example the one on the temperature goes from 1.0% with one tracer to 0.72% with two and to 0.46% with five. The marginal errors improve much more with the second tracer, by a factor between three and six, while the following tracers give a smaller improvement, even slower than $\sqrt{n}$, since they are placed between the heights already sampled. What matters is therefore the separation between the heights of the tracers more than their number.

n	M_*	M_d	R_c	$T_{\mathrm{mid},100}$	q_{mid}	$r_{T_{\mathrm{mid},100}\,q_{\mathrm{mid}}}$	$r_{M_*\,T_{\mathrm{mid},100}}$	$r_{M_d\,T_{\mathrm{mid},100}}$	$r_{M_*\,M_d}$
1	5.3	144	119	264	361	−0.98	+0.95	+0.85	+0.42
2	1.4	29	34	46	91	−0.87	+0.84	+0.67	0.00
3	1.2	26	30	42	82	−0.88	+0.85	+0.68	+0.04
4	1.1	23	27	38	74	−0.89	+0.85	+0.69	+0.05
5	1.0	21	25	35	67	−0.89	+0.85	+0.70	+0.06

Table 4.3: Relative marginal errors, in per cent, and marginal correlations of the parameters of the isothermal model as a function of the number of tracers n. Every value is the median over the 3^9 combinations of the sampled space. The first row corresponds to Table 4.2.

The reason is the way the height enters the isothermal rotation curve. In Equation 2.63 the pressure term does not depend on z and, apart from the self-gravity, the height appears only in the term multiplied by q_{mid}, which grows approximately as $(z/R)^2/2$. For $q_{\mathrm{mid}} = 0$ the only dependence on the height left would be the one of the self-gravity. The difference between two curves at different heights is therefore a direct measurement of the slope. However this difference is small: for surfaces at $z/R = 0.3$ and 0.18 it is about $0.03\,|q_{\mathrm{mid}}|$ in units of v_k^2, which for $q_{\mathrm{mid}} = -0.23$ changes the velocity by 0.3%, the same order as the errors of the extracted curve. The degeneracy between $T_{\mathrm{mid},100}$ and q_{mid} is therefore only weakened, with a correlation going from −0.98 to −0.87, and the one between M_* and $T_{\mathrm{mid},100}$ survives almost unchanged.

The disc mass improves as well, from 144% to 29% with two tracers. This improvement comes from the weakening of the degeneracy with the temperature, whose correlation with the disc mass goes from +0.85 to +0.67. For the same reason the marginal correlation between the two masses drops from +0.42 to about zero, while their partial correlation stays at −0.97 for any number of tracers. Combining the kinematics of several lines has been proposed by Andrews et al. [14] as a way to reduce the degeneracies of the problem, and multi-molecule fits of the rotation curve have been performed by Pezzotta et al. [9]. Even with five tracers, however, the correlations between $T_{\mathrm{mid},100}$ and q_{mid}, between M_* and $T_{\mathrm{mid},100}$ and between R_c and q_{mid} are still −0.89, +0.85 and +0.66, and the temperature is recovered with a median error of 35%.

These results rely on three simplifications. First, we assign the same errors to every tracer, while the lines of the rarer isotopologues, which emit from deeper layers, are fainter and noisier. Second, the independence of the errors of different tracers holds for the thermal noise, but not for the systematic errors of the extraction, which can be common to all the curves. Third, the surfaces of the additional tracers are not measured but built by lowering the first one by factors evenly spaced between one and 0.6, as described at the beginning of this section, therefore their separation in height is set by the ratio observed between the ^{13}CO and ^{12}CO surfaces.

4.2.3 Radial range

In this section we discuss the effects of setting a radial range outside of which we discard the data points for the fits. This analysis is needed because, as seen in Section 3.2, the amplitude of the uncertainties and the presence of a measuring bias depend on the radial distance. As we will see however different parameters are better constrained by certain radial regions of the disc. The balance between these two concepts is what determines the chosen radial range.

In order to carry out this analysis we used the usual radial grid, parameter space and errors. Then we tried changing the external radius from 600 au to 400 au, removing one beam size at a time, with the internal one fixed at 50 au, since this is the range where the external cuts usually fall. Afterwards we did the same for the internal radius changing it from 1 to 6 beam sizes with the external one fixed at 600 au. We fitted the relative increase of the errors we got from this analysis with a power law of the radius, $(R_{\rm out}/600\,{\rm au})^{\alpha}$ for the external cuts and $(R_{\rm in}/50\,{\rm au})^{\alpha}$ for the internal ones, to quantify how much each parameter depends on the cuts. The exponents α, negative for the external cuts and positive for the internal ones, are reported in the legends of Figures 4.2 and 4.3, which show the results for the external and the internal radius.

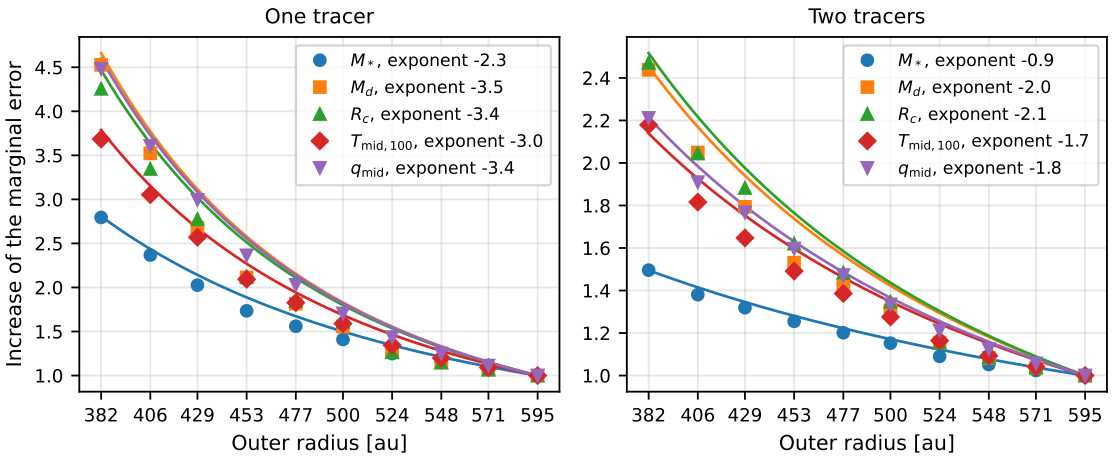


Figure 4.2: Increase of the marginal error of each parameter as a function of the outer radius, with respect to the curve extending to 600 au, for one tracer on the left and two tracers on the right. The points are medians over the sampled space and the lines are the fitted power laws, whose exponents are reported in the legend. The spread between the 16th and 84th percentiles is about a factor of two.

It is evident from the graphs in this section that the uncertainties of the parameters depend on different radial regions. The stellar mass for example is constrained mostly by the inner radial data while the scale radius and disc mass need data from the external region. For example, cutting the curve at 500 au increases the marginal errors by a factor between 1.4 and 1.7 with one tracer and between 1.2 and 1.3 with two, while starting it at 100 au increases the error on the stellar mass by a factor of 2.6 with one tracer and 1.7 with two. These factors describe only the statistical errors, while a cut can also remove the part of the curve where the extraction is less reliable.

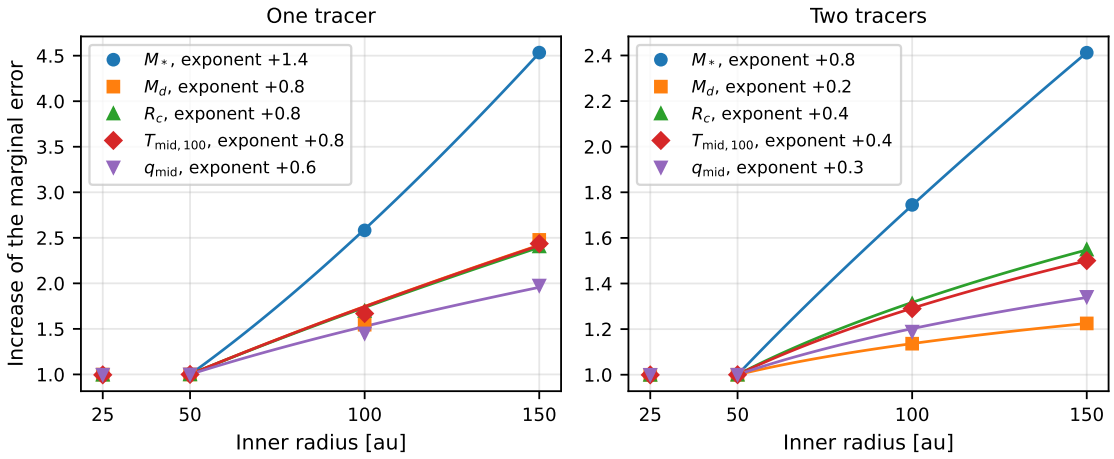


Figure 4.3: Same as Figure 4.2 as a function of the inner radius, with respect to the curve starting at 50 au. The point at 25 au coincides with the one at 50 au because the error on the innermost point is too large to add information.

It is worth comparing these radial ranges with the extent of real curves. For LkCa 15, the disc used as an example in Chapter 3, the emission surface of the ^{12}CO stops rising at 488 au, while that of the ^{13}CO already stops at 222 au, from the parameters of Table 3.2. The outer cuts explored in this section therefore cover the range within which a real analysis has to choose, and for the second tracer that choice is often made by the data rather than by us.

4.3 Stratified case

We now apply the same procedure to the stratified model, with the same emission surfaces, radial grid and errors of Section 4.2. The intervals of the parameters are listed in Table 4.4. With nine parameters of the model and four of the emission surface, a grid with three values for each interval would contain more than a million combinations, therefore we sample the space with a Latin hypercube[5] of 3^9 points, the same number used for the isothermal case. The logarithm of $f\,g$ is computed with Equations 2.67 and 2.68.

Fitting the same curve with the stratified model adds four thermal parameters without adding any data. We therefore study it first with the atmosphere fixed and then with the atmosphere free, knowing that adding free parameters to Fisher's matrix can only increase the marginal errors. We keep the atmospheric parameters fixed because they can be measured from the brightness temperature of optically thick lines [15, 37], while the midplane ones cannot.

[5]A sampling in which the interval of every parameter is divided into N equal parts and each part is used exactly once, so that the points cover every interval evenly without filling the whole grid.

Parameter	Symbol	Interval
Stellar mass	M_* [$M_\odot$]	0.47 – 2.03
Disc mass	M_d [$M_\odot$]	0.04 – 0.16
Scale radius	R_c [au]	90 – 370
Midplane temperature at 100 au	$T_{\mathrm{mid},100}$ [K]	13 – 32
Midplane temperature slope	q_{mid}	−0.37 – −0.08
Atmospheric temperature at 100 au	$T_{\mathrm{atm},100}$ [K]	34 – 67
Atmospheric temperature slope	q_{atm}	−0.87 – 0
Transition height at 100 au	Z_0 [au]	10 – 87
Transition height slope	β	0 – 1.21

Table 4.4: Intervals of the parameters of the stratified model. The first five rows are the ones of Table 4.1, while the atmospheric temperature and the transition height cover the two-dimensional temperature structures of the ten exoALMA discs measured by Galloway-Sprietsma et al. [37] and reported by Longarini et al. [18]. The transition heights, given in arcseconds, are converted to au with the distances of Teague et al. [39].

4.3.1 Degeneracies

The marginal and conditional errors are reported in Table 4.5, while Figure 4.4 shows the ellipses of the same pairs of parameters of Figure 4.1.

Parameter	Marginal error		Conditional error	Ratio	
	fixed	free		fixed	free
M_*	4.7	100	0.11	42	900
M_d	68	5800	1.5	40	3600
R_c	39	4500	2.5	15	2100
$T_{\mathrm{mid},100}$	166	8000	1.7	79	3600
q_{mid}	447	17600	5.9	61	2300
$T_{\mathrm{atm},100}$	–	12900	2.5	–	5000
q_{atm}	–	61	0.020	–	3000
Z_0	–	7700	7.4	–	1200
β	–	47	0.073	–	650

Table 4.5: Marginal and conditional errors of the parameters of the stratified model and their ratio, for one tracer, with the atmosphere fixed and free. The errors are relative and in per cent, except for q_{atm} and β, whose intervals contain zero and whose errors are absolute. Every value is the median over the Latin hypercube, and the conditional errors do not depend on which parameters are free.

With the atmosphere fixed the conditional errors of the midplane parameters are larger than in the isothermal case, 1.7% against 1.0% for $T_{\mathrm{mid},100}$ and 5.9% against 3.9% for q_{mid}. The emitting layer sits in the warmer part of the disc, where the temperature is set mostly by the atmosphere, and the midplane parameters act on the curve only through H_{mid} and the shape of f. The degeneracies are however weaker, with a ratio between marginal and conditional error of 79 for the temperature against 317, and

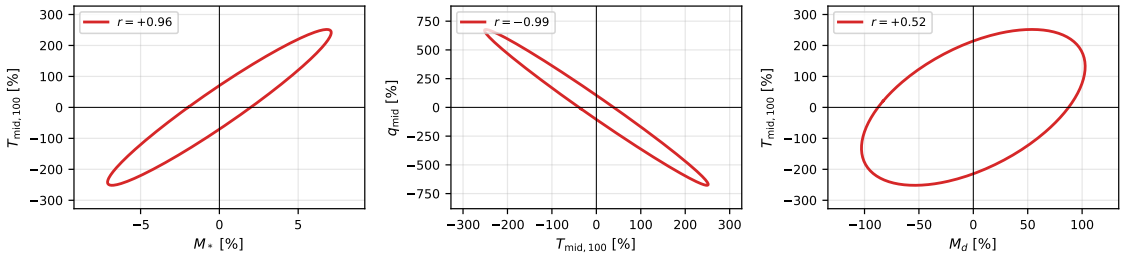


Figure 4.4: Same as Figure 4.1 for the stratified model with the atmosphere fixed.

the marginal errors of the disc mass, of the scale radius and of the temperature are lower than in the isothermal case, with 68%, 39% and 166% against 144%, 119% and 264%. In Equation 2.62 the term which contains the scale radius is multiplied by f, which is larger than one in the atmosphere, and this helps to separate R_c from the other parameters. The slope of the midplane temperature is the only parameter which gets worse, from 361% to 447%, and its degeneracy with the temperature becomes the strongest of the model, with $r = -0.99$. The other pairs of Figure 4.4 keep the sign they have in the isothermal case, while the correlation between the disc mass and the scale radius changes sign, from +0.58 to −0.69. In both models their partial correlation is negative, −0.80 and −0.82, which shows that the marginal correlation of these two parameters is set by the other degeneracies.

Leaving the atmosphere free changes the picture completely. The conditional errors do not change, since they do not depend on the other parameters, but the marginal ones grow by factors between 20 and 120: the stellar mass goes from 4.7% to 100%, the midplane temperature from 166% to 8000%, and none of the other parameters is constrained. Four new degeneracies appear, which involve the atmosphere, three of which are shown in Figure 4.5: between Z_0 and β, with $r = -0.94$, between $T_{\mathrm{atm},100}$ and Z_0, with $r = +0.94$, between $T_{\mathrm{atm},100}$ and q_{atm}, with $r = -0.87$, and between $T_{\mathrm{mid},100}$ and $T_{\mathrm{atm},100}$, with $r = +0.83$. The first and the third have the same origin as the one between $T_{\mathrm{mid},100}$ and q_{mid}, since in both cases a normalisation at 100 au and a slope describe the same radial profile. The second one comes from the fact that raising the transition height brings colder gas into the emitting layer, and a hotter atmosphere compensates for it.

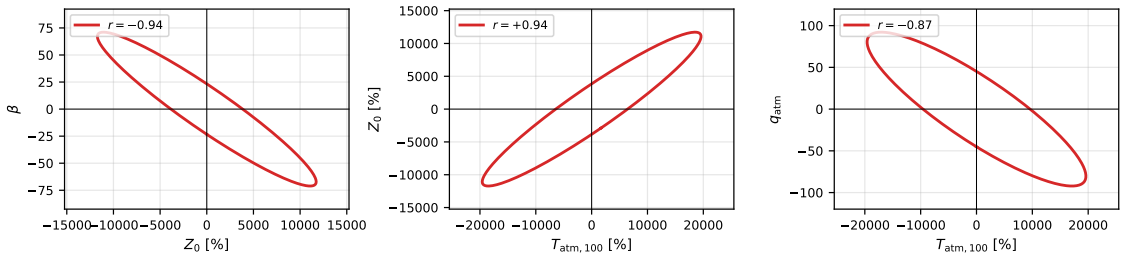


Figure 4.5: Confidence ellipses at the 68% level for three pairs of atmospheric parameters of the stratified model with the atmosphere free, built as in Figure 4.1. The variations are relative for $T_{\mathrm{atm},100}$ and Z_0 and absolute for q_{atm} and β, whose intervals contain zero.

4.3.2 Number of tracers

We repeat the analysis of Section 4.2.2 with the same prescription for the additional surfaces. The results with the atmosphere fixed are reported in Table 4.6.

n	M_*	M_d	R_c	$T_{\mathrm{mid},100}$	q_{mid}	$r_{T_{\mathrm{mid},100}\,q_{\mathrm{mid}}}$	$r_{M_*\,T_{\mathrm{mid},100}}$	$r_{M_d\,T_{\mathrm{mid},100}}$	$r_{M_*\,M_d}$
1	4.7	68	39	166	447	−0.99	+0.96	+0.52	−0.22
2	1.2	20	12	49	125	−0.99	+0.85	+0.68	−0.18
3	1.0	17	10	44	112	−0.99	+0.87	+0.69	−0.11
4	0.94	16	9.2	40	102	−0.99	+0.88	+0.70	−0.09
5	0.86	14	8.4	37	94	−0.99	+0.88	+0.70	−0.07

Table 4.6: Same as Table 4.3 for the stratified model with the atmosphere fixed. Every value is the median over the Latin hypercube.

The second tracer reduces the marginal errors by a factor between three and four, similarly to the isothermal case, and the following ones give a smaller improvement. With five tracers the midplane temperature is recovered with a median error of 37%, close to the isothermal 35%, while the scale radius reaches 8% against 25%. The main difference is that the degeneracy between $T_{\mathrm{mid},100}$ and q_{mid} is not weakened at all, with $r = -0.99$ for any number of tracers. This probably happens because in the stratified model the dependence of the curve on the height is dominated by the transition between the midplane and the atmosphere, therefore the difference between two tracers carries information on f more than on the slope of the midplane temperature. With the atmosphere free the second tracer reduces the errors by factors between about 20 and 50, but only the stellar mass becomes constrained, with 3.8% for two tracers and 2.4% for five. The midplane and atmospheric temperatures remain above 150% even with five tracers, and the correlation between $T_{\mathrm{atm},100}$ and Z_0 grows to +0.98. The simplifications listed in Section 4.2.2 hold also for these results.

4.3.3 Radial range

With the same cuts of Section 4.2.3 and the atmosphere fixed we obtain Figures 4.6 and 4.7. With one tracer the exponents lie between −2.5 and −3.6 for the outer radius and between +0.8 and +1.3 for the inner one, close to the isothermal ones, therefore the predictions for the radial cuts of Section 4.2.3 hold for both models. The scale radius becomes the parameter most sensitive to the outer radius, with an exponent of −3.6. With the atmosphere free the dependence is much steeper, with exponents between −5.7 and −8.1 for the outer radius with one tracer and between −1.9 and −3.4 with two, therefore every beam removed from the outer disc costs much more than in the isothermal case.

4.4 A simple test

In the last section of this chapter we verify the results of Fisher's analysis. We do this by computing a rotation curve from our model, adding Gaussian noise and then

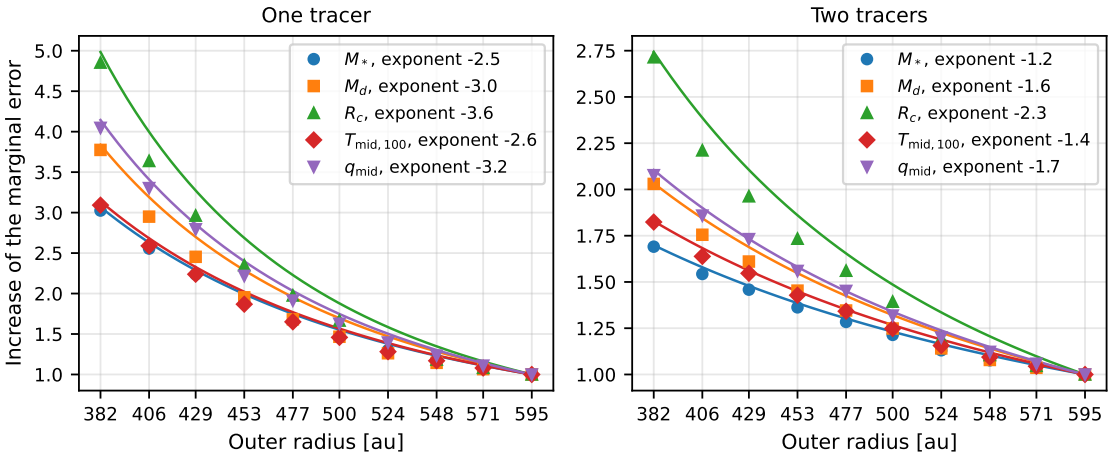


Figure 4.6: Same as Figure 4.2 for the stratified model with the atmosphere fixed.

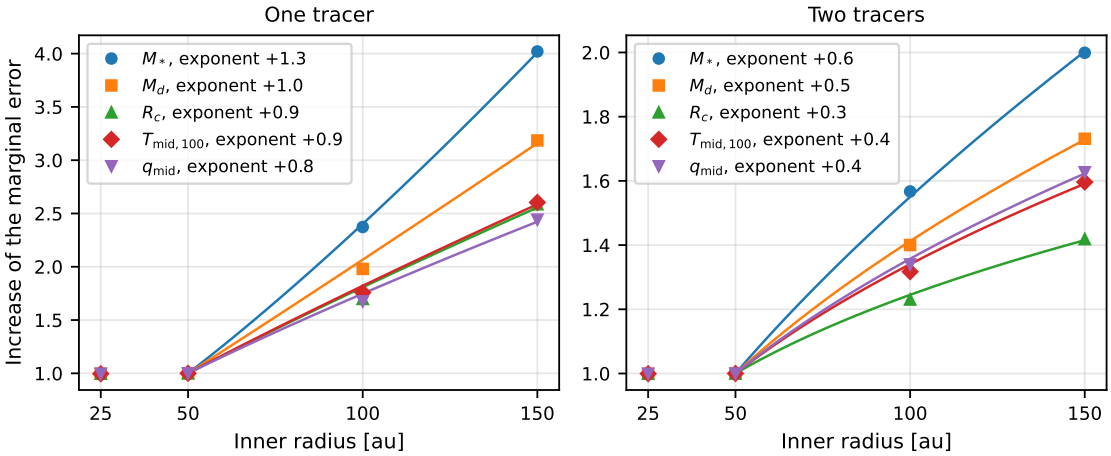


Figure 4.7: Same as Figure 4.3 for the stratified model with the atmosphere fixed.

fitting it with an MCMC. If the analysis is correct, the results of the fit should show the same degeneracies and errors. We use the usual radial grid and the error profile of Section 4.2, both for the noise added to the curve and in the likelihood, and we test the disc of LkCa 15, the same one used as an example in Chapter 3, whose parameters are listed in Table 3.2. The error profile is the one estimated by Discminer on the curves extracted by our pipeline, but the noise we add is Gaussian and independent from point to point, which is exactly the hypothesis on which Fisher's analysis is built. The test therefore shows the behaviour of a fit whose errors satisfy these hypotheses, while the residuals of the curves actually extracted by the pipeline are correlated and contain a systematic bias, as discussed at the end of this section. We test both the isothermal and the stratified case, using the ^{12}CO and ^{13}CO molecules. In both cases the free parameters are M_*, M_d, R_c, $T_{\mathrm{mid},100}$ and q_{mid}, with uniform priors, while in the stratified case the atmosphere is kept fixed. The fit is performed with 64 walkers, 3000 burn-in steps and 30000 production steps.

Figures 4.8 and 4.9 show the chains of the walkers. After the burn-in the chains are stationary and the integrated autocorrelation time is at most 271 steps in the isothermal case and 156 steps in the stratified one, therefore the production phase covers at least 111 autocorrelation times.

The results are shown in Figures 4.10 and 4.11, where we overlay the prediction of Fisher's analysis computed at the parameters of Table 3.2, and they are summarised in Table 4.7. All of the parameter values are recovered inside an interval of 1.7σ in the isothermal case and of 0.3σ in the stratified one, which shows that with Gaussian noise alone and two tracers every parameter of the model is recovered within its errors. The errors of the fit and the ones predicted by Fisher's analysis agree within 10 per cent in the isothermal case and within 11 per cent in the stratified one. The correlations agree as well where the degeneracies are strong, -0.98 against -0.99 for $T_{\mathrm{mid},100}$ and q_{mid} in the stratified case and -0.78 against -0.80 between the two masses. In the isothermal case that same degeneracy is broken, -0.15 in the fit against -0.03 predicted, because the two tracers emit from heights which differ by a factor of two and the difference between their curves measures the slope directly. The one quantity on which the two methods disagree is the correlation between the disc mass and the scale radius, $+0.78$ in the fit against $+0.25$ predicted. The errors of this disc differ also from the median ones of Tables 4.3 and 4.6, since the analysis of the previous sections describes the median behaviour over the whole parameter space, while the test concerns a single disc with its own parameters and emission surfaces.

We note that the fit with only one tracer does not work, as foreseen in Sections 4.2.2 and 4.3.2. Repeating the test with ^{12}CO alone, Fisher's analysis at the same parameters predicts an error on the temperature of 32 K in the isothermal case and of 14 K in the stratified one, both larger than the 10.1 K standard deviation of the uniform prior, so the curve of a single tracer carries no information on the temperature at all. The posteriors are indeed narrower than that only because the prior truncates them, with standard deviations of 7.3 K and 5.9 K. In this regime the likelihood is far from Gaussian and the result is set by the prior, therefore the true values fall inside the errors only because the posterior fills the whole prior.

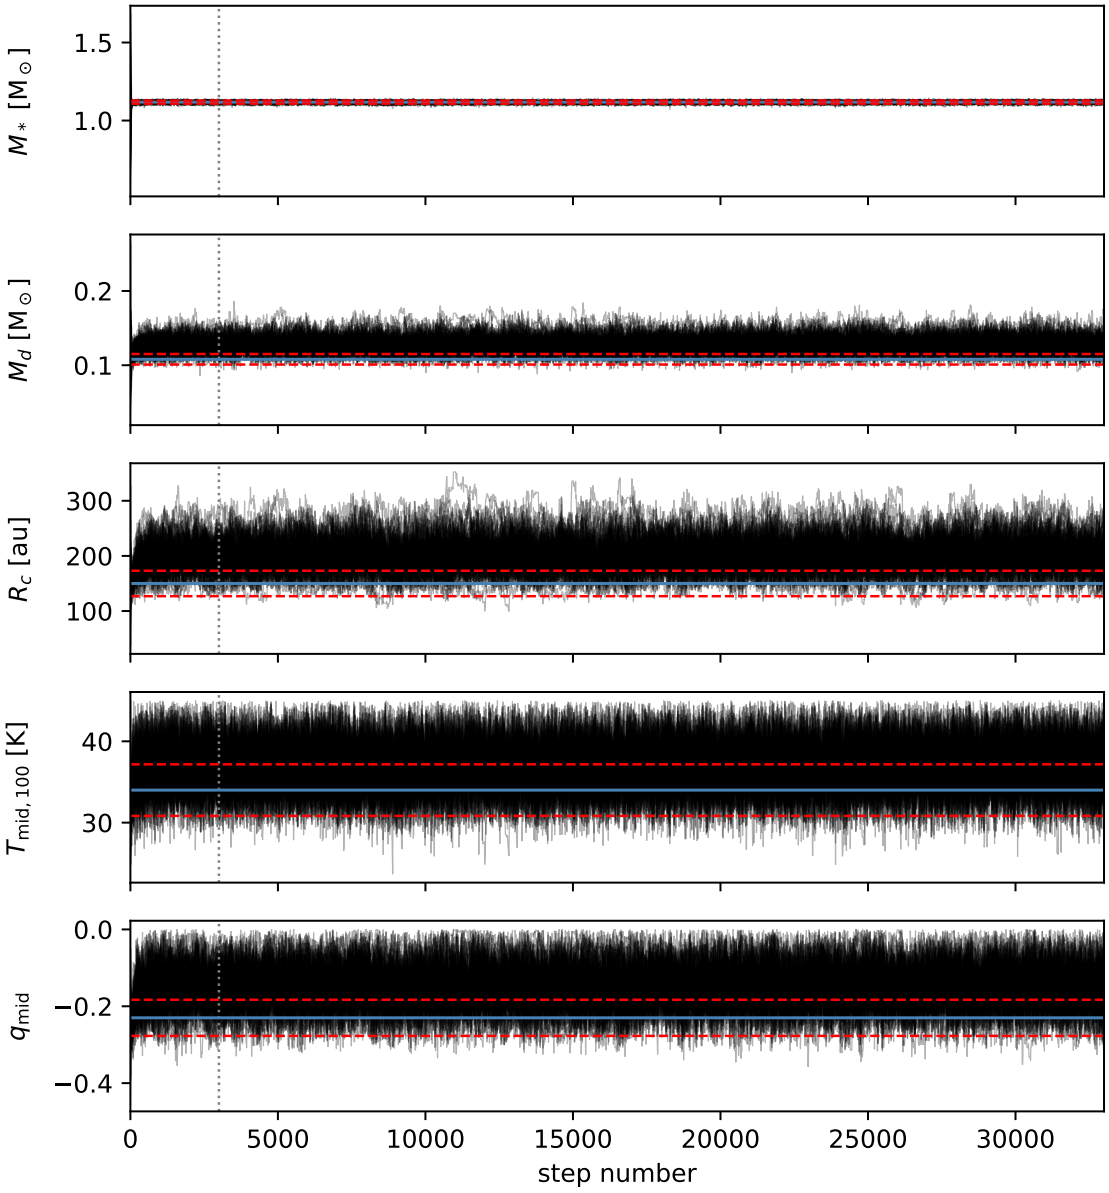


Figure 4.8: Chains of the 64 walkers for the fit of the isothermal synthetic curves of LkCa 15 with ^{12}CO and ^{13}CO. The dotted line marks the end of the burn-in, the blue lines the true values and the red dashed lines the true values plus or minus the errors predicted by Fisher's analysis at the parameters of Table 3.2.

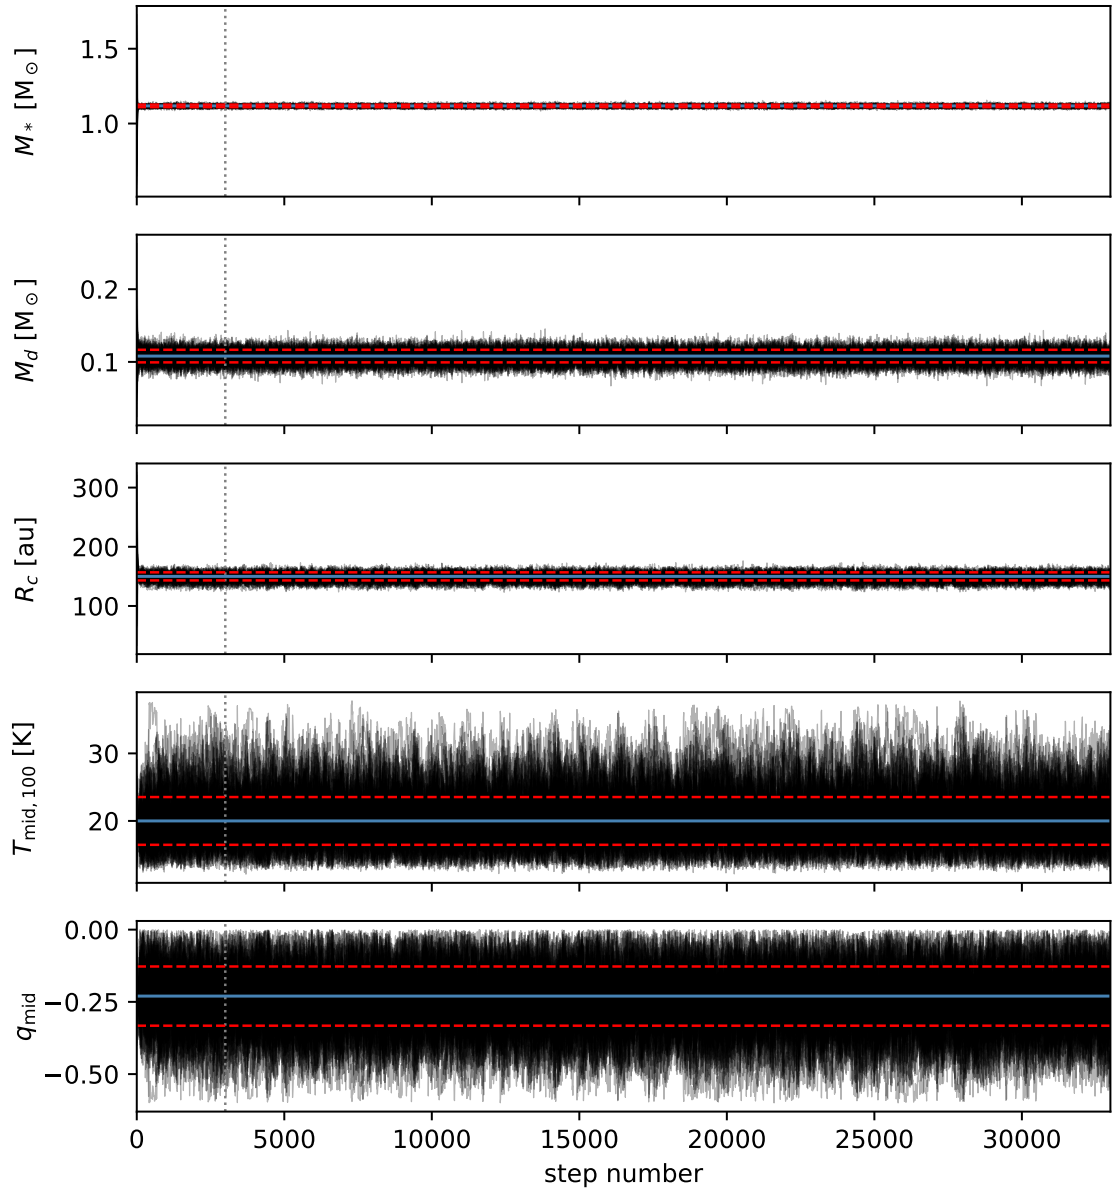


Figure 4.9: Same as Figure 4.8 for the stratified case with the atmosphere fixed.

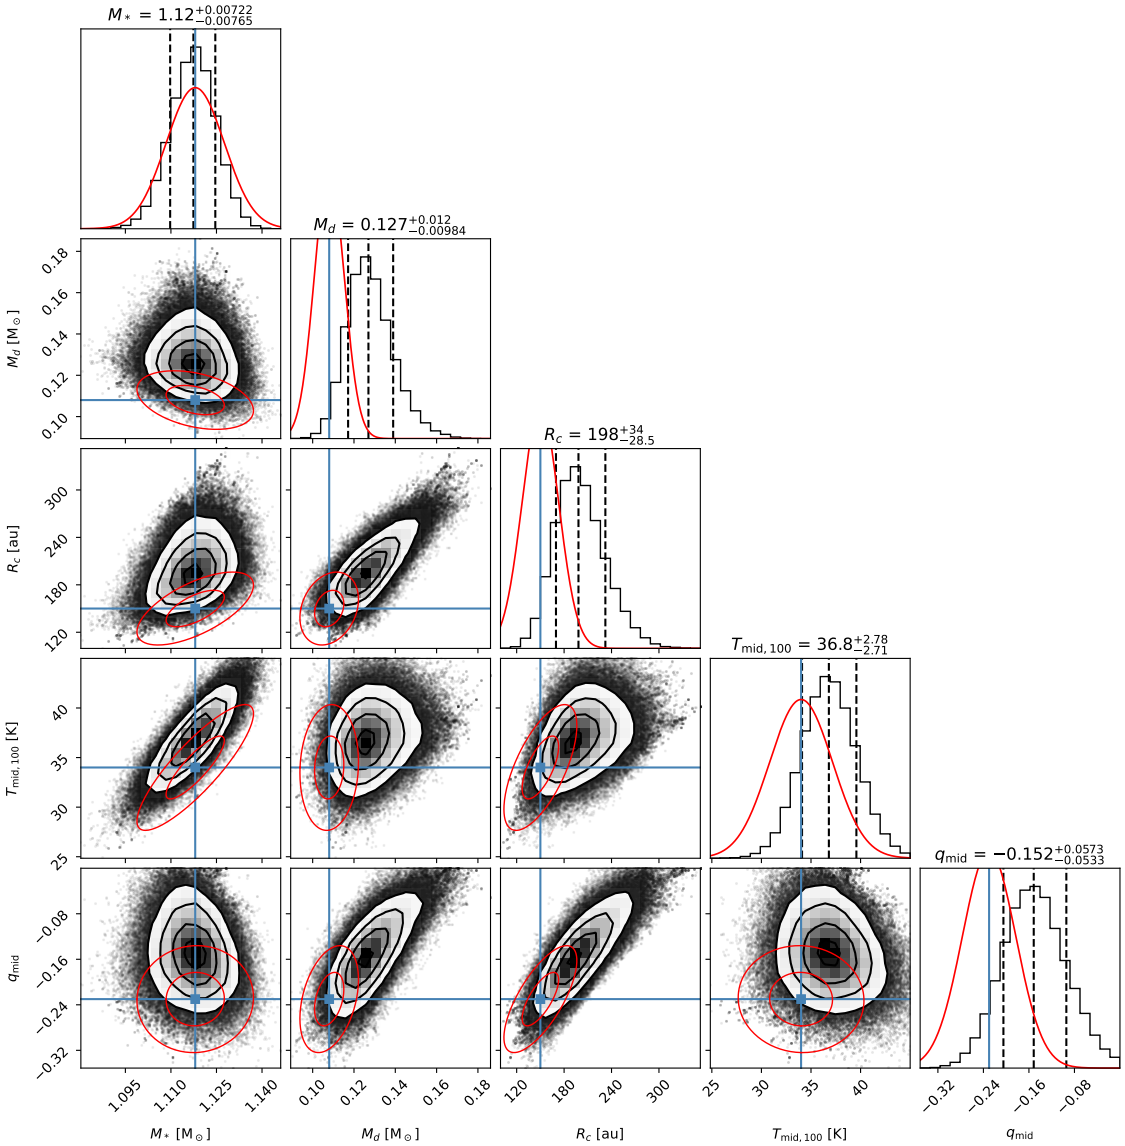


Figure 4.10: Posterior distributions of the fit of the isothermal synthetic curves of LkCa 15 with ^{12}CO and ^{13}CO. The contours are at 0.5, 1, 1.5 and 2σ, the dashed lines mark the 16th, 50th and 84th percentiles and the blue lines the true values. The red ellipses, at 1σ and 2σ, and the red curves on the diagonal show the prediction of Fisher's analysis computed at the parameters of Table 3.2, centred on the true values.

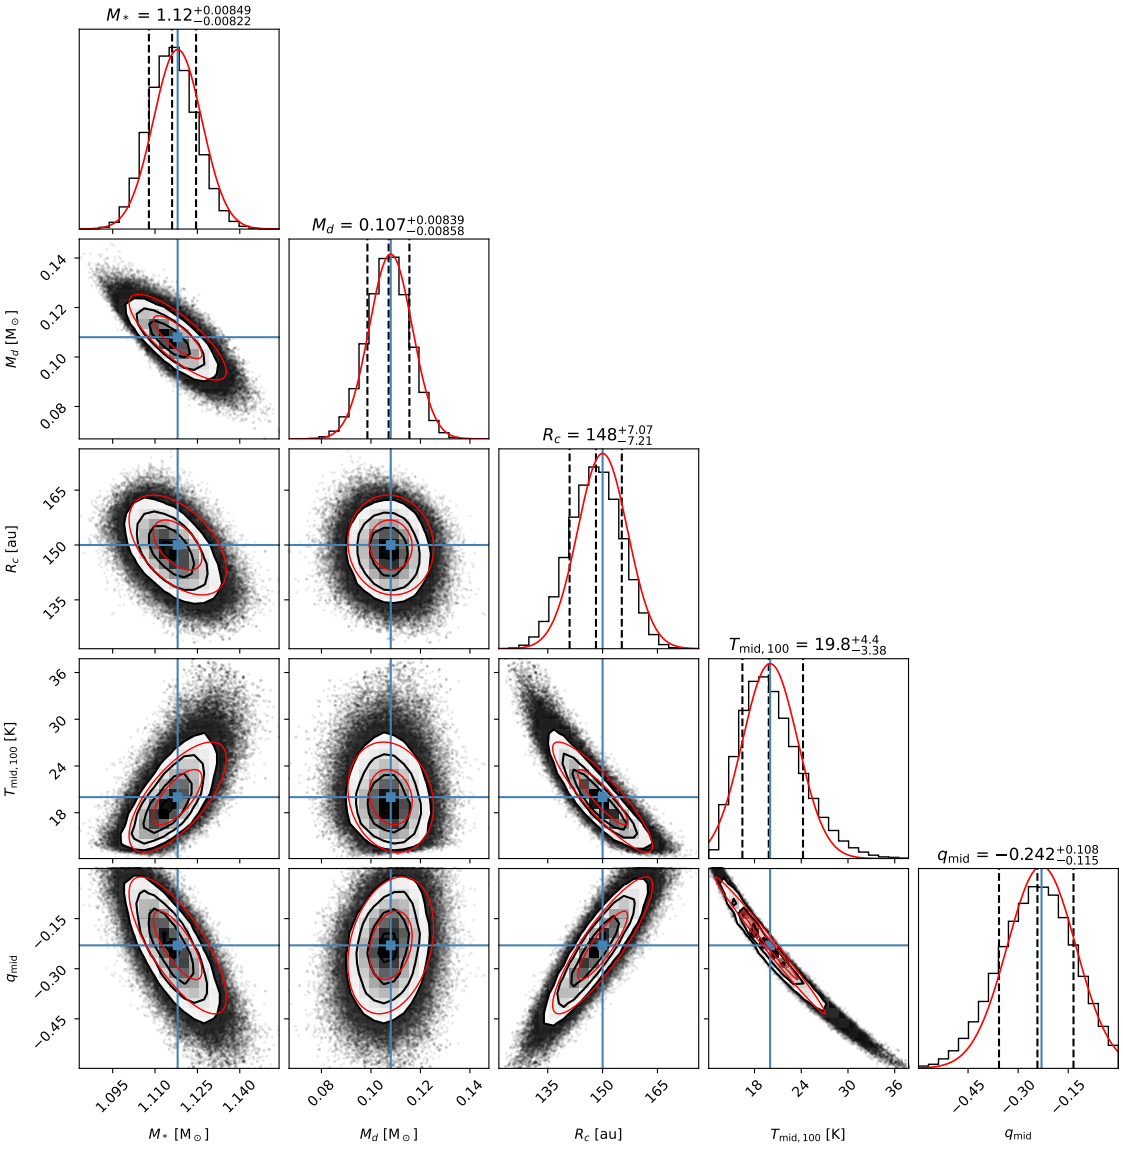


Figure 4.11: Same as Figure 4.10 for the stratified case with the atmosphere fixed.

		Isothermal			Stratified		
Parameter	True	Median	MCMC [%]	Fisher [%]	Median	MCMC [%]	Fisher [%]
M_* [$M_\odot$]	1.118	1.117	0.7	0.7	1.116	0.8	0.8
M_d [$M_\odot$]	0.108	0.127	10	11	0.107	8.0	8.1
R_c [au]	150	198	21	21	148	4.7	4.5
$T_{\mathrm{mid},100}$ [K]	34/20	36.8	8.1	8.2	19.9	20	18
q_{mid}	−0.23	−0.15	24	25	−0.24	47	45

Table 4.7: Results of the fits of the synthetic curves of LkCa 15 with ^{12}CO and ^{13}CO: median of the posterior, relative standard deviation of the posterior and relative error predicted by Fisher's analysis at the parameters of Table 3.2. The midplane temperature is 34 K in the isothermal case, which is the vertical average of the measured structure, and 20 K in the stratified one, where the atmosphere is described separately. The relative errors are given with respect to the true value.

Fisher's analysis gives therefore a reliable description of the ideal case, but it does not contain three effects which are present in the curves extracted from real or synthetic observations. First, the residuals of the extraction are correlated over about 50 au, that is two beams, while we treat every point as independent, which overestimates the information. Second, the amplitude of the errors depends on the extraction: the profile we use has a median of 0.31%, while other extractions of our pipeline give errors between 0.7% and 0.9%, and the marginal errors scale with this amplitude. Third, the errors of Fisher's analysis are statistical, while the extracted curve is also affected by a systematic bias, which does not enter the matrix but moves the recovered parameters. The effect of these limits on the recovery of the thermal parameters is the subject of Chapter 5.

Chapter 5

Results

This chapter combines everything that was written up until now. We take the model introduced in Chapter 2, specifically for the isothermal case because, as seen in Chapter 4, the thermal parameters in the stratified case are intrinsically harder to retrieve. We apply the pipeline introduced in Chapter 3 to generate a sample of synthetic data in order to verify the method under different parametrical conditions. We take the intervals of Table 4.1 as dimensions for our parameter space and we add a further dimension relative to the inclination of the disc from 30 to 65 degrees which covers the MAPS sample. From this parameter space we take 40 random points and build a disc for each of them. Having proved in Chapter 4 that the real lever in retrieving the parameters is adding a second well-distanced in height tracer, we simulate the two most usual molecules[1] for each disc for a total of 80 cubes. To streamline the simulation of data and the extraction of the rotational curve, as explained in Chapter 3, we need to precompute the self-gravity and thermal terms of the model. However, as shown in Chapter 2, the thermal term has a closed-analytical form for the isothermal case. Therefore, we only precompute a grid for the self-gravity term over the intervals of Table 4.1. In the next sections of this Chapter we will study the extraction of the rotational curve and the consequent retrieval of the parameters.

5.1 The extraction bias

We show the results of the extraction of the rotational curve over our sample. Every number in this section is the mean over the sample with an error of one standard deviation. In Table 5.1 we report the residuals[2] for different radial regions. Between 24 and 50 au the beam smearing reduces the velocity but we recover 65 per cent of the residual by applying the beam correction explained in Section 3.2.3. After 50 au the residuals grow linearly and around 565 ± 93 au they switch from negative to positive.

Section 4.2.3 fits the increase of the marginal errors with the outer cut as a power law, and with two tracers the exponents of the thermal parameters are -1.7 for the

[1] The ^{12}CO and the ^{13}CO.

[2] The difference between the extracted data and the generating model.

Region	Extracted [%]	Beam corrected [%]	ξ	Reduction [%]
24 – 50 au	-21.5 ± 10.5	-7.5 ± 8.4	1.20	65
50 – 150 au	-7.3 ± 4.2	-4.6 ± 3.0	1.03	37
150 – 400 au	-4.4 ± 2.5	-3.0 ± 2.1	1.02	32
400 – 550 au	-2.1 ± 3.1	-0.9 ± 3.3	1.01	59
550 – 700 au	$+2.1 \pm 4.6$	$+3.6 \pm 5.0$	1.01	—

Table 5.1: Residual of the extracted rotational curve with respect to the model, before and after the beam smearing correction of Section 3.2.3. Every entry is the mean over the eighty curves of the sample with its standard deviation, ξ is the mean correction factor of the region and the last column is the reduction of the mean, which is not quoted in the outermost region because there the correction moves the mean further from zero. The outer disc is split at 550 au because the residual changes sign inside it, so that a single mean over the whole range would average a negative half against a positive one and return a number close to zero which describes neither. Each column is an independent mean over the eighty curves, so the corrected column is not the product of the extracted one and ξ.

temperature and -1.8 for its slope, so that cutting the curve at 400 au doubles both errors, and extrapolating the same law down to 300 au, below the range over which it was calibrated, would multiply them by about three and a half. The thermal information therefore lives in the outer disc, which is also where the bias of the extraction is smallest in absolute value, -2.1 per cent between 400 and 550 au against -7.3 inside 150, so that how far out we are willing to follow the curve decides how much thermal signal we keep. Further out the residual changes sign and grows again, which is what makes the choice of the outer cut delicate.

This makes the choice of the outer cut a part of the method and it has to be made without knowing the truth. We first cut each tracer where its fitted emission surface stops rising, at $R_t(\psi/q_t)^{1/q_t}$, beyond which the height is an extrapolation. That rule is observable but badly conditioned, because it depends on the tapering exponent of the surface, which is the worst determined parameter of the discminer fit: on our sample it is generated at 6.8 for the ^{12}CO and 4.7 for the ^{13}CO and recovered at 1.5 for both, and the radius which follows from it spreads by a factor of five between the discs of the sample. We therefore cut instead at the first radius where the azimuthally averaged peak brightness of the line falls below 4.5 K, which is three times the noise of the cube for our beam and frequencies, and which is measured rather than extrapolated. With this rule the usable curve of the ^{13}CO reaches 492 au on the median disc instead of 318, and thirty three discs out of forty keep a second tracer beyond 400 au instead of ten.

Inside this chosen external radius, the residual of the extraction is negative, -1.9 per cent on the ^{13}CO, in the first fifty au outside it is $+0.6$ per cent, and further out it grows to $+5.8$ per cent with a scatter four times as large. The cut falls where the bias changes sign, so the points we discard would contribute a systematic of the opposite sign and considerably more noise.

The comparison which decides whether the thermal parameters are within reach is

between this residual bias and the change that the thermal parameters themselves produce on the curve, measured on the same points, at the same heights and over the same window. Averaged between 100 and 600 au, which is the range the fits use, the bias left by the beam smearing correction is −2.0 per cent, while a change of one standard deviation of our grid[3] moves the curve by 1.1 per cent for the midplane temperature and by 0.6 per cent for its slope.

The systematic error left by the extraction is therefore twice the signature of the midplane temperature and more than three times that of its slope. This is the central result of this work. At this level of precision in the extraction of the rotational curve the thermal parameters cannot be retrieved, and the slope is the first of the two to be lost, which is what the fits of the next section show. It is a systematic statement and not an instrumental one, because a longer integration reduces the error bars of the curve but leaves this bias where it is.

5.2 Retrieval of the parameters

In this section we show the results of different attempts at retrieving the thermal parameters together with the scale radius, the mass of the star and the disc.

We call a parameter recovered when five conditions hold together. The true value has to fall inside three effective errors, where the effective error is the larger of the standard deviation of the posterior and five per cent of the true value. The standard deviation of the posterior has to be smaller than fifteen per cent of the width of the prior, otherwise the parameter is not constrained by the data but by the prior. The marginal posterior must not be bimodal and the chain must have converged. Moreover, fewer than twenty per cent of the samples may sit within two per cent of either edge of the prior, otherwise the fit is pressed against a wall and its median is not a measurement. The first two conditions separate two different failures, a value which is wrong and an uncertainty which is not honest. The priors are the intervals of Table 4.1, with one exception: the prior on the scale radius is limited to 400 au, which is the extent of the precomputed grid of the self-gravity term and not a physical choice.

5.2.1 Without any correction

We start from the curves as they come out of the extraction. Table 5.2 reports, for each parameter, the median of the absolute relative deviation from the true value, the median of the effective pull and the number of discs which satisfy all five conditions, over the forty discs of the sample. Figure 5.1 shows which discs those are, in the plane of the inclination and of the true value of each parameter.

Only the stellar mass is recovered in a sizeable fraction of the sample, and even there the median deviation is 16 per cent. The disc mass and the slope are wrong by more than a factor of two, the temperature and the scale radius by about a half, and the

[3] The forty discs sample the temperature at 100 au uniformly between 15 and 32 K, slightly narrower than the interval of Table 4.1 at the cold end, and its slope between −0.37 and −0.08, so that one standard deviation is 4.9 K and 0.084 respectively.

Parameter	Deviation [%]	Pull	Recovered
M_*	16^{+9}_{-11}	$3.2^{+1.6}_{-2.2}$	14/40
M_d	116^{+123}_{-78}	$7.9^{+14.3}_{-5.4}$	8/40
R_c	37^{+31}_{-27}	$3.5^{+5.5}_{-2.6}$	4/40
$T_{\mathrm{mid},100}$	59^{+41}_{-40}	$2.9^{+9.1}_{-2.2}$	3/40
q_{mid}	159^{+193}_{-86}	$14.1^{+24.5}_{-7.6}$	0/40

Table 5.2: Retrieval of the parameters from the extracted curves, without any correction. The deviation is the relative distance of the median of the posterior from the true value, in absolute value, and the pull is the same distance in units of the effective error. Each entry is the median over the forty discs and the two numbers which follow it are the distances to the sixteenth and to the eighty fourth percentile of the sample, so that they enclose 68 per cent of the discs and can be read as the one sigma interval of a Gaussian would be. They are quoted instead of a standard deviation because these quantities are positive and strongly skewed, as in Chapter 4. The last column counts the discs which satisfy the five conditions of recovery.

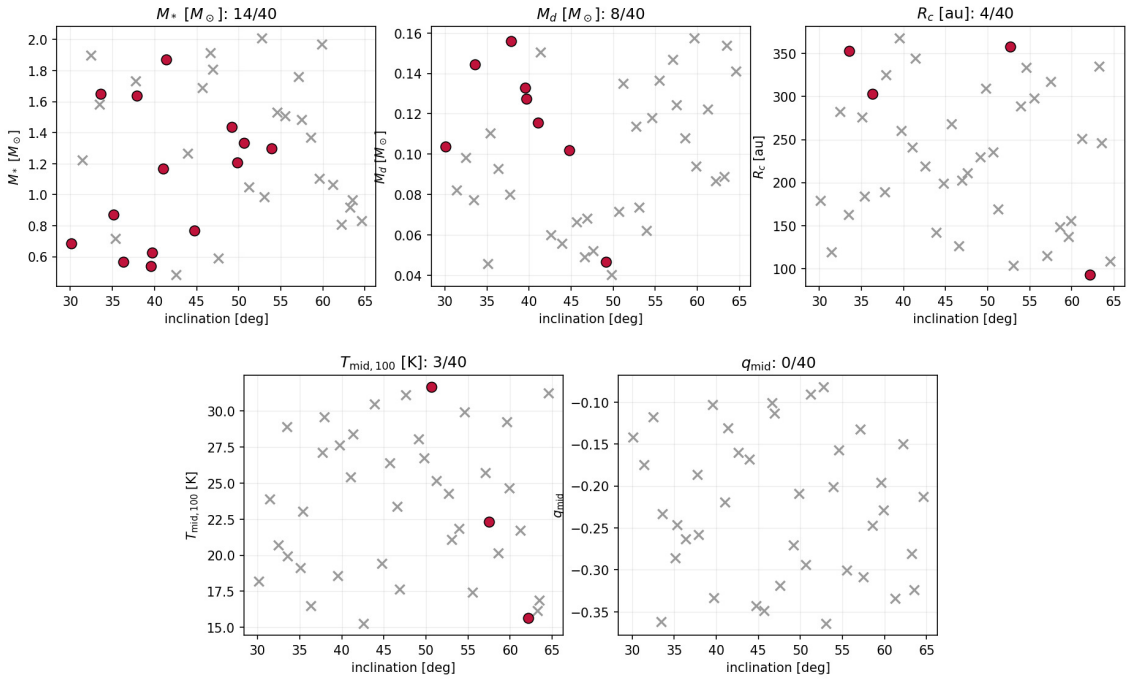


Figure 5.1: Discs in which each parameter is recovered, without any correction, as a function of the inclination and of the true value of the parameter. Filled circles are the recovered ones and crosses the others.

pulls run from 3 to 14, which means that the uncertainties the fits report are far too small, for the two reasons we come back to in Section 5.2.3. The slope is recovered in no disc at all. Nothing is retrieved in the sense in which Chapter 4 retrieved everything on ideal data, and the difference between the two chapters is entirely the extraction.

5.2.2 With the beam smearing correction

We now repeat the fits on the curves corrected as described in Section 3.2.3, which also lets us start the fit at one beam instead of two and therefore adds the innermost points.

Parameter	Deviation [%]	Pull	Recovered
M_*	11^{+9}_{-7}	$2.0^{+2.0}_{-1.2}$	23/40
M_d	73^{+133}_{-52}	$2.8^{+5.8}_{-1.6}$	15/40
R_c	57^{+65}_{-37}	$3.6^{+7.0}_{-1.8}$	4/40
$T_{\mathrm{mid},100}$	54^{+45}_{-20}	$3.7^{+7.5}_{-3.1}$	4/40
q_{mid}	149^{+196}_{-78}	$14.6^{+22.9}_{-10.7}$	3/40

Table 5.3: As Table 5.2, for the curves corrected for the beam smearing and extended inwards to one beam.

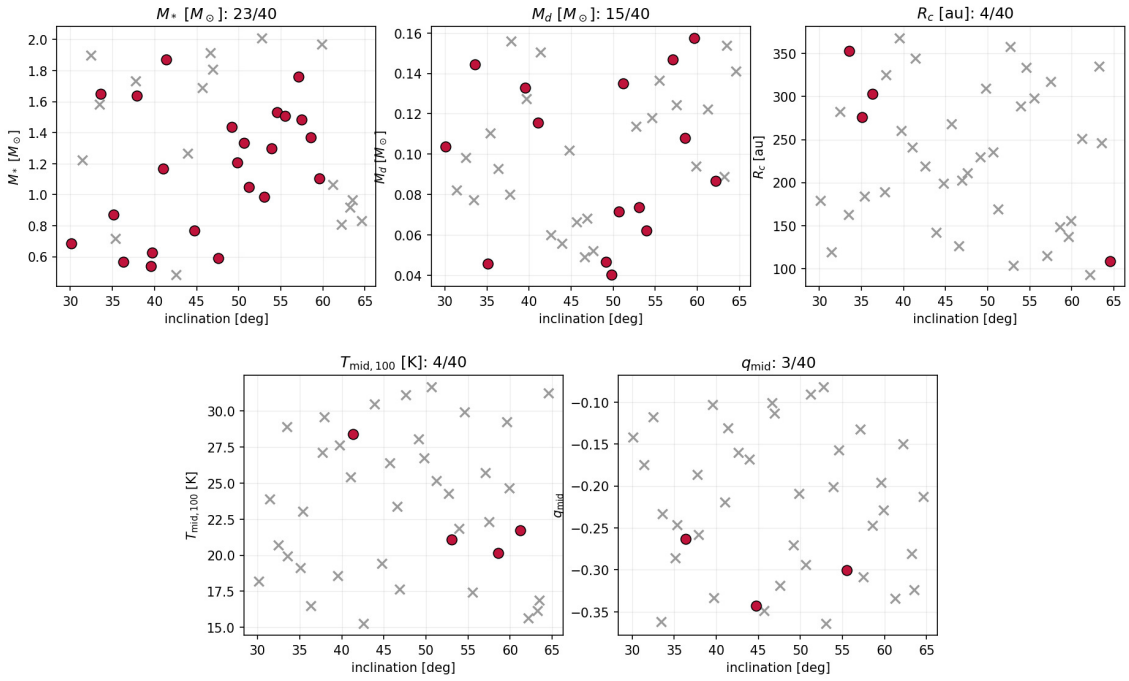


Figure 5.2: As Figure 5.1, for the curves corrected for the beam smearing.

The stellar mass improves from 16 to 11 per cent and its successes go from fourteen to twenty three, and the disc mass almost doubles its own. This is the expected gain, since the points which the correction makes usable are the innermost ones, where the curve is steepest and the gravity of the star dominates over the pressure. What improves most, however, is not the accuracy but the convergence. Figure 5.3 shows

this on the disc where the change is largest. The correction does not help for the case of the scale radius and it remains unrecovered. The thermal parameters barely move, and the slope passes the conditions in three discs, when before it passed in none.

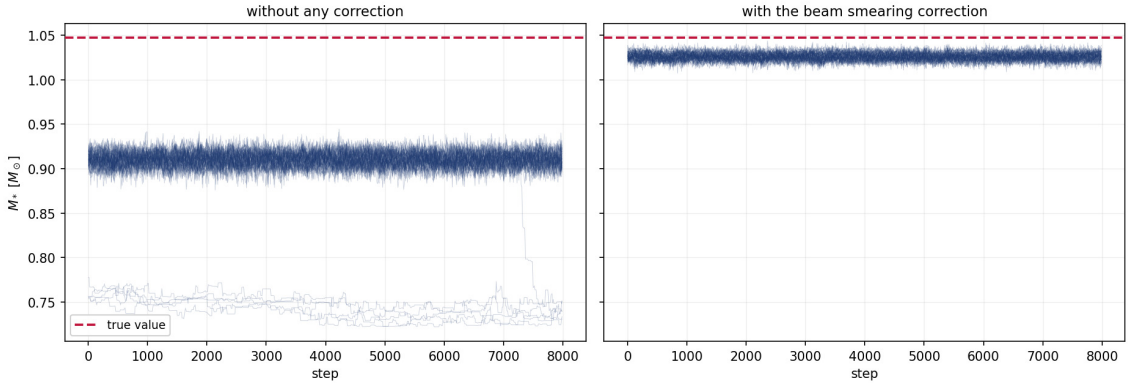


Figure 5.3: Chains of the stellar mass for the same disc, fitted on the extracted curve on the left and on the curve corrected for the beam smearing on the right, with the same vertical scale and the true value as a dashed line.

5.2.3 What the rotation curve measures

Table 5.4 collects, for each parameter, the number of discs in which it satisfies the five conditions, the error which Chapter 4 predicts on ideal data, and the error and the deviation from the true value which our fits reach on the discs where the parameter is recovered. The predicted error is the marginal one of Table 4.3 with two tracers.

Parameter	Recovered	Predicted [%]	Our error [%]	Our deviation [%]
M_*	23/40	1.4	0.5	5.5
M_d	15/40	29	17.8	21.8
R_c	4/40	34	13.0	28.2
$T_{\mathrm{mid},100}$	4/40	46	20.1	36.3
q_{mid}	3/40	91	22.3	57.8

Table 5.4: What the rotation curve measures. The second column counts the discs which satisfy the five conditions of recovery, the third is the marginal error predicted by Fisher's analysis for two tracers on ideal data, from Table 4.3, and the last two are the median standard deviation of the posterior and the median distance of its median from the true value, both relative to the true value and both computed on the discs where the parameter is recovered.

Only the stellar mass and the disc mass are recovered in a sizeable fraction of the sample. The stellar mass is the parameter which the curve constrains best, by more than an order of magnitude over every other one, and it is recovered in more than half of the discs. The disc mass passes in fifteen, and the threshold below which it cannot be measured [13, 14], which Table 4.1 uses to choose the intervals, appears in our sample as well: the ten discs whose M_d/M_* falls below 0.05 have a median deviation of 236 per cent and the disc mass is recovered in three of them, against 69

per cent and twelve recoveries for the thirty above it. The scale radius, the midplane temperature and its slope are recovered in four, four and three discs out of forty.

Our errors are smaller than the ones predicted for ideal data, by factors between 1.6 and 4.1, and this is not an improvement of the fit but a defect of it. The rotation curve is sampled every 5.9 au while the beam is 24 au wide, so there are four points per resolution element, and the likelihood treats them as independent. The posteriors come out narrower than they should by the square root of that ratio, a factor of two. The number of independent points we have, 47, is the same which Chapter 4 assumed, 24 for each of the two tracers, so once the factor of two is removed the width of our posteriors agrees with the prediction of Fisher's analysis. The information content of our fits is therefore the one predicted on ideal data.

What separates the fits of Chapter 4 is a bias. In every line of Table 5.4 the deviation is larger than the error, from a factor of 1.2 for the disc mass to a factor of eleven for the stellar mass, which is to say that the fits are precise and wrong rather than imprecise. That the residual of the extraction is a systematic error and not noise can be seen from three properties of the eighty curves of the sample. The first is its sign, since every single curve lies predominantly below the model, with 86 per cent of the points below it on the median curve, where independent noise would put half of them on each side. The second is its size, since the residual of a point is on average 3.2 times the error bar which the extraction assigns to that point, while thermal noise would give about 0.8. The third is its regularity, since two adjacent points of the residual are correlated with $r = +0.90$, whereas noise would make them independent. A quantity which is always of the same sign, several times larger than the uncertainty quoted on it and smooth in radius is a systematic error, and Section 5.1 measured how large it is next to the signal we are looking for.

Chapter 6

Conclusions

6.1 Summary

This work started from a line of papers which showed that the rotation curve of a protoplanetary disc is a measurement instrument, and that comparing it with a curve derived from the conservation of mass and the Navier-Stokes equations gives the mass of the star, the mass of the disc and the scale radius [9, 11, 12, 18]. In all of them the thermal structure is the one ingredient which is not measured dynamically but taken from outside, from the brightness temperature of optically thick lines [15, 19], a procedure which never reaches the midplane and which therefore extrapolates exactly the quantity that sets the pressure support. Section 1.3 asked whether the same curve which gives the masses can also give the thermal structure, and this thesis is the attempt to answer that question.

Chapter 2 derived the model which the answer rests on, from the conservation of mass and the Navier-Stokes equations up to the rotation curve of a vertically stratified disc. What matters for the question is that the thermal structure enters the curve only through the pressure gradient, which is a correction of the order of one per cent to Keplerian rotation, so the measurement is always the measurement of a small difference between two large numbers.

Chapter 3 built the instrument needed to test it. A synthetic disc is given to the radiative transfer code MCFOST [38], the resulting cube is degraded to the resolution and the noise of a real ALMA observation, and the rotation curve is extracted from it with discminer [17, 41] exactly as it would be extracted from real data.

Chapter 4 asked what the curve can constrain in principle, before any observational effect. Fisher's method showed that the information is there but that it is locked in degeneracies, between the stellar mass and the temperature, between the temperature and its slope, and between the disc mass, the scale radius and the temperature, so that the conditional errors are of the order of a few per cent while the marginal ones explode. It also showed the two levers which open those degeneracies, a second tracer emitting from a different height, which is the one that matters, and the radial extent of the curve, since the thermal signal grows with radius. Fitting synthetic curves with

Gaussian noise alone recovered every parameter within its errors, which proved that the limit is not in the model.

Chapter 5 ran the whole chain on a sample of forty discs covering the parameter space of Table 4.1, with two tracers each. The extraction turned out to carry a systematic error of -2.0 per cent on average over the fitting window, always of the same sign, which is twice the change that one standard deviation of the midplane temperature produces on the rotation curve and more than three times that of its slope. Only the stellar mass and the disc mass survived it. The comparison with Fisher's predictions in Section 5.2.3 showed that the fits lose no information with respect to ideal data, since the number of independent points of our curves is the one which Chapter 4 assumed and the width of our posteriors agrees with the prediction once the sampling of the curve is taken into account. What separates our results from that prediction is the bias alone, which for every parameter moves the answer further from the truth than the uncertainty the fit reports, by a factor of eleven in the case of the stellar mass.

The answer to the question of Section 1.3 is therefore negative. However our analysis indicates where to act. Radial extent alone is not the answer, since widening the window of the second tracer from 318 to 492 au on the median disc left the thermal parameters exactly where they were. What Chapter 4 identifies as the lever is a second tracer well separated in height from the first, which adds signal because the thermal term is measured by the difference between two curves emitted at different heights, although it leaves the bias where it is since that affects each curve on its own, and a more in-depth treatment of the chemistry of the disc is what could provide one which is at the same time low and still defined at large radii.

6.2 Future work

The most severe simplification of this work is chemical. The abundances of the molecules are imposed as constants and the regions where a molecule cannot survive are cut with thresholds, while the gas temperature is prescribed rather than computed. A thermochemical model would give the abundance structure and the gas temperature self consistently [8, 21], and this is not a cosmetic refinement, because the abundance structure is what decides where a molecule emits from, and therefore which tracers are available at all.

That this is the crux can be seen from the tracer problem itself, which is the single largest obstacle to the thermal parameters. The thermal signal grows with radius, but the second tracer runs out before the first one does, and beyond that radius the curve carries a single tracer, for which Table 4.3 gives errors three to five times larger. Within a model with constant abundances this cannot be repaired. Raising the abundance of a molecule extends its surface outwards but raises it at the same time, so the vertical separation which is the actual lever is lost in exchange for the radial extent. Choosing a lower rotational transition does not help either, because at the temperatures of these discs the population of the rotational levels peaks around the transition we already use, so a lower one is both lower and less extended. Real discs are not bound by this, and the comparison is instructive: in HD 163296 the CN emits from the same

height as the ^{13}CO but its surface is traced about twenty per cent further out [49], and that gain comes from a chemistry which keeps the molecule abundant where the CO isotopologues are already gone. Identifying tracers which are both well separated in height and defined out to large radii is therefore the same problem as modelling the chemistry properly, and solving it would raise the thermal signal, which is the quantity the bias of the extraction is currently larger than. It would not reduce that bias, so this direction and the next one have to be pursued together.

The second direction is the extraction. The bias of a few per cent which Section 5.1 measures is what costs the stellar mass a factor of eight with respect to what the data allow, and it is a systematic error, so it cannot be reduced by observing longer. The most promising place to look is the emission surface, which the curve is deprojected with and which enters the velocity through the term of order $(z/R)^2$. On our sample the height at 100 au recovered by the fit differs from the true one by about ten per cent for the ^{12}CO and thirty for the ^{13}CO, and the exponent which describes where the surface turns over is essentially not measured at all, which is why the outer cut of Section 5.1 had to be taken from the brightness of the line rather than from the surface. Determining the surface better, or propagating its uncertainty into the fit instead of treating it as exact [9, 14], would act directly on the dominant term. A second possibility is to generalise the correction of Section 3.2.3, which already calibrates one step of the chain by passing a known model through it, to the whole extraction rather than to the beam alone.

A further direction would use information which the method currently discards. Section 3.2.1 measures the temperature of the gas from the brightness of the lines, at the height from which each tracer emits, but the fits of this work make no use of it and leave the thermal parameters free with uniform priors. Building an informed prior on the thermal structure from those brightness temperatures would let the kinematics refine a structure which photometry has already constrained, instead of determining it from nothing, and it requires no new observations.

A further direction concerns the dynamical structure. The discs used here are axisymmetric and in steady state, and their self-gravity enters only through the axisymmetric term of the rotation curve. Generating the density and velocity fields from a smoothed particle hydrodynamics simulation [50] would make it possible to test the method on self-gravitating and non axisymmetric discs, where spiral structure perturbs the curve in a way which the present model cannot represent.

Finally, and only once the previous points have been addressed, the bias of the extraction could be attacked directly. Section 5.2.3 showed that the residual is systematic rather than random, since it is always of the same sign, several times larger than the uncertainty quoted on each point, and smooth enough that adjacent points of the curve are correlated with $r = +0.90$. A quantity with these properties is in principle learnable, and a decomposition of the residual into a small number of components, with their coefficients predicted from observable quantities, is the natural way to try. We regard this as a last resort rather than a solution, because it would be the only step of the method which is calibrated on data rather than derived from physics, and because a correction learned on synthetic discs is only as good as the synthetic discs

it was learned on, which brings the argument back to the chemistry.

Appendix A

SPH simulation

This appendix describes a simple smoothed particle hydrodynamics (SPH) simulation of a viscous disc which I wrote from scratch. It is an original work aimed at understanding the way these simulations work and has served as the final project of the "Problem Solving" exam of my Laurea Magistrale. Much more advanced codes like Phantom [50] are SPH simulations actually used in research for model validation, as in Martire et al. [10]. The paper I followed to build this simulation is Lodato and Price [51]. This work also conveniently serves as visual explanation for Section 2.1.5.1 which otherwise might be too abstract to grasp.

A.1 What is an SPH simulation

An SPH simulation of a fluid adopts a Lagrangian approach to describe the time evolution. It is therefore particle centred instead of using a spatial grid like in an Eulerian approach. Each particle of the fluid carries a mass, position and velocity. Fields like density or pressure are obtained by averaging over the neighbouring particles with a smoothing kernel, and the width of this kernel, which is called smoothing length, sets the resolution of the simulation.

This approach is convenient for a disc [50] because it easily adapts to problems where the spatial scale evolves along time. It is also numerically convenient since there is no grid, and the material does not have to be transported from one cell to the other. Another interesting property of SPH simulations is that they conserve momentum and energy by construction.

A.2 The code

The simulation is two dimensional for complexity reasons and therefore the density is a surface one. The disc is viewed face-on and is made up of $N = 1000$ particles that move in its plane. These particles are of equal mass and orbit a star of one solar mass which is a fixed centre of force. The temperature is a power law of the radius and does not change in time:

$$T(R) = T_{100}\left(\frac{R}{100\,\mathrm{au}}\right)^{-q} \tag{A.1}$$

with $T_{100} = 20$ K and $q = 0.23$. The equation of state is therefore locally isothermal:

$$P = c_s^2(R)\,\rho, \qquad c_s^2 = \frac{k_B T(R)}{\mu m_p} \tag{A.2}$$

The code can start from two different initial conditions. The first one is a narrow ring where the particles are distributed with a uniform surface density between 180 and 220 au and orbit at the keplerian velocity. The second one is the self-similar solution defined in Equation 2.49 with $R_c = 150$ au and $\gamma = 1$. The radii of the particles are drawn from this profile between 50 and 600 au, and the velocity is given by the solution of the radial component of the Navier-Stokes equation in two dimensions, analogous to the one shown in Section 2.1.7:

$$v_\phi^2 = v_k^2 + \frac{R}{\rho}\frac{\partial P}{\partial R} = v_k^2 - c_s^2\left[\gamma + (2-\gamma)\left(\frac{R}{R_c}\right)^{2-\gamma} + q\right] \tag{A.3}$$

The mass of the disc does not enter the dynamics because the self-gravity is neglected for simplicity.

Each particle has its own smoothing length h_a, which adapts to the local density because some particles might otherwise have too few or too many neighbours during the simulation. The density is given by the sum of the masses of the neighbours weighted by the smoothing kernel W:

$$\rho_a = \sum_b m_b\, W(|\mathbf{r}_a - \mathbf{r}_b|, h_a) \tag{A.4}$$

We chose a cubic spline kernel, which is equal to zero beyond a distance $2h$. The smoothing length is then updated as $h_a = \eta(m_a/\rho_a)^{1/2}$ with $\eta = 1.2$, which corresponds to about eighteen neighbours for each particle. This adaptability makes h a function of the density and therefore we have a system of two coupled equations whose solution we get by iteration.

The acceleration of a particle is computed as the sum of three terms, the pressure, the viscosity and the gravity of the star:

$$\frac{d\mathbf{v}_a}{dt} = -\sum_b m_b\left(\frac{P_a}{\rho_a^2} + \frac{P_b}{\rho_b^2} + \Pi_{ab}\right)\nabla_a W_{ab} - \frac{GM_*}{|\mathbf{r}_a|^3}\mathbf{r}_a \tag{A.5}$$

Written in this way the term inside the sum is symmetric in a and b so as to respect Newton's third law. The kernel here is evaluated on the average of the smoothing length of the considered couple of particles.

The viscosity is contained in the term Π_{ab}, which is different from zero only when two particles are approaching each other to avoid compenetration:

$$\Pi_{ab} = \frac{-\alpha_{\mathrm{AV}} \bar{c}_s \mu_{ab} + \beta_{\mathrm{AV}} \mu_{ab}^2}{\bar{\rho}_{ab}}, \qquad \mu_{ab} = \frac{\bar{h}_{ab}\, \mathbf{v}_{ab} \cdot \mathbf{r}_{ab}}{|\mathbf{r}_{ab}|^2 + \epsilon \bar{h}_{ab}^2} \tag{A.6}$$

where the bar is the average between the two particles and we use $\alpha_{\mathrm{AV}} = 1$, $\beta_{\mathrm{AV}} = 2$ and $\epsilon = 0.01$. This term was introduced to treat the shocks but in a disc it behaves as a physical viscosity, which is what makes our disc evolve. The topic is complex and Lodato and Price [51] treat it in detail, showing that an artificial viscosity of this kind corresponds to a kinematic viscosity $\nu = \alpha c_s H$, as in the prescription of Shakura and Sunyaev [52], where:

$$\alpha \simeq \frac{1}{10} \alpha_{\mathrm{AV}} \frac{\langle h \rangle}{H} \tag{A.7}$$

We have to be careful with this relation because it was obtained in three dimensions and for a form of the term which is applied to all the couples of particles and has no β_{AV}, so for our code it only gives the order of magnitude of α.

A.3 The time loop

At every step the code performs the following operations.

1. It chooses the time step as the smallest of three limits: the time a sound wave needs to cross a smoothing length, the time needed to move by a smoothing length with the current acceleration, and a viscous limit which depends on the signal velocity of the particles so as to capture every physical phenomenon.
2. It updates the velocities over half a time step with the old accelerations.
3. It moves the particles over a whole time step with the updated velocities.
4. It removes the particles that are closer than 30 au to the star, which are considered accreted.
5. It computes the densities again, the smoothing lengths and the accelerations at the new positions.
6. It updates the velocities over the second half of the time step with the new accelerations.

This evolution scheme is called kick drift kick leapfrog and it is possible to prove that this scheme exactly conserves energy because it is symplectic. However this only holds for a fixed timestep while ours is adaptive. The introduced error is nonetheless negligible as we will show in Section A.6. Every 50 years the code saves the positions and the densities of the particles together with the total energy and angular momentum, until the final time of 5000 years. The whole simulation takes less than a minute on Colab.

A.4 The spreading of a ring

Figure A.1 shows the evolution of the first initial condition. The ring is the simplest most spectacular case since the self-similar one is, unremarkably, self-similar and does not change its form in time. In the first hundreds of years the edges of the ring are smoothed by the pressure: the particles of the inner edge have no neighbours on the side of the star and they are pushed towards it, and the same happens at the outer edge in the opposite direction. After this transient the viscosity takes over. The ring spreads in both directions, its density drops by almost an order of magnitude and a few particles reach the sink around the star as shown in the top half of Figure A.1. The picture is the one described by the diffusion equation: the particles that move inwards give their angular momentum to the ones that move outwards and the total angular momentum is conserved. Part of the spreading is also the result of the pressure, which the diffusion equation does not contain. The particles pushed by the sharp edges at the beginning end up on eccentric orbits, and this is the reason why the outer part of the ring comes back towards the star between 2500 and 5000 years. This can be seen in the bottom half of Figure A.1.

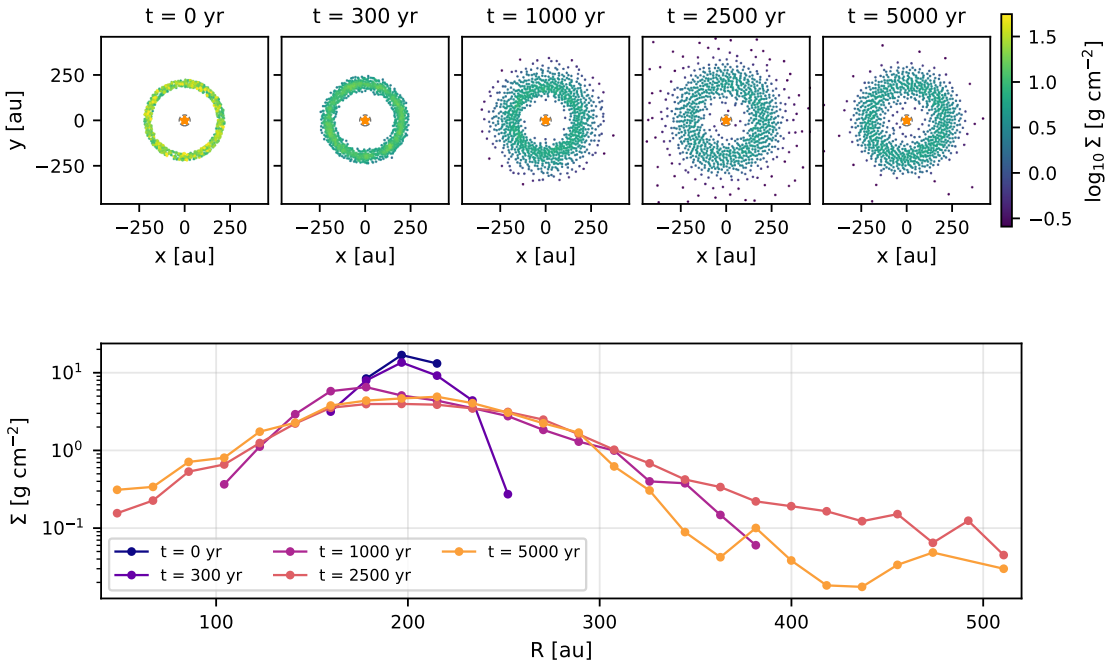


Figure A.1: Evolution of the ring at five different times. The particles are coloured according to their surface density, the dashed circle is the sink of radius 30 au and the star sits at the centre. At the beginning the particles are between 180 and 220 au and they orbit at the keplerian velocity. The lower panel shows the surface density measured in radial bins at the same five times.

A.5 The self-similar disc

Figure A.2 shows the evolution of the second initial condition at the same five times of the ring, and below it the surface density, which we measure by counting the particles

in rings equally spaced in the logarithm of the radius. The disc keeps its shape, which is what we expect from a solution built to evolve in a self-similar way. Between 100 and 500 au the measured profile stays on the analytic one and fitting it with the self-similar function we get $R_c = 150 \pm 9$ au at the beginning and $R_c = 155 \pm 10$ au at the end, while the number of particles in this region goes from 676 to 683. What changes are the two edges of the initial distribution, for the reasons explained in Section A.6.

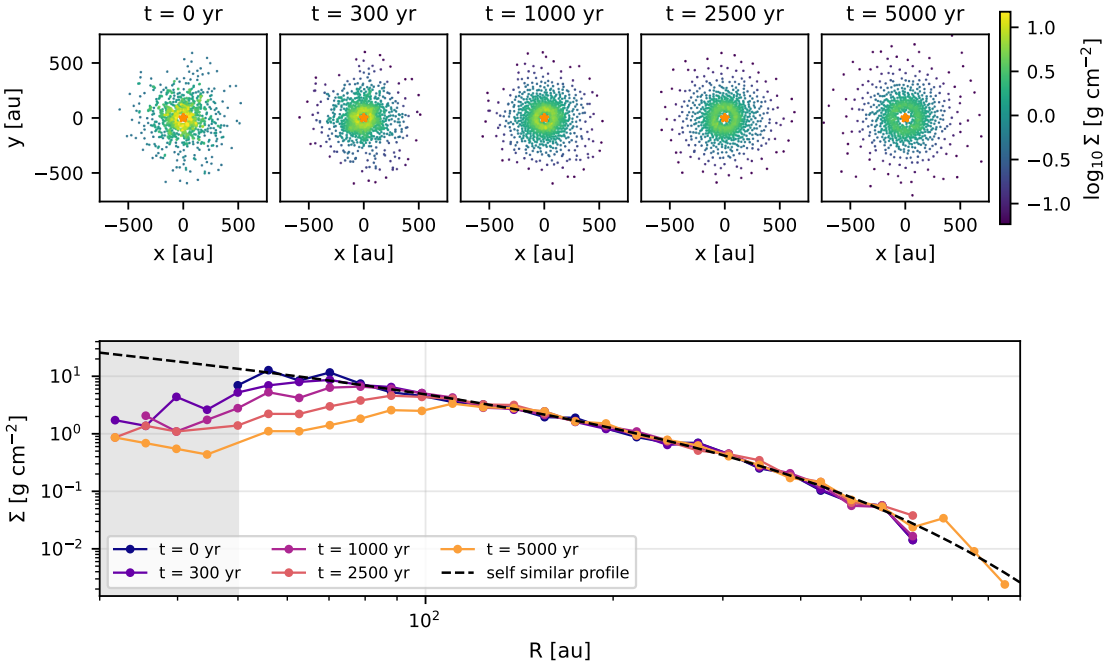


Figure A.2: Evolution of the disc which starts from the self-similar profile, with the same colours and the same times of Figure A.1. The lower panel shows its surface density at those times, compared with the self-similar profile of Equation 2.49 with $R_c = 150$ au and $\gamma = 1$. The grey band is the region between the sink and the inner edge of the initial distribution.

The time scales tell us how much this result is worth. The effective viscosity of our simulation gives a viscous time of about 7×10^4 years at R_c, which is fifteen times longer than the simulation itself, so the self-similar solution predicts a change of R_c of a few per cent and our errors are larger than that. What we can see in the figure is that a disc with the self-similar surface density and a rotation curve which is not keplerian, because it contains the pressure, stays as it is for several orbits.

This is useful for our model, where the self-similar solution of Section 2.1.5.1 is used for a disc that does not rotate at the keplerian velocity. To measure how much this matters we ran the same simulation starting from keplerian velocities, and after 5000 years the two surface densities differ on average by two per cent between 100 and 500 au, while the scatter between the bins is of about twelve per cent. The simulation is therefore not able to tell the two rotation laws apart, which is another way of saying that a deviation of order $(H/R)^2$ from the keplerian velocity does not change the profile.

A.6 Known errors and how we handle them

The first quantity we check is the angular momentum. All the forces of the code are central, so it can change only when the sink removes a particle. We therefore record the angular momentum that every accreted particle takes away. The lower row of Figure A.3 shows the result. The angular momentum of the disc alone drops, while the sum with the accreted one is conserved to one part in 10^{16} in both simulations. That drop is the effect of the accretion.

The energy behaves in a different way. The equation of state is locally isothermal, so the gas radiates immediately the heat produced by the viscosity and by the compressions, and the code does not follow the internal energy. The upper row of the same figure shows the energy with its components. The total energy of the disc grows because the sink removes the particles which are the most bound. Adding the energy of the accreted particles the total decreases by 35 per cent in the self-similar case and by 12 per cent in the ring. This is the energy that the gas loses while the particles move towards the star. In a real disc we would see it as the accretion luminosity of Equation 2.1.

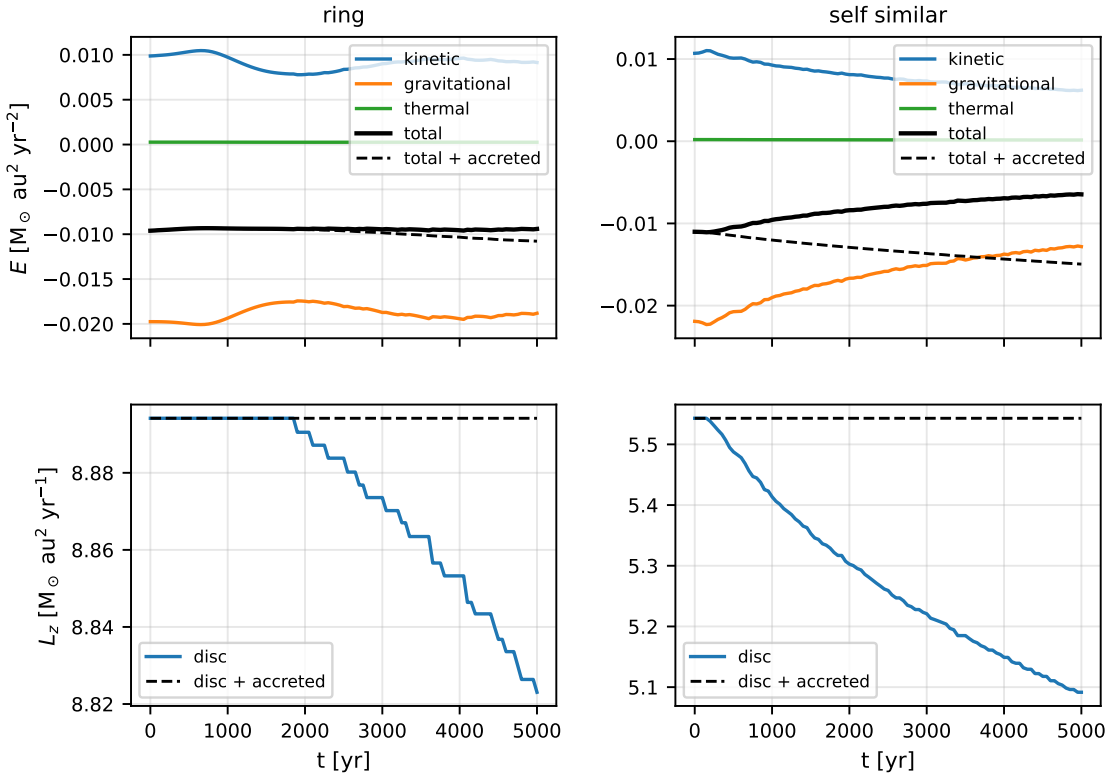


Figure A.3: Conservation tests for the ring, on the left, and for the self-similar disc, on the right. The upper row shows the energy of the disc with its kinetic, gravitational and thermal components, together with the total energy of the disc and the one obtained by adding the energy of the particles removed by the sink. The lower row shows the angular momentum of the disc and the sum of it with the angular momentum carried away by the same particles.

These tests do not show the other two limits of the simulation. The first one is at the edges of the disc. A particle which sits at an edge has neighbours only on one side, so its kernel is half empty and its density is underestimated. The pressure force between two particles is always repulsive, and at an edge there is nothing on the other side to balance it. Both edges are therefore pushed away from the rest of the material. The outer one spreads beyond 600 au and the inner one loses its particles into the sink, 19 per cent of the mass in 5000 years. This is why in Section A.5 we compare the simulation with the analytic profile only between 100 and 500 au. The second limit is the resolution. With 1000 particles the smoothing length is everywhere comparable to the scale height of the disc or larger, so the viscosity of Section A.2 corresponds to an α between 0.1 and 0.2. This is much larger than the viscosity of a real disc and it makes our discs evolve faster than they should.

Bibliography

[1] ALMA Partnership et al. “The 2014 ALMA Long Baseline Campaign: First Results from High Angular Resolution Observations toward the HL Tau Region”. In: *ApJL* 808 (2015), p. L3.

[2] NASA Exoplanet Archive. *Confirmed Planets Table*. https://exoplanetarchive.ipac.caltech.edu. accessed in August 2026. 2026.

[3] W. Zhu and S. Dong. “Exoplanet Statistics and Theoretical Implications of Planet Formation”. In: *ARA&A* 59 (2021), pp. 291–336. DOI: 10.1146/annurev-astro-112420-020055.

[4] J. Ji et al. “CHES: a space-borne astrometric mission for the detection of habitable planets of the nearby solar-type stars”. In: *Research in Astronomy and Astrophysics* 22 (2022), p. 072003. DOI: 10.1088/1674-4527/ac77e4. arXiv: 2205.05645.

[5] A. Wootten and A. R. Thompson. “The Atacama Large Millimeter/Submillimeter Array”. In: *Proceedings of the IEEE* 97 (2009), p. 1463. DOI: 10.1109/JPROC.2009.2020572. arXiv: 0904.3739.

[6] C. Pinte et al. “Direct mapping of the temperature and velocity gradients in discs. Imaging the vertical CO snow line around IM Lupi”. In: *A&A* 609 (2018), A47.

[7] R. Teague et al. “Molecules with ALMA at Planet-forming Scales (MAPS). XVIII. Kinematic Substructures in the Disks of HD 163296 and MWC 480”. In: *ApJS* 257 (2021), p. 18. DOI: 10.3847/1538-4365/ac1438. arXiv: 2109.06218.

[8] K. Zhang, A. S. Booth, C. J. Law, et al. “Molecules with ALMA at Planet-forming Scales (MAPS). V. CO Gas Distributions”. In: *ApJS* 257 (2021), p. 5.

[9] V. Pezzotta et al. “The two-dimensional pressure structure of the HD 163296 protoplanetary disk as probed by multi-molecule kinematics”. In: *A&A* 694 (2025), A108. DOI: 10.1051/0004-6361/202451307. arXiv: 2501.05517.

[10] P. Martire et al. “Rotation curves in protoplanetary disks with thermal stratification”. In: *A&A* 686 (2024), A9. arXiv: 2402.12236.

[11] G. Lodato et al. “Dynamical mass measurements of two protoplanetary discs”. In: *MNRAS* 518 (2023), p. 4481. arXiv: 2211.03712.

[12] B. Veronesi et al. “A Dynamical Measurement of the Disk Mass in Elias 2-27”. In: *ApJL* 914 (2021), p. L27. DOI: 10.3847/2041-8213/abfe6a.

[13] B. Veronesi et al. “Weighing protoplanetary discs with kinematics: physical model, method and benchmark”. In: *A&A* 688 (2024), A136. arXiv: 2405.15944.

[14] S. M. Andrews et al. “On Kinematic Measurements of Self-gravity in Protoplanetary Disks”. In: *ApJ* 970 (2024), p. 153. DOI: 10.3847/1538-4357/ad5285. arXiv: 2405.19574.

[15] C. J. Law et al. “Molecules with ALMA at Planet-forming Scales (MAPS). IV. Emission Surfaces and Vertical Distribution of Molecules”. In: *ApJS* 257 (2021), p. 4. DOI: 10.3847/1538-4365/ac1439. arXiv: 2109.06217.

[16] R. Teague. “eddy: Extracting Protoplanetary Disk Dynamics with Python”. In: *JOSS* 4 (2019), p. 1220.

[17] A. F. Izquierdo et al. “The Disc Miner. I. A parametric modelling technique to detect substructures and kinematic perturbations”. In: *A&A* 650 (2021), A179.

[18] C. Longarini et al. “exoALMA. XII. Weighing and sizing exoALMA disks with rotation curve modelling”. In: *ApJL* 984 (2025), p. L17. arXiv: 2504.18726.

[19] C. P. Dullemond et al. “Midplane temperature and outer edge of the protoplanetary disk around HD 163296”. In: *A&A* 633 (2020), A137. arXiv: 1911.12434.

[20] C. J. Lada. “Star formation: from OB associations to protostars”. In: *Star Forming Regions*. Vol. 115. IAU Symposium. 1987, p. 1.

[21] A. Miotello et al. “Setting the Stage for Planet Formation: Measurements and Implications of the Fundamental Disk Properties”. In: *Protostars and Planets VII*. Vol. 534. ASP Conference Series. 2023, p. 501. arXiv: 2203.09818.

[22] J. Frank, A. King, and D. Raine. *Accretion Power in Astrophysics*. 3rd ed. Cambridge University Press, 2002. ISBN: 9780521629577.

[23] S. J. Kenyon and L. Hartmann. “Spectral Energy Distributions of T Tauri Stars: Disk Flaring and Limits on Accretion”. In: *ApJ* 323 (1987), p. 714. DOI: 10.1086/165866.

[24] E. I. Chiang and P. Goldreich. “Spectral Energy Distributions of T Tauri Stars with Passive Circumstellar Disks”. In: *ApJ* 490 (1997), p. 368.

[25] L. Hartmann et al. “Accretion and the Evolution of T Tauri Disks”. In: *ApJ* 495 (1998), p. 385. DOI: 10.1086/305277.

[26] C. F. Manara, M. Ansdell, G. P. Rosotti, et al. “Demographics of Young Stars and their Protoplanetary Disks: Lessons Learned on Disk Evolution and its Connection to Planet Formation”. In: *Protostars and Planets VII*. Vol. 534. ASP Conference Series. 2023, p. 539. arXiv: 2203.09930.

[27] S. M. Andrews. “Observations of Protoplanetary Disk Structures”. In: *ARA&A* 58 (2020), p. 483. DOI: 10.1146/annurev-astro-031220-010302.

[28] P. J. Armitage. *Astrophysics of Planet Formation*. 2nd ed. Cambridge University Press, 2020. ISBN: 9781108420501.

[29] J. E. Pringle. “Accretion discs in astrophysics”. In: *ARA&A* 19 (1981), p. 137. DOI: 10.1146/annurev.aa.19.090181.001033.

[30] G. Bertin and G. Lodato. “A class of self-gravitating accretion disks”. In: *A&A* 350 (1999), p. 694.

[31] L. D. Landau and E. M. Lifshitz. *Fluid Mechanics*. 2nd ed. Vol. 6. Course of Theoretical Physics. Pergamon Press, 1987. ISBN: 9780750627672.

[32] D. Lynden-Bell and J. E. Pringle. “The evolution of viscous discs and the origin of the nebular variables”. In: *MNRAS* 168 (1974), p. 603.

[33] A. Erdélyi et al. *Tables of Integral Transforms*. Vol. 2. Bateman Manuscript Project. New York: McGraw-Hill, 1954.

[34] E. Dartois, A. Dutrey, and S. Guilloteau. “Structure of the DM Tau outer disk: probing the vertical kinetic temperature gradient”. In: *A&A* 399 (2003), p. 773.
[35] A. Toomre. “On the gravitational stability of a disk of stars”. In: *ApJ* 139 (1964), p. 1217.
[36] K. Kratter and G. Lodato. “Gravitational Instabilities in Circumstellar Disks”. In: *ARA&A* 54 (2016), pp. 271–311. DOI: 10.1146/annurev-astro-081915-023307.
[37] M. Galloway-Sprietsma et al. “exoALMA. V. Gaseous Emission Surfaces and Temperature Structures”. In: *ApJL* (2025). arXiv: 2504.19902.
[38] C. Pinte et al. “Monte Carlo radiative transfer in protoplanetary disks”. In: *A&A* 459 (2006), p. 797.
[39] R. Teague et al. “exoALMA. I. Science Goals, Project Design and Data Products”. In: *ApJL* (2025). arXiv: 2504.18688.
[40] K. M. Flaherty, A. M. Hughes, K. A. Rosenfeld, et al. “Weak Turbulence in the HD 163296 Protoplanetary Disk Revealed by ALMA CO Observations”. In: *ApJ* 813 (2015), p. 99. DOI: 10.1088/0004-637X/813/2/99.
[41] A. F. Izquierdo, L. Testi, S. Facchini, et al. “The Disc Miner. II. Substructures and velocity perturbations in the MAPS disks”. In: *A&A* 674 (2023), A113.
[42] A. F. Izquierdo, J. Stadler, J. Bae, et al. “exoALMA. III. Line-intensity Modeling and System Property Extraction from Protoplanetary Disks”. In: *ApJL* (2025). DOI: 10.3847/2041-8213/adc439. arXiv: 2504.19986.
[43] C. J. Law, R. Teague, R. A. Loomis, et al. “CO Line Emission Surfaces and Vertical Structure in Mid-Inclination Protoplanetary Disks”. In: *ApJ* (2022). arXiv: 2205.01776.
[44] C. J. Law et al. “Mapping Protoplanetary Disk Vertical Structure with CO Isotopologue Line Emission”. In: *ApJ* 948 (2023), p. 60. DOI: 10.3847/1538-4357/acb3c4. arXiv: 2212.08667.
[45] E. Weaver, A. Isella, and Y. Boehler. “Empirical Temperature Measurement in Protoplanetary Disks”. In: *ApJ* 853 (2018), p. 113. DOI: 10.3847/1538-4357/aaa481. arXiv: 1801.03478.
[46] M. Tegmark, A. N. Taylor, and A. F. Heavens. “Karhunen-Loève Eigenvalue Problems in Cosmology: How Should We Tackle Large Data Sets?” In: *ApJ* 480 (1997), pp. 22–35. DOI: 10.1086/303939.
[47] D. Coe. “Fisher Matrices and Confidence Ellipses: A Quick-Start Guide and Software”. In: *arXiv e-prints* (2009). arXiv: 0906.4123.
[48] A. Heavens. “Statistical Techniques in Cosmology”. In: *arXiv e-prints* (2009). arXiv: 0906.0664.
[49] T. Paneque-Carreño, A. Miotello, E. F. van Dishoeck, et al. “Directly tracing the vertical stratification of molecules in protoplanetary disks”. In: *A&A* 669 (2023), A126. arXiv: 2210.01130.
[50] D. J. Price et al. “Phantom: A Smoothed Particle Hydrodynamics and Magnetohydrodynamics Code for Astrophysics”. In: *PASA* 35 (2018), e031. DOI: 10.1017/pasa.2018.25. arXiv: 1702.03930.
[51] G. Lodato and D. J. Price. “On the diffusive propagation of warps in thin accretion discs”. In: *MNRAS* 405 (2010), p. 1212. DOI: 10.1111/j.1365-2966.2010.16526.x. arXiv: 1002.2973.

[52] N. I. Shakura and R. A. Sunyaev. “Black holes in binary systems. Observational appearance”. In: *A&A* 24 (1973), p. 337.

Acknowledgements

Voglio ringraziare sinceramente il professor Lodato per il tempo e l'attenzione continua che mi ha dedicato. Durante i nostri incontri settimanali, andati avanti per mesi, mi sono reso conto di quanto si possa essere appassionati del proprio lavoro e, in particolare, dell'Astrofisica. Non sono meno importanti i contributi di Cristiano, che è sempre stato disponibile ad aiutarmi, dall'impostazione del codice fino alla correzione finale di questa tesi. Quiero también agradecer a Andrés por haberme dedicado parte de su agosto ayudándome con las curvas de rotación, extraídas con su increíble software. Ringrazio anche Viviana, che mi ha fatto conoscere Discminer e le piccolissime deviazioni dalla keplerianità. Grazie a questo gruppo ho potuto vivere una genuina esperienza di ricerca indipendente senzaperò essere mai lasciato completamente a me stesso, il che mi ha aiutato a capire che la ricerca è ciò che mi piacerebbe davvero fare.

Scrivo questi ringraziamenti in una sera di Londra, mentre papà è sicuramente più ansioso di me perché deve andare a stampare la tesi e io mi sono ridotto all'ultimo come al solito. Nel frattempo mamma pensa al completo che ovviamente non ho portato in tempo in lavanderia. Mi sembra quindi doveroso cominciare dai miei genitori, che nonostante tutta la truppa trovano sempre il tempo di aiutarmi e che con il loro supporto mi permettono di fare tutto ciò che mi metto in testa. Mi auguro di poter un giorno restituire tutto, anche se probabilmente non funziona così. Grazie Emma, per seguirmi sempre in giro per il mondo senza farmi mai sentire solo e per l'esempio di testardaggine che mi dai nel vivere la vita come vuoi tu. Ringrazio le bimbe per farmi sentire ancora bambino quando giochiamo insieme, e spero di essere, insieme all'altra sorellona, la prova che possiamo essere quello che vogliamo. Non vedo l'ora di vedere tutto quello che farete nella vita. Poi ci sono i miei nonni, che sono tutti una grandissima motivazione a dare il meglio: sono molto fortunato ad avervi conosciuti. Non sono sicuro poi che il ringraziamento per nonno Sergio vada fra quelli della famiglia, dato che è più che altro un correlatore dall'infinita esperienza. Dopo ogni esame, che sia andato bene o male, sono sempre felice di chiamarti.

Ringrazio anche tutti gli amici che sono fortunato ad avere. Sbronks, compagno di avventure che devono ancora realizzarsi, sempre pronto a vedere un horror e ad ascoltare tutte le mie crisi. Fede, l'altra decana, anche lei sempre pronta ad ascoltarmi. Tengo tantissimo a tutti i momenti che passiamo insieme, anche quando parliamo dei gossip più scemi. Poi i miei amici pugliesi, per non dire di peggio, Luca e Claudio, che a forza di taralli e ospitalità si sono fatti voler bene. Ringrazio tanto anche i miei compagni di suonate, Ange, Ludo ed Ennio la *****. Siete fra le persone più appassionate che conosca e i concerti insieme sono fra i ricordi più belli di questi anni. Non

sarei poi dove sono ora senza le capacità di talent manager della queen Cedro, per cui ti ringrazio. In generale sono grato a tutti i miei compagni d'anno, ognuno ha contribuito a suo modo a rendere bellissimi questi anni. Lo stesso vale per tutti gli altri collegiali, che ogni giorno rendono viva la nostra speciale comunità.

È poi obbligatorio ringraziare l'anziano Seppia. Non è esagerato dire che senza i tuoi appunti e consigli di carriera non mi sarei mai laureato. Un grazie va per forza anche a Davide, il mio amico di più lunga data, perché in questo momento simbolico mi rendo conto di quante cose abbiamo passato insieme e di quante ancora ne passeremo. I am also grateful to my Erasmus friends, none of whom was actually doing an Erasmus, for showing me that special people can always come into your life unexpectedly. También les doy las gracias a mi familia y a mis amigos chilenos. Me parece increíble que, a lo largo de tantos años, me pueda sentir en casa en la otra parte del mundo. Un ringraziamento speciale va poi a Laura, ci sarebbero mille motivi ma in questa occasione mi piace ricordare quando premevi invio per me alle domande di Claude sul lavoro di tesi.

Ormai qui a Londra si è fatto tardi, spero di non essermi dimenticato nessuno ma se l'ho fatto è perché ho un ridottissimo span di attenzione e non perché non ci tengo.