Star Formation studies using the Herschel-SPIRE Imaging FTS

by

Adam Rykala

A thesis submitted to

Cardiff University

for the degree of

Doctor of Philosophy

September 2011

DECLARATION

This work has not previously been accepted in substance for any degree and is not being concurrently submitted in candidature for any degree.

Signed .....

Date .....

STATEMENT 1

The work presented in this thesis is all my own work carried out under the supervision of Prof. D. Ward-Thompson, Prof. P. Ade and Dr. K.G. Isaak with the exception of the following:

LoCi hardware was constructed by the Astronomical Instrumentation Group's workshop here at Cardiff University, and the choice of hardware and purchasing was performed before the start of this student project. The MPFIT IDL fitting routine used by SLIDE for greybody fitting, line contamination fitting and for prestellar core fitting was produced by Markwardt (2009). The examination of the fringing produced by SPIRE was undertaken using a fringing function produced by Giorgio Savini for his own work. The SPIRE photometer RSRF function was supplied by Locke Spencer and Kate Isaak from joint work. SPIRE observations of the sources for the line contamination chapter were supplied by Kate Isaak and Ed Polehampton, and the work on line contamination was undertaken by myself working with Dr Isaak. SPIRE observations of the prestellar cores were processed in HIPE by myself, using a method for telescope correction suggested by Ed Polehampton. The data came from the SPIRE Guaranteed Time Key Programme on the interstellar medium proposed by SPIRE SAG4, which was led by A. Abergel and provided to me by D. Ward-Thompson and J. Kirk.

A bibliography is appended.

Signed .....

Date .....

STATEMENT 2

I hereby give consent for my thesis, if accepted, to be available for photocopying and for inter-library loan, and for the title and summary to be made available to outside organisations.

Signed .....

Date .....

ACKNOWLEDGMENTS

To Φ_L, my wife and my true universal constant. As a mature student, this body of work would not exist without your un-ending support. This work is dedicated to the memory of Czeslaw and Coral Rykala, my parents. Although you are no longer with me in body, I owe all of who I am and what I achieve to your love and guidance.

I would like to offer my heartfelt thanks to Prof. Derek Ward-Thompson, whose support and guidance has shone bright light into darkness and helped transform this thesis into a body of work I can be proud of. I would also like to thank Prof. Peter Ade, Locke Spencer, Giorgio Savini and Dr Kate Isaak for additional help and guidance during Chapters 1 and 2 (Prof. Ade, Locke Spencer, Giorgio Savini), and Chapters 3 and 4 (Dr Isaak, Locke Spencer, Giorgio Savini). The work presented in this PhD thesis was performed at the Astronomy and Physics School in Cardiff University, and I would like to thank all staff and students there for their help and support.

Working with the SPIRE FTS data has been made infinitely more enjoyable by the SPIRE FTS research teams and support staff in RAL, and I would like to thank them for the help provided when needed. I would like to extend this heartfelt thanks to members of the various research groups I have worked with, whose names would create a document that would rival the size of this work. The design, testing and implementation of SLIDE was greatly helped by the guidance from the various testers who suggested features, worked through bugs and then used the code to produce results for their own work. I would like to thank David Naylor and his colleagues at the University of Lethbridge for guidance and help involving SLIDE and LoCi. SLIDE exists in its current state because of the help and collaboration of Locke Spencer, Naseem Rangwela, Kate Isaak, Nicola Sacchi, Mikako Matsuura, Mike Barlow, Brad Gomm, and many others.

I would like to personally thank my office colleagues, Jason Kirk, Lucy Wilcock, Ciara Quinn and Sarah Harry, for everything they have done to make me feel welcome in my new home. My time at Cardiff University has been one of great pleasure, and I wish to thank the Star Formation group for a lively, interesting and fascinating environment in which to work. On a personal level, I would like to add special thanks to Dr B. Richardson, Dr C. Tucker and Prof A. Whitworth for moral support. Thanks to them, a candle has illuminated many dark places.

'*For me, it is far better to grasp the Universe as it really is than to persist in delusion, however satisfying and reassuring.*' - Carl Sagan

ABSTRACT

The study of low mass star formation in our local Galaxy is particularly suited to HERSCHEL. The SPIRE spectrometer and photometer aboard the spacecraft operate in the $\sim 200 - 600\mu m$ range and are well suited to probe the cold, dusty environments in molecular clouds where prestellar cores reside. The SPIRE FTS spectrometer is an interferometer, and this instrument design has strengths and weaknesses which are important to understand when using data from the instrument. Herschel is set to continue groundbreaking work in the infrared, building upon earlier work from ISO, IRAS, and SPITZER, probing deep into star forming regions and improving our knowledge of the processes within.

In this PhD thesis, we outline the current body of knowledge in low mass star formation. We examine the properties of the SPIRE FTS as a spectrometerusing a small, laboratory designed desktop FTS. We study the intrinsic properties of the instrument, as a way of understanding issues we are likely to see when using the SPIRE FTS in flight.

With these issues firmly in mind, we examine the creation and use of SLIDE - an interactive IDL-based tool for processing SPIRE FTS data. SLIDE can extract line and continuum information from SPIRE FTS SEDs. We outline the creation, testing and use of SLIDE and provide examples of the use of SLIDE in astronomy with some examples from the literature.

We then use the line information we extract from a variety of sources with the spectrometer, to examine how SED fitting from photometer data could be affected by line contamination. We simulate a wide range of greybodies with noise and line contamination and examine how SED fitting is affected. Our simulations conclude that line contamination is not enough to affect the recovery of temperature and spectral index β significantly.

Finally we use the information we have deduced to examine SPIRE FTS SEDs of L1689B – a prestellar core located in Ophiuchus. Our SED fitting of the core confirms that this core is starless with no internal heating source, and the spectral index profile over the core morphology is consistent with an increasing density of fractal aggregrate grains towards the centre. The increase in grain density and spectral index profile is also in agreement with previous CO depletion data. Fractal grain growth of this nature is consistent with dust grain models.

Contents

List of Figures

List of Tables

.

Chapter 1

Low Mass Star Formation

1.1 An introduction to low mass star formation

Developments in sub-mm and infrared astronomy have meant that the processes involved in low mass star formation are beginning to be understood. As star formation is believed to begin in dense molecular clouds, the process is hidden from view from conventional optical astronomy. The cold temperatures of the initial molecular cloud (at around 10K-20K) mean that the bulk of any emission is in the sub-mm and infra-red. Advances in ground-based and space-based infrared astronomy have allowed detailed observations of sites believed to be star-forming regions. These observations have allowed the development of star-formation models and led to classifications of observations into a chronological order. These classifications start from the initial higher density region in a molecular cloud and chart the evolution of the object into a pre-main sequence star. The earliest identifiable stage in the formation of a low-mass star is known as the prestellar core stage. The definition of a prestellar core is that it is the phase in which a gravitationally bound core has

formed in a molecular cloud, but no central hydrostatic protostellar object exists yet within the core, (Andre, Ward-Thompson, & Barsony, 2000; Ward-Thompson et al., 2007).

In this pre-stellar evolutionary phase, there is a central clump that is bound gravitationally. This clump evolves towards greater central density until it collapses into a proto-star. For this evolutionary phase, the infall of mass would be dictated by the central density and so an in-depth study of L1689B was seen as important in defining characteristics of the pre-stellar phase. This will be described in Chapter 6.

We shall now examine their detection and identification. We shall consider their physical properties, and examine how low mass star formation is believed to start within molecular clouds.

1.2 Prestellar and starless cores

Within a molecular cloud, we see cores that do not contain any protostars. These are called starless cores. The subset of starless cores that are gravitaionally bound are called prestellar cores. These are the earliest observed stage of low-mass star formation. A prestellar core has an increased density relative to the parent cloud, and an extent of approximately 0.1 pc, with a temperature of aound 10K and volume density $n(H_2) \sim 10^4 - 10^6\ cm^{-3}$ (Myers & Benson, 1983; Myers, Linke, & Benson, 1983). Prestellar cores are gravitationally bound and may be supported by magnetic fields, thermal support or turbulence, or some combination of these (Mouschovias, 1991; Bonnell et al., 1997; Nakano, 1998). Gravitational collapse of the core following an instability was modelled as early as Larson (1969) from

isothermal contraction leading to a protostellar object.

Early work on starless cores was led by Myers and coworkers throughout the 1980's. Myers et al. (1983) involved an initial ground-based survey of dark molecular clouds in ^{13}CO and $C^{18}O$ as possible sites of low mass star formation. This survey concentrated on local clouds (where local was defined as being within 500 parsecs), of small size (less than approximately a tenth of a parsec) and being cold and dense. Hence, T is $\sim 10K$ and gas number density is $\sim 3 \times 10^4 cm^{-3}$ for this sample. The line choices were made as they are tracers of high column density in the clouds, and the two isotopes of CO where chosen since in these clouds the ^{13}CO line is optically thick and the $C^{18}O$ line is optically thin, so the combination of the two should be very useful as a probe. The position of these dense cores was selected by using the Palomar Sky Atlas to identify such regions.

Ninety sources were studied by Myers et al. (1983) in total, seventy of which where within the 500 pc limit. They concluded that from this survey, the typical $C^{18}O$ source was found to be a third of a parsec in size with a mass of approximately 30 solar masses. Typical temperatures were around 10K. They also concluded that the observed sizes and temperatures were not consistent with equilibrium models, but were consistent with some collapse models. They suggested that turbulence was at least partly responsible for explaining the results. Over half of the observed sources showed non-Gaussian line shapes indicative of turbulence.

Myers & Benson (1983) described the results from mapping over one hundred similarly defined visually opaque regions, in NH_3. Here, the aim was to identify the cores closest to low-mass star formation. From these one hundred opaque regions, twenty seven dense cores were successfully mapped by Myers & Benson (1983). These were found to have mean properties of: Temperature $\sim 11K$, diameter

~ 0.1 pc and mass $\sim 4M_{\odot}$. Six of these sources had such strong lines that the line shape was examined. From this, they deduced that the line shapes were Gaussian and that thermal motions within the core were combined with either small-scale turbuence or collapse. In the case of collapse, it was deduced that the collapse was quite young and that star formation had not started within the core.

Properties of these cores were compared to a model of a spherical, isothermal nonuniform cloud, surrounded by a constant pressure medium. From this it was deduced that if the clouds where thermally supported only, then they are either collapsing or in an unstable equilibrium. In conclusion, they deduced that nearly all of the observed cores are due to form low-mass stars within a few free-fall times.

The properties of the cores considered by Myers et al. (1983), Benson & Myers (1989) and Myers & Benson (1983) were further studied by Beichman et al. (1986), who compared them with the IRAS all-sky survey to classify these regions. It was found that when comparing the ratios of cores with infrared counterparts to those without, the ratio was ~ 1 to 2. Beichman et al. (1986) concluded that the cores with an infrared but no optical counterpart either contained an embedded T Tauri star or an even younger object. As a result of the IRAS comparison, the cores were classified as either 'starless' or 'star-containing'.

At such low temperatures, starless cores are only visible from their sub-mm dust emission. Dependent on their parameters, these cores may disperse back into the cloud medium or become gravitationally bound. The cores that become bound via self-gravity become prestellar cores, and their evolution into main sequence stars is still not completely understood although much progress has been made in charting the progress. Prestellar cores are the first stage of the evolutionary process that forms low mass stars. For low mass star formation, di Francesco et al.

(2007) defines their mass as below $10M_{\odot}$ and we follow this definition. In this chapter, we shall examine their observational, dynamical and chemical properties as well as consider prestellar core models.

1.3 Prestellar core properties

The initial, starless core is believed to be the first stage for low mass star formation inside molecular clouds. Magnetic effects (Mouschovias, 1991), turbulence, dissipation (Nakano, 1998) and external influences (Bonnell et al., 1997) are all considered significant in the development and evolution of prestellar cores. In this section, we shall review the properties of prestellar cores with respect to these evolutionary effects.

1.3.1 Core modelling

In Section A.2 we derive an equation of state for an isothermal, spherically symmetric distribution of matter in hydrostatic equilibrium as

$$\frac{1}{r^2}\frac{d}{dr}\left[r^2\frac{d\ ln\ \rho(r)}{dr}\right] = -\frac{4\pi G\rho(r)}{a^2}. \tag{1.1}$$

Here r is the radius, $\rho(r)$ is the density at a radius r and a is the isothermal sound speed in the medium. We note that this equation is a second order differential equation, and as such needs two boundary conditions to solve. Choices for these boundary conditions could include:

- central density ρ_0 and the slope of the density (for example, a power law defining density with respect to radius)
- overall mass M and the external pressure in which the distribution is located P_{out}. This would define the pressure at the radius $P(r)$ of the sphere.
- Overall radius r_{tot} and the external pressure P_{tot}

In modelling a prestellar core in a molecular cloud we consider it a matter distribution located within a molecular cloud. In that case, we wish to consider the mass of the prestellar core dependent on the pressure of the cloud in which it is located. Ebert (1955) and Bonnor (1956) considered solutions of this equation of state and as a result the solutions are commonly referred to as Bonnor-Ebert spheres. A simple solution can be obtained by considering an infinite radius sphere, which we show in Section A.3 results in a singular solution where

$$\rho(r) = \rho_0 \left(\frac{r}{r_0}\right)^{-2} = \frac{a^2}{2\pi G r^2}, \tag{1.2}$$

and

$$M(r) = 2\frac{a^2 r}{G}. \tag{1.3}$$

This simpler solution was the one proposed by Shu (1977) as the initial state of the core prior to collapse. It is termed in the literature as a 'Singular Isothermal Sphere' as the solution has an infinite central density, even though the mass contained within any radius r is finite. Solutions where the central density is a finite value ρ_c in a medium of pressure P_{out} yield solutions where there is a two-tier structure,

where a central region of radius r_{flat} has an approximately flat radial density profile, bounded by a region where the density follows a profile similar to the Singular Isothermal Sphere solution. Figure 1.1 from Shu (1977) shows the radial density profiles for the Singular solution, as well as a number of other Bonnor-Ebert spheres with a finite central density.

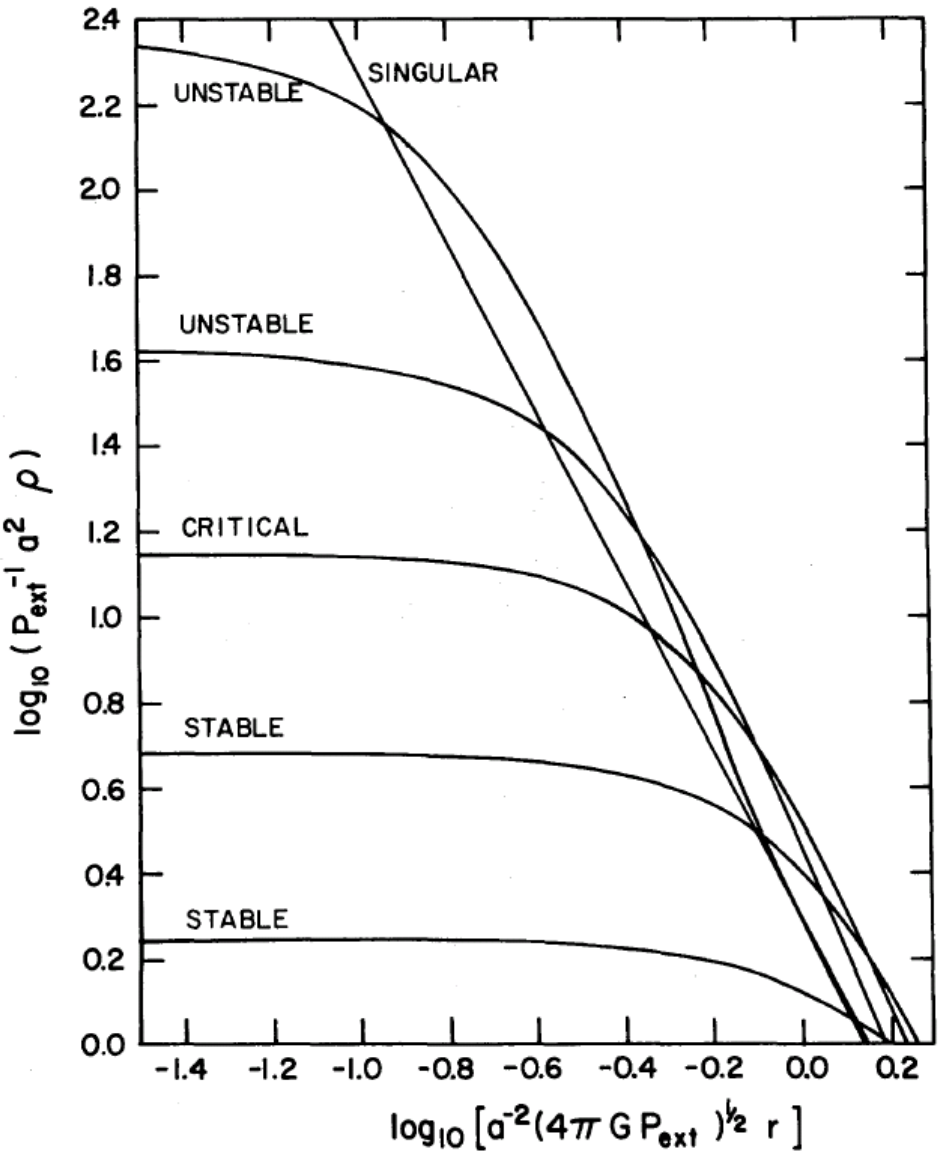

Figure 1.1: Radial density profile for various solutions of the Bonnor-Ebert sphere, from Shu (1977).

We note in Figure 1.1 that the line marked 'Singular' is the radial density profile for a Singular Isothermal Sphere, with a simple power law of order 2. The other lines reference a set of masses for the Bonnor-Ebert sphere. The line marked 'Critical' references a Bonnor-Ebert sphere whose mass is the maximum allowed for a stable configuration. This mass can be derived from

$$r_c = 0.49\sqrt{\frac{kT}{\bar{m}G\rho_0}} \equiv 0.49a\frac{1}{\sqrt{G\rho_0}} \tag{1.4}$$

and

$$\frac{M(r_c)}{r_c} = 2.4\frac{kT}{\bar{m}G} \equiv 2.4\frac{a^2}{G}. \tag{1.5}$$

$$\therefore M(r_c) = 2.4\frac{a^2 r(c)}{G}. \tag{1.6}$$

Here we note that Equations 1.4 and 1.5 are from Bonnor (1956). r_c is the critical radius for a Bonnor-Ebert sphere of a given temperature T, mean particle mass $\bar{m}$ and central density ρ_0. It should be noted here that this Bonnor-Ebert critical radius is comparable to the Jeans length

$$R_{Jeans} = a\sqrt{\frac{15}{4\pi G\rho}}, \tag{1.7}$$

which we define in Section A.4.2 of the Appendix. Combining the Equations 1.4 and 1.5 with the definition for the isothermal sound speed a and the central pressure $P_0 = \rho_0 a^2$ yields a maximum mass for a stable Bonnor-Ebert sphere of

$$M_{maxBE} = m_1 \frac{a^4}{P_0^{1/2} G^{3/2}} \equiv m_1 \frac{a^3}{\sqrt{G^3 \rho_0}}. \tag{1.8}$$

where $m_1 = 1.18$. This is shown in Figure 1.2. We also note this is comparable to the Jeans Mass, as in Section A.4.2 of the Appendix.

$$M_{Jeans} = \sqrt{\frac{375}{4\pi}} \frac{a^3}{\sqrt{G^3 \rho}}. \tag{1.9}$$

It is possible to express the Bonnor-Ebert solutions in dimensionless variables of mass and density , and this plot from Dullemond (2005) shows the maximum value graphically where the dimensionless mass is defined as

$$m \equiv \frac{\sqrt{PG^3} M}{a^4} \tag{1.10}$$

and the dimensionless density is simply the ratio of the central density to the edge density.

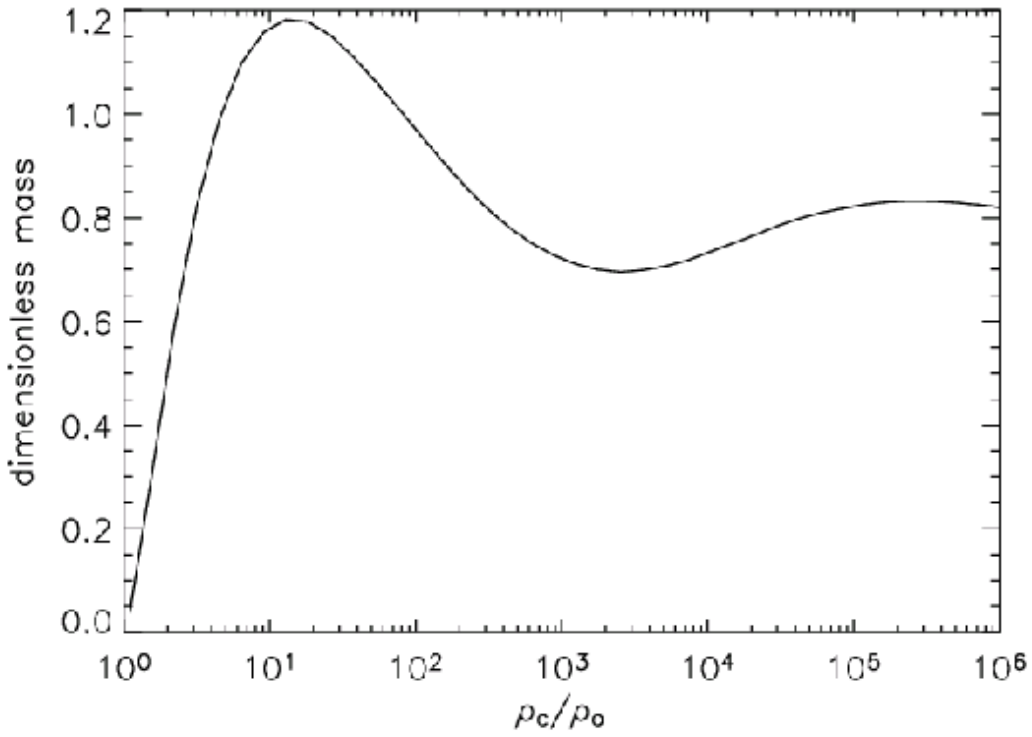

Figure 1.2: Plot of dimensionless mass and density variables, showing the critical mass from Dullemond (2005).

This leads us to two possible solutions for prestellar core structure. The singular solution has an infinite central density, and an infinite radius, although in practice the radius is taken to be that at which the density approximates that of the region in which it is located. The non-singular Bonnor-Ebert solutions have a finite central density and radius and have an external pressure from the molecular cloud as a boundary condition. However, the Bonnor-Ebert solutions are unstable if the external pressure or the core mass are above critical values. At some critical pressure, the central density increases and the radius decreases until gravitational collapse occurs. The solutions are also unstable if the centre to edge density ratio within the sphere is over 14.1 (Shu, 1977).

We note from Figure 1.1 that as the mass increases over this critical value, the radial density profile tends to approximate to that of the singular solution. As a result, Shu (1977) considered the singular solution as the one to use for modelling prestellar cores, as the unstable solutions approximate to the singular solution. However, as we shall see below, the two-tier density profile of the non-singular Bonnor-Ebert solution seems to be supported by observations of many prestellar cores.

Another model used for prestellar cores is the simpler Plummer model (Plummer, 1911). This model was originally chosen to model star density in globular clusters, and here we shall merely touch upon the mathematical properties of the model. The model is derived from considering the globular cluster as a ideal gas. We note that from Dejonghe (1987) the form of the density of this model in normalised units is simply

$$\rho(r) = \frac{3}{4\pi}(1 + r^2)^{-5/2} \tag{1.11}$$

and graphically the density distribution is shown in Figure 1.3. We note that it shows a flat central density profile, with a rapid dropoff of density over a critical radius. A Plummer-like model is used in modelling prestellar core density profiles, and Whitworth & Ward-Thompson (2001) use a model of the form

$$\rho(r, t = 0) = \rho_{flat} \left[\frac{R_{flat}}{\sqrt{R_{flat}^2 + r^2}} \right]^{\eta}. \tag{1.12}$$

where R_{flat} is the radius of the central, flat-density profile region, and η is a fit parameter such that $\eta = 5$ is a true Plummer model.

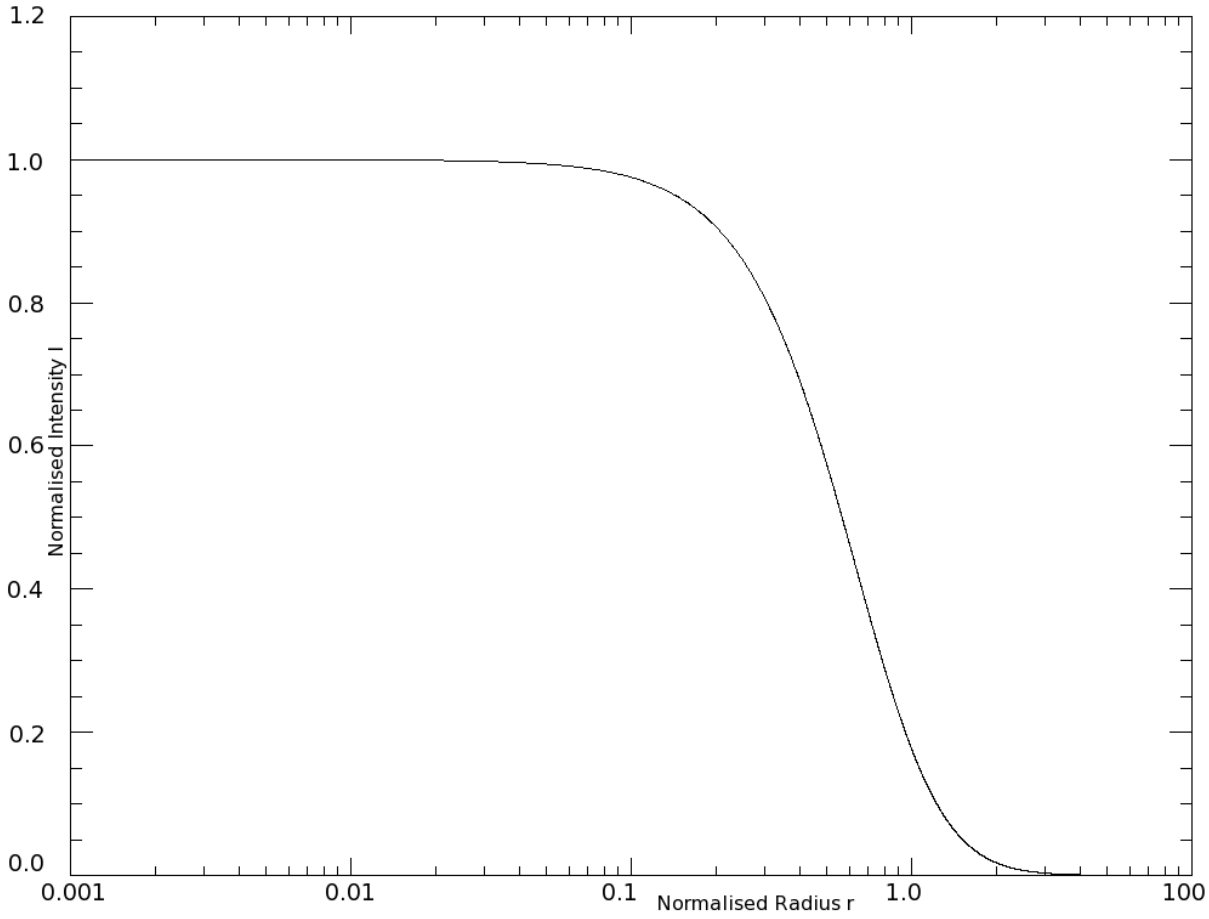

Figure 1.3: Two tier radial density profile structure of the Plummer model.

1.3.2 Core temperatures

Prestellar cores lack a central heating source, and so they are externally heated by the ambient Interstellar Radiation Field (ISRF). As a result, the core is warmer on the outside, and the cooler interior is shielded by the outer layers of the core. Core temperatures have been observed to be around 10K, for example Stamatellos, Whitworth, & Ward-Thompson (2007) concluded that core temperatures in the ρ Oph cloud were of the order of 10-11K. Myers et al. (1983),Myers & Benson (1983) obtained kinetic temperatures for CO, NH_3 repectively of this value, and Ward-

Thompson, André, & Kirk (2002) obtained dust temperatures from SED fitting of ISO observations of a similar value. Sub millimetre mapping of the core traces the properties of the dust within the core, and at the dust temperatures we seek the peak wavelength of emission is around 200 microns. As a result, the dust is optically thin to its own emission and so the core is cooler on the inside (Stamatellos et al., 2007).

1.3.3 Core sizes and morphology

Myers et al. (1991) detailed the studies of maps of sixteen cores in $NH_3, C^{18}O$ and CS from Benson & Myers (1989). This study found that non-circular core shapes were common. This suggested that the elongation is not a selection effect, or a result of particular environmental factors such as chemistry or excitation. Over half of the cores were found to be prolate, where their polar axis is greater in size than the equatorial axis. Elongation was seen in cores with and without embedded stars – suggesting this was a prior condition and not a result of star formation. The elongation also suggested that the available models at the time of an equilibrium between self-gravity and isotropic internal random motions did not fully explain the core, since in that case neither oblate or prolate was favoured over the other.

Ward-Thompson et al. (1994) concentrated on 21 of these starless cores, including L1689B, and studied them using the JCMT. This involved mapping the core central region in $^{13}CO\ J = 2 \rightarrow 1$. The cores selected had no infrared counterpart (Beichman et al., 1986), and from these cores a detailed survey in the sub-millimetre region was prepared. The aim was to identify peak positions, which would then be used for photometry of the local continuum and in some cases for continuum mapping. This involved the assumption that the CO peaks were repre-

sentative of column density peaks in the core, by assuming the cores are optically thin in CO. Timescales did not permit the mapping of large areas in a possibly more optically thin line.

The survey found that CO peaks did not necessarily coincide with continuum peaks, although the degree of seperation was slight. In some cases, there were more than one CO peak in the core but there was always the sole continuum peak. For each cloud, it was found that the continuum peak was different from the CO peaks. As dust emission is optically thin in this wavelength region, the paper concluded that the continuum peak represents the true peak in column density within each cloud. For each core, they deduced that the morphological structure was symmetric but non-spherical, with single or multiple peaks and some degree of elongation seen – a conclusion which agrees with Myers et al. (1991). 17 cores out of the sample of 21 showed a peak position in the continuum, with many of the cores showing multiple CO peak positions. The presence of a single continuum peak was confirmed by deeper continuum integration observations of the CO peaks of the seventeen cores. It was found that for each core there was a single peak that showed continuum emission.

It was found that if a core showed a detection in NH_3 then it was also detected by continuum emission (for all but one of the cores). The conclusion was that the underlying sub-millimetre continuum emission was just as reliable as NH_3 as a tracer for dense cores, and that there was no correlation between the NH_3 detections and the ^{13}CO clumps seen.

Core modelling performed by various authors such as Andre et al. (1996) and Ward-Thompson et al. (1999) for determining radial density structure assume ellipsoid shapes. Ward-Thompson et al. (1999) noted that when discussing models

of the prestellar core properties, that cores often show structure bordering on the filamentary when is difficult to explain. Goodwin, Ward-Thompson, & Whitworth (2002) noted from modelling core morphology that there is a difference in the morphology of starless and protostellar cores, with starless cores showing a flatter shape then protostellar cores.

1.3.4 Core mass and density

Core masses can be estimated using an integrated flux in the sub-millimetre by the following expression. This can be derived from radiative transfer and the derivation appears in the Appendix in Section A.1, and is

$$M_\nu = \frac{S_\nu^I D^2}{B_\nu(T_M)\kappa}. \tag{1.13}$$

M_ν is the core mass calculated using the integrated flux S_ν^I at a frequency ν. D is the distance to the core, $B_\nu(T_M)$ is the Planck blackbody function calculated at a temperature T_M for this frequency, and κ is the dust opacity value. This assumes the dust is optically thin. It is typical to assume that the core temperature is constant across the core. Typically the value of κ assumes a gas-to-dust mass ratio of 100, and its value at the relevant frequency is decided on a case-by-case basis in the literature by the author, dependent on how they interpret the dust within the core. The value is chosen by considering factors such as the core density, and whether grains are likely to have a mantle of ice.

Values of κ relevant to prestellar core models are discussed in a review by Henning, Michel, & Stognienko (1995). Observational and modelled opacity val-

ues for prestellar and protostellar cores give $\kappa \approx 1.1 g^{-1} cm^2$ at $\lambda = 1.3mm$, as recommended for dense cores by Ossenkopf & Henning (1994). This is relevant for silicate and carbon grains with ice mantles. Ossenkopf & Henning (1994) used coagulation models derived from simulations performed by Ossenkopf (1993) to tabulate a range of dust opacities at different frequencies, modelling physical conditions believed to be representative of protostellar cores. It is also noted here that this dust opacity is dependent on an ice mantle, and Ossenkopf & Henning (1994) warn that the dust opacity could be up to five times larger without ice. However, they recommend the lower value for warm discs. The caveat here is that an assumption in the model is that the core region is shielded from the interstellar radiation field (ISRF) and is undisturbed by local star formation. However, when considering prestellar cores, it is believed that this is the case, as the high density of the core shields the inner regions from the ambient ISRF. This results in colder temperatures within the core of around 10K, as there is no central luminosity source present.

Nutter, Ward-Thompson, & André (2006) used a value of κ for $850\mu m$ from Andre & Montmerle (1994) by rescaling the value for $1.3mm$ from Henning et al. (1995) for cloud envelopes. Setting $\beta = 1.5$ (Andre & Montmerle, 1994) they rescale the dust opacity by noting that it is proportional to ν^{β}. They also assume a gas-to-dust ratio of 100 to further rescale the value of κ. This results in their value for dust opacity being $\kappa_{850\mu m} = 0.01 g^{-1}\ cm^2$. Using this value, and a temperature of 12K for prestellar cores, the authors obtain prestellar masses of 0.1 - 16 $M_{\odot}$ for the L1689 region in Ophiuchus. Enoch et al. (2008) also derived their value of κ from Henning et al. (1995), assuming a value consistent with dust grains with thin ice mantles, with a similar gas-to-dust ratio. For prestellar cores in the Perseus region they derive masses between 0.3 - 4 $M_{\odot}$. These masses show

the range typical of prestellar cores.

Observational evidence for the radial density structure of prestellar cores suggests a two-tier density structure (Ward-Thompson et al., 1994). Up to a radius R_{flat} the radial density is approximately constant with a shallow slope, and greater than this radius the radial density is approximately $\propto r^{-2}$ (Ward-Thompson et al., 1994). This structure is not that of the Singular Isothermal Sphere suggested by Shu (1977) as the initial state prior to collapse. This profile has been found by emission (Ward-Thompson et al., 1994) and absorption (Bacmann et al., 2000) surveys of starless cores.

Emission studies assume that as the dust is optically thin, the sub-millimetre continuum emission is proportional to the mass of dust (Ward-Thompson et al., 1994). It is then possible to convert column densities measured in this method to radial density profiles by assuming the core is isothermal, and by assuming spherical symmetry.

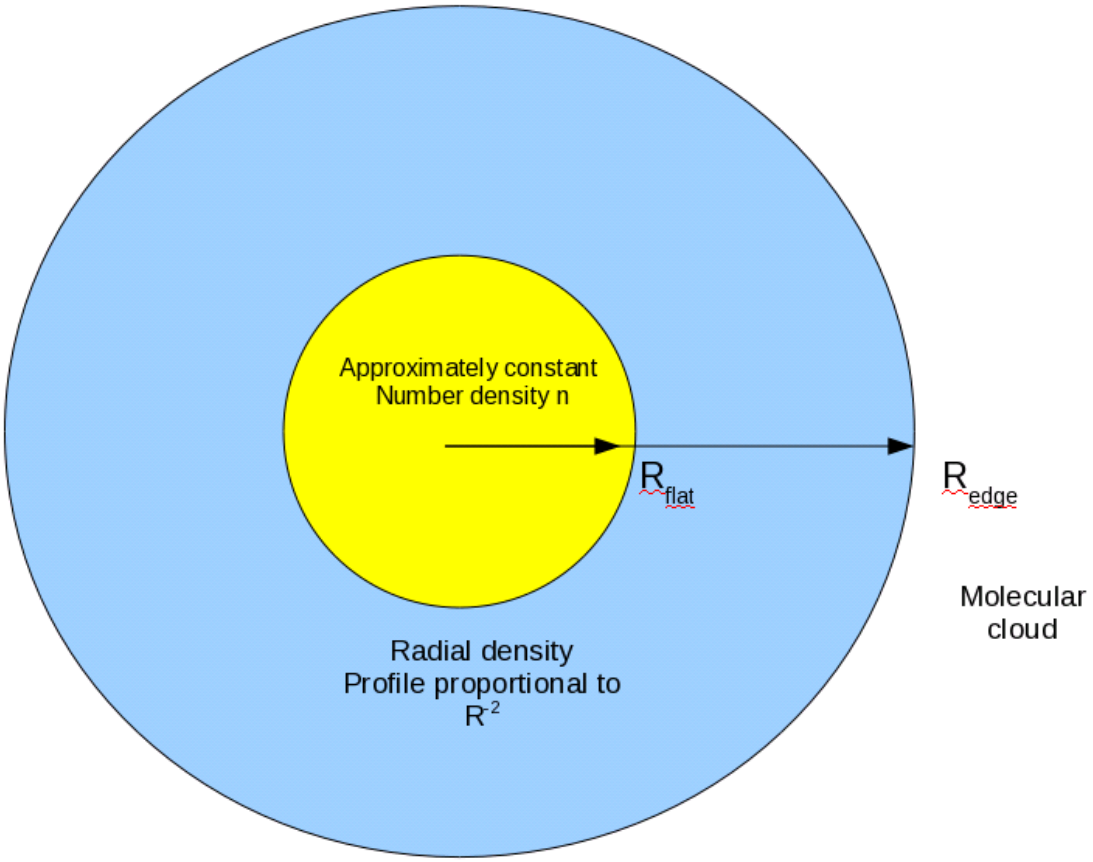

Figure 1.4: Two tier radial density profile structure for prestellar cores

Ward-Thompson et al. (1994) mapped four prestellar cores in continuum emission, and found that none were consistent with a single power-law relationship for the radial density profile. The inner regions of the core have a flatter density profile than the $\rho(r) \propto r^{-2}$ which would be seen in the Singular Isothermal Sphere. This paper suggests that the findings were consitent with earlier prestellar core observations, and that the increased resolution they were using made the detection of the structure possible. A graphical view of this two-tier structure is shown in Figure 1.4.

L1689B was examined by Andre et al. (1996), and is discussed in depth in Chapter 6. Since Ward-Thompson et al. (1994) had found that prestellar cores

were larger in size but comparable in mass to the envelope surrounding a protostar, it was reasonable to suggest that prestellar cores were the precursor stage. The protostar was believed to have started collapse. Here, the two tier structure was again detected. The authors generated a radial density profile of this core by averaging either circular or elliptical annuli over the core. For the elliptical annuli, an aspect ratio of 0.7 was used which was suggested by the mapping of the core. This was believed to represent a prolate spheroid inclined to the line of sight (Andre et al., 1996). A variety of artificial observations using models of the structure were generated, allowing for instrumental effects.

Andre et al. (1996) found a core structure for L1689B that is in essence identical to the conclusions of Ward-Thompson et al. (1994) – a two-component structure with a flat density profile in the central regions. Similarly, they concluded that this structure is consistent with a magnetically supported structure.

In this model, the initial state of the core is such that it is self gravitating. Contraction along the field lines rapidly leads to a balance between gravity and thermal pressure forces. Perpendicular to the field lines, contraction is slowed by magnetic forces. As ambipolar diffusion is related to the interaction between the intrinsic magnetic field and the number concentration of ions within the core, the timescale t_{ad} is related to the ionization fraction x. Nakano (1979) gives an estimate for the timescale as

$$t_{ad} = 2 \times 10^6 \times \frac{x}{10^{-7}}\ years. \tag{1.14}$$

The balance between gravitational contraction and magnetic support can be considered by the ratio of the mass of the core and the field strength resisting con-

traction. At some stage of the contraction, there is a critical value for this ratio, where the gravitational force is just enough to overcome the magnetic support. Mouschovias (1991) defines this value as

$$\left(\frac{M}{\Phi}\right)_{crit} = \frac{1.5}{\sqrt{63G}}, \tag{1.15}$$

where Φ is the magnetic flux providing support to the core. If a core exceeds this ratio it is considered magnetically 'supercritical' as it can no longer be supported by magnetic forces, although there may be an envelope of material surrounding the core that is still supported. Collapse of the core is nevertheless magnetically controlled. The morphology of the resultant structure is oblate with an inner core region of relatively flat density (Mouschovias, 1991).

Ward-Thompson et al. (1999) surveyed a number of prestellar cores with the JCMT. None of the cores were spherical. Comparing synthetic observations of models to actual observations, the two tier density structure was suggested. Figure 1.5 shows on each plot the dotted solid line as the radial profile. The dotted line represents a two-tier structure, with a flatter inner profile to the $\rho(r) \propto r^{-2}$ outer structure. The dashed line is a single power-law strucure, and the thin dashed line is the simulated model profile. It can be seen from this figure that the two-tier structure is the better fit. Here, the authors again discount magnetic support due to low ambient field strength in the regions.

Kirk, Ward-Thompson, & André (2005) mapped and considered a larger number of bright prestellar cores. Cores were considered bright if the flux density at $850\mu m$ was above a 10σ level. In all, thirteen cores were considered and simulated observations were compared to the real cores to ascertain density profile parame-

ters. The two-tier structure was found to be the best fit to the observational data and none of the cores were circular.

The synthetic profiles generated were compared to critical Bonnor Ebert spheres, and only one core was found to match. For a Bonnor-Ebert sphere, we can define a parameter $\xi = 2R_{flat}/R_{edge}$ where R_{flat} is the radius of the inner flatter density section of the core and R_{edge} is the outer radius. ξ has a critical maximum value 6.5 for Bonnor-Ebert spheres (Ebert, 1955; Bonnor, 1956). Above this critical value, the core is unstable. For all but one core, the value of ξ_{max} exceeded this value. This was also seen in the fitted temperature in the modelling, which exceeded the fitted temperature from the core SEDs.

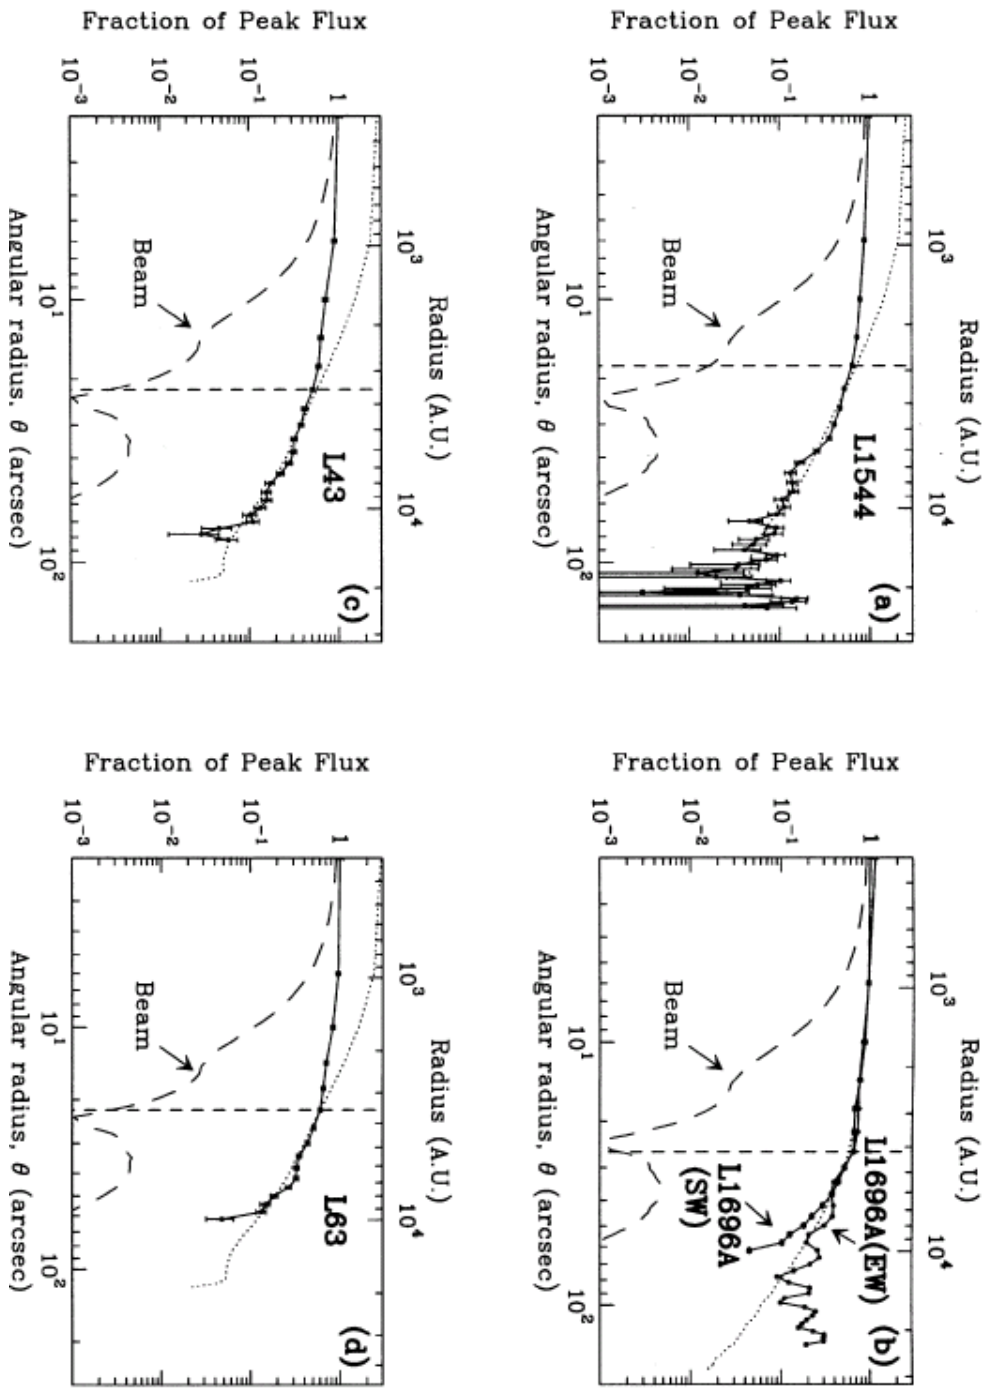

Figure 1.5: Results from radial density modelling from Ward-Thompson et al. (1999).

Bacmann et al. (2000) mapped prestellar cores in absorption against the diffuse mid-infrared background. Using ISOCAM, prestellar core profiles were mapped out to a radius of around $10,000AU$. This greater range is possible to due to increased sensitivity of this technique. As the dust emission is weak, the emission detection method is not sensitive to the outlying structure of the core. The absorption method also can derive the outlying column density of the material without any dust temperature assumptions. Of the 24 cores mapped using this technique, all show the two-tier structure proposed by Ward-Thompson et al. (1994). They also noted that three of the cores showed a sharp discontinuity between themselves and the ambient cloud. These cores are decoupled from the cloud and this may have repercussions on the amount of mass available for star formation.

From the initial 24 cores, 11 were identified with distinct compact absorption typically believed from a dense core. To this sample L1544 was added. Absorption observations were compared to IRAM-based dust emission observations of these cores and it was noted that peaks in mid-infrared absorption and dust emission were coincident. This suggests that the features are truly that of dense cores and not background fluctuations. Modelling of the prestellar core properties was performed to generate artificial observations. Modelling used a slightly more advanced structure, where the outer layer of the core was split into two regions. This split region had an inner region with a power law density profile, and an outer one with a steeper region that eventually merges into the cloud. This split region was used to include the cores which did not show a sharply defined edge, but instead merged into the cloud medium at some radius. The model follows a pressure bounded, Bonnor Ebert sphere. Their results confirmed the view of a two-tier structure, and noted that three of the cores seemed to be isolated from the local medium.

One result of note is that from their models, the contrast ratio of the density

between the centre and the edge exceeds the maximum contrast for a Bonnor Ebert sphere. This suggests that some support mechanism is in place otherwise they would be gravitationally unstable. They suggest that magnetic support and the ambipolar diffusion model is the most likely. They admit that a dynamic, turbulent magnetic field may need to be present in the outer layers of the core to explain the filamentary structure of some of the cores.

The effect of a composite structure to the density gradient of the core has repercussions for the evolutionary development of the prestellar core. For example, one must conisider how this would affect the acretion rate of material as the core evolves towards the main sequence. Bacmann et al. (2000) discuss that prior to core collapse, the radial density profile dictates the rate at which mass infall occurs. The outer density profile also dictates the amount of mass available to the forming star. If the outer density is too low, then this dictates when mass infall stops. This sets the final mass of the star.

We can examine suitable mass scales we would expect to see for prestellar cores, and one example is the Jeans mass (Jeans, 1928). This is derived in section A.4. This mass considers the balance between self-gravity and internal and external pressure to consider a length scale R_{Jeans} and a mass M_{Jeans}

$$R_{Jeans} = a\sqrt{\frac{15}{4\pi G\rho_0}} \tag{1.16}$$

and

$$M_{Jeans} = \frac{4}{3}\pi R_{Jeans}^2 \rho \approx 5.5\frac{a}{G^{3/2}\rho_0^{1/2}}. \tag{1.17}$$

This is the radius and mass for a spherical cloud to have roughly comparable contributions from self gravity, internal and external pressure. This results in a stable cloud, and this size condenses out of the medium. In considering stable masses of cores we have the Virial mass, as defined in Section A.4.3 of the Appendix

$$M_{vir} \approx \frac{5R\sigma^2}{aG}, \tag{1.18}$$

and

$$M_{maxBE} = m_1 \frac{a^4}{P_0^{1/2} G^{3/2}} \equiv m_1 \frac{a^3}{\sqrt{G^3 \rho}}, \tag{1.19}$$

the critical Bonnor-Ebert mass. It is therefore of interest to consider how prestellar core masses compare to these three measures of stability for a self-gravitating spherically symmetric mass in hydrostatic equilibrium. Given that they are stable masses for some idealised situations, we would expect for stable cores that their masses are all in the order of these three limits.

Ward-Thompson et al. (1994) compared core masses obtained from sub-millimetre dust emission with the virial mass. They calculate the virial mass using linewidths from CO (Myers et al., 1983) and NH_3 (Myers & Benson, 1983) observations. For five cores, all but one had a core mass greater than the virial mass, but when factoring in error bars they concluded that the core masses could best be described as comparable to the virial mass, suggesting they are gravitationally bound. Kirk et al. (2005) used a similar comparison of dust emission maps to NH_3 linewidth data from Myers & Benson (1983). They also generated synthethic Bonnor-Ebert sphere observations to compare against the density profiles they mapped for the

cores under examination. Although the virial masses were again comparable to the dust masses obtained, they found that the fitting of Bonnor-Ebert sphere data did not support a critical Bonnor-Ebert sphere interpretation for these cores. They found that the temperature of the sphere was at least a factor of two higher than the SED fitted temperature for the core. Significantly, they found that the parameter ξ was greater than its limiting value for this model of 6.5 (Ebert, 1955; Bonnor, 1956). This parameter is proportional to the radius of the flat central profile to the total radius.

Nutter et al. (2006) considered the L1689 region in Ophiuchus, comparing dust emission masses to the critical Bonnor-Ebert mass. Following the definition from Ward-Thompson et al. (2007) they considered a core in the region to be starless if its mass was less than the critical value, or prestellar if above. They found relative numbers of 13 to 3 for prestellar to starless cores, although masses were comparable.

Enoch et al. (2008) considered the largest sample size of 92 starless cores, with cores from Serpens, Ophiuchus and Persues. They used NH_3 surveys to infer an H_2 dispersion velocity and obtain a virial mass. Figure 1.6 is a plot of $1/\alpha \equiv M_{dust}/M_{virial}$ against M_{dust} for the 92 starless cores, and the 108 protostellar cores they also considered. The two dotted lines for $1/\alpha = 1, 0.5$ seperate cores that are in virial equilibrium and self gravitating – see Section A.4.3 of the Appendix. As shown, the vast majority of the starless cores considered were gravitationally bound, signifying they are prestellar (Ward-Thompson et al., 2007).

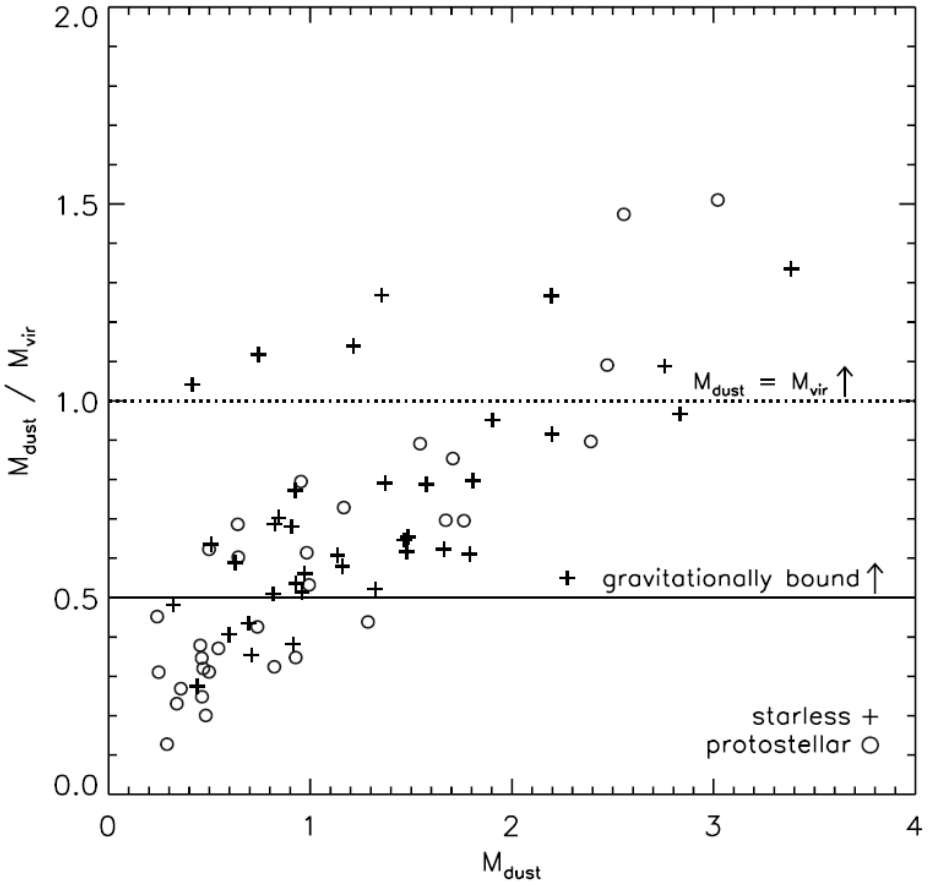

Figure 1.6: Comparison of Virial parameter to dust mass for 92 cores from Enoch et al. (2008).

Therefore, we see that, as expected, prestellar core masses tend to be of the order of being gravitationally bound. This is necessary if they are to ultimately evolve into stellar objects.

1.3.5 Core lifetimes

In Section A.4.1 of the Appendix, we define a freefall time for a core to collapse. This is for a core that is only affected by gravity, as it condenses out of the local

medium. As a result, if a prestellar core is in freefall collapse, the lifetime of the core must be in order of this freefall time

$$t_{ff} = \sqrt{\frac{\pi^2 R_0^3}{8GM}} \equiv \sqrt{\frac{3\pi}{32G\rho}}. \tag{1.20}$$

The observed timescales for prestellar core lifetimes are around two to five times greater than the freefall timescale (Andre et al., 2000; Ward-Thompson et al., 2007). This implies some additional support mechanism may be in place for prestellar cores.

Using the Benson & Myers (1989) observations of the cores, where half of the cores within 500 parsecs were detected in NH_3, Ward-Thompson et al. (1994) deduced that sub-millimetre continuum observations and ammonia detections were equivalent. Using this ratio, it was claimed that for the lifetime of the core with CO clumping half of the time would be spent with no detectable central clump. After that, the core would contain a submillimetre-continuum detectable central clump whose density they approximated as of the order of $10^6\ cm^{-3}$. This extended lifetime for CO clumping proved problematic for the Shu (1977) model. Using the clump density quoted, the free fall time for the clump is given by Equation 1.20 and is of the order of a few tens of thousands of years. By necessity, a mechanism for prolonging the lifetime of the clump is required. If we consider the ambipolar diffusion lifetime for a typical clump, we find that

$$t_{ad} \approx 8 \times \frac{R_{core}}{\delta v} \tag{1.21}$$

where δv is the velocity dispersion, of the order of 0.6 kms^{-1} (Shu, 1977). From

this a timescale of a few million years was concluded, which is consistent with the lifespan of prestellar evolution from core formation to the development of an infrared detectable source in the core. The Ward-Thompson et al. (1994) study concluded that magnetic support is present, and that the cores are in the ambipolar diffusion stage just before the protostellar phase. Examining the 20 cores without IRAS sources and finding a fifth of these had no evidence of CO clumping lead to the conclusion that a fifth of this lifetime is spent without CO clumping of any significance. The remainder of this time was spent with some CO clumping before an acreting proto-star is formed.

Beichman et al. (1986) compared a number of dense cores in star-forming regions, and used the ratio of starless cores to cores with IRAS sources to deduce the lifetime of the starless core phase relative to the lifetime of the cores with IRAS sources. By assigning a lifetime of the IRAS associated stage of 10^6 years (based on timescales for the T Tauri phase) they obtained a prestellar lifetime of also around 10^6 years.

Lee & Myers (1999) obtained a lifetime for prestellar cores based upon the ratio of cores observed with an embedded stellar object to those without. The ratio for their survey was 0.3 for a survey size of 406, and by applying the relationship that

$$\frac{N(embedded)}{N(starless)} \approx \frac{t_{protostar}}{t_{starless}} \tag{1.22}$$

together with a lifetime of the protostar stage of around $1 - 5 \times 10^5$ years, they obtained a value for the prestellar core lifetime of $0.3 - 1.6 \times 10^6$ years. This lifetime does not take into account intrinsic properties of the core such as central density, and is obtained purely from observational ratios. From their survey, they

obtained a mean density of $n \approx 7 \times 10^3\ cm^{-3}$ and this corresponds to a freefall time of around 4×10^5 years. The lifetime obtained by Lee & Myers (1999) is around one to four times the freefall timescale. They also found that their timescale for evolution was too fast to support ambipolar diffusion, and some other support method was needed.

Jessop & Ward-Thompson (2000) surveyed a catalogue of 60 cores, which were in medium opacity molecular clouds. These cores had lower volume densities than other samples, and they found a lower percentage had formed protostars. Using a similar method to Beichman et al. (1986) to calculate prestellar lifetimes on the ratio of cores with embedded sources, to those without, they obtained a lifetime of order 10^6 years. The authors then went on to compare their data with those from a selection of other papers – namely, Bourke, Hyland, & Robinson (1995) , Bourke et al. (1995), Clemens & Barvainis (1988), Clemens, Yun, & Heyer (1991), Wood, Myers, & Daugherty (1994), Myers & Benson (1983) and Myers et al. (1983). They found that the prestellar lifetime is shorter for higher volume densities. This means that increasing volume density leads to a faster collapse.

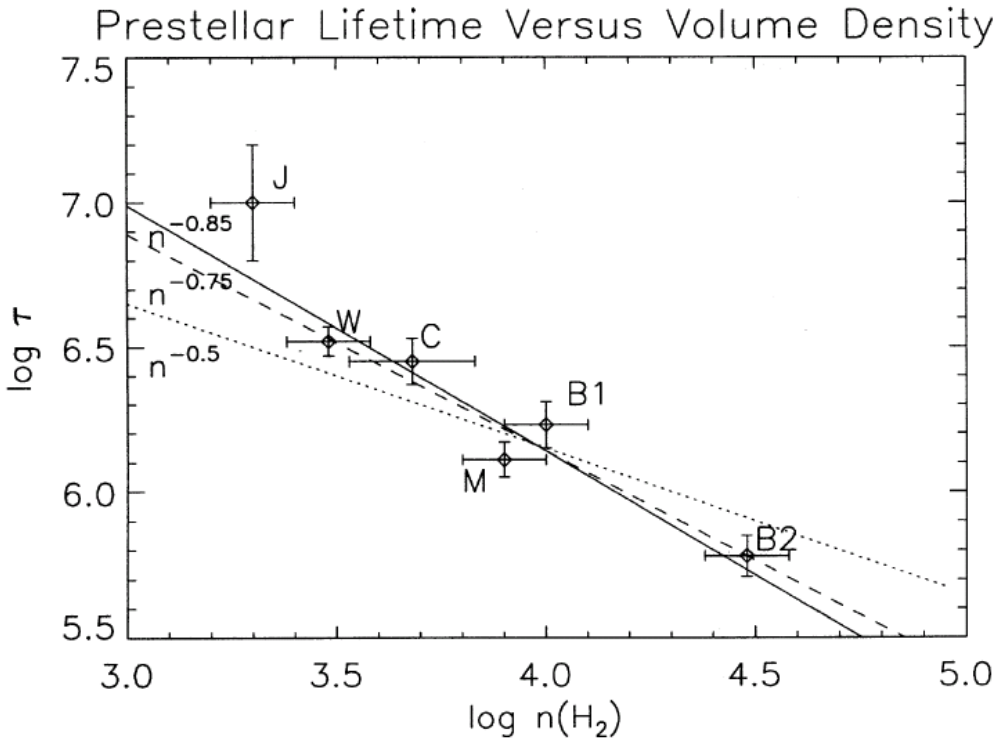

Figure 1.7: Comparison of prestellar lifetimes with volume density (Jessop & Ward-Thompson, 2000).

Figure 1.7, from Jessop & Ward-Thompson (2000), shows the prestellar lifetime plotted against the volume density. They note that it is possible to fit a power law profile to Figure 1.7 with a slope of -0.85.

Kirk et al. (2005) considered a number of cores surveyed using SCUBA, and used the ISO data from Ward-Thompson et al. (2002) to obtain prestellar core lifetimes. Of the cores surveyed, they considered 27 out of the 50 to be sub-millimetre detected if the SNR at $850\mu m$ was greater than 3σ. They categorised the cores as 'bright' or 'intermediate' depending on a flux density threshold of $170mJy\ beam^{-1}$ at this wavelength. Using the same method as Beichman et al.

(1986) they obtained a prestellar lifetime of 0.6×10^6 years for the cores. They then obtained lifetimes of $\approx 3 \times 10^5$ years for the 27 sub-mm detected cores, and a value of around half that for the bright phase. In both cases, this lifetime is a few times the free fall time for the relevant volume densities. They also noted that their volume densities and lifetimes seem to obey the same power law as seen by Jessop & Ward-Thompson (2000).

Crapsi et al. (2005) examined a number of cores using N_2H^+ and N_2D^+. Using statistical comparisons of their results, they obtained a prestellar lifetime of 6×10^5 years which is consistent with other timescales obtained. With lifetimes exceeding the freefall timescale, it is clear that some extra mechanism is providing support to these prestellar core structures. In the next section, we consider the possible mechanisms.

1.4 Core Dynamics

In this section we consider the dynamical processes within cores that shape the evolution of a prestellar core towards the protostellar phase. It is common to use the line profiles of molecular species such as CO to examine core dynamics – e.g. di Francesco et al. (2007). Examination of the line profile can show whether there is rotation, infall or outflow. Infall is revealed by a characteristic asymmetric double peaked profile (Shu, Adams, & Lizano, 1987).

Ballesteros-Paredes et al. (2007) suggest turbulence as an explanation for the formation of Molecular Clouds (MCs), and Giant Molecular Clouds (GMCs). GMCs are formed within the spiral arms of the Galaxy by gravitational interaction, and the smaller scale MCs are formed by turbulent compression, perhaps

by supernovae. Turbulence can lead to compression, and the formation of local density enhancements.

The presence of a magnetic field threading the core offers an extra level of support for material, and the long lived nature of ambient magnetic fields offer a compelling argument for this (Shu et al., 1987). The movement of ionised particles is limited by magnetic field lines. The ratio of the populations of ionised and neutral material affects core support. Magnetic support leads to an oblate morphology for the core as the support is primarily perpendicular to the field direction – see Figure 1.8. Neutral particles feel the effect of the field through particle collisions, hence the coupling between the neutral particles and the field is weak if the ionization fraction is low.

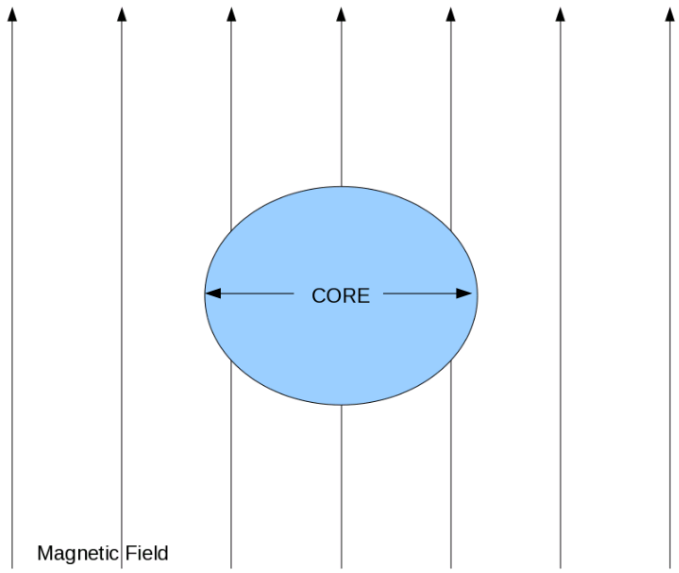

Figure 1.8: Magnetic support for a prestellar core.

Ambipolar diffusion is the process by which the coupling is broken. Neutral

particles fall to the centre and eventually the central gravitational potential energy becomes great enough to overcome the support and the remainder of the core can collapse. This was first considered by Mestel & Spitzer (1956). Eventually the increase in central core mass means a critical mass M_{crit} is achieved where the gravitational potential energy overcomes the magnetic support and the core becomes supercritical. This mass is dependent on the field strength and so it is more accurate to speak of a critical mass to flux ratio

$$\left(\frac{M}{\Phi_B}\right)_{crit} = \frac{0.53}{3\pi}\sqrt{\frac{5}{G}} \tag{1.23}$$

(Mouschovias & Spitzer, 1976; Mouschovias, 1991). The core is stable until the mass to flux ratio exceeds this critical value. This eventually forms a pre-Main Sequence star.

Cores are located in molecular clouds and so there is some shielding from the ambient Interstellar Radiation field. This decreases the ionisation fraction of the material and this allows the ambipolar diffusion process to occur on a timescale comparable to that which is observed. The ambipolar diffusion timescale is believed to be approximately ten to twenty times greater than the freefall timescale (Ciolek & Mouschovias, 1995).

Crapsi et al. (2005) studied 31 low mass starless cores in N_2H^+ and N_2D^+, to examine the extent of deuterium fractionation in the core. Studying these line transitions in detail allowed the authors to deduce a set of markers for identifying more evolved starless cores – assumed to be prestellar. They attempted to find statistical correlations between observed properties of the prestellar cores they studied, and they deduced the properties of evolved cores. These properties in-

cluded H_2 column density and R_{flat}. Their conclusions were that a combination of parameters could provide a means of identifying the more evolved cores within their sample.

Lee & Myers (2011) compiled a survey of starless cores comprised of a subset of the 306 cores optically selected by Lee & Myers (1999), comparing line and mapping observations in a variety of molecules (Lee, Myers, & Tafalla, 1999; Lee, Myers, & Plume, 2004; Sohn et al., 2007; Lee, Myers, & Tafalla, 2001). Line profile information was used to obtain a normalised velocity value

$$\delta v = \frac{v_{thick} - v_{thin}}{\Delta v_{thin}} \tag{1.24}$$

where v_x is the peak velocity of the thick or thin line respectively and Δv_{thin} is the full width at half maximum of the optically thin line profile. This parameter is negative for infall and these profiles were termed 'blue' profiles, with outflows being termed 'red' profiles. This is because in the case of infall, optically thick spectral lines will show a characteristic asymmetry in the form of a double peak, and for infall, the shorter wavelength peak will be stronger. For optically thin lines, the shape is a simple Gaussian.

1.5 Core Composition

We now consider the chemical properties of prestellar cores within a molecular cloud. Two important molecular constituents of prestellar cores are CO and NH_3. Initial surveys by Myers et al. (1983) and Myers & Benson (1983) mapped a number of dense cores in these molecules and obtained important information about their

physical properties. Ward-Thompson et al. (1994) used the linewidths from these observations to obtain core virial masses. NH_3 was used as a tracer for dense cores, rather than CO. This is because in comparison to the sub-millimetre continuum emission, there was a correlation between a detection for NH_3 and a detection by continuum emission. Myers et al. (1991) also mapped cores in in $NH_3, C^{18}O$ and CS and noted that the data enabled them to deduce non-circular shapes for dense cores, and they suggested that prolate cores were the norm.

N_2H^+, N_2D^+, CS and HCN data for a number of cores were studied by Crapsi et al. (2005) to deduce infall motion. Possible molecular differentiation within a dense core was outlined by di Francesco et al. (2007), and is shown graphically in Figure 1.9.

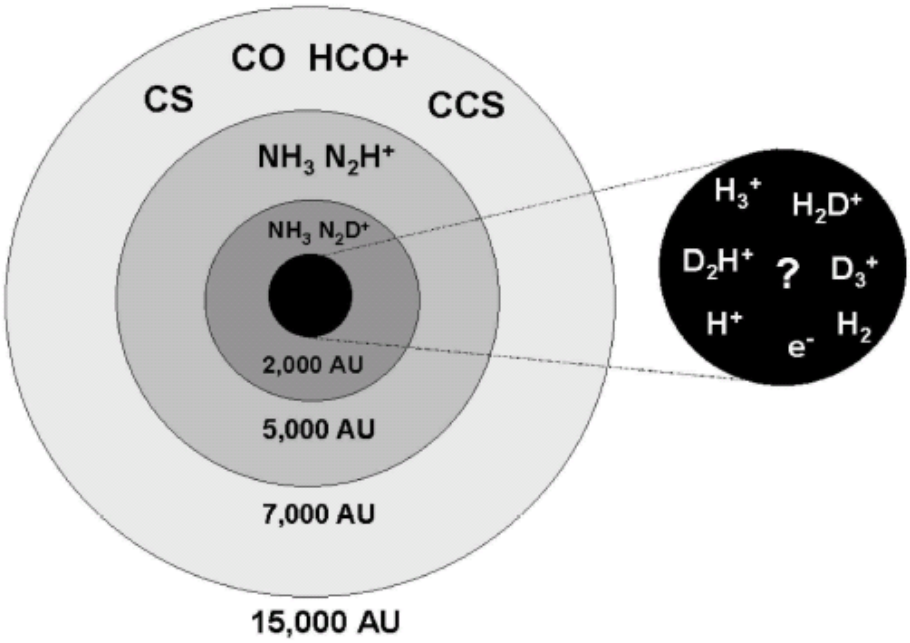

Figure 1.9: A typical dense core composition (di Francesco et al., 2007).

Deeper within the core, CO and CS are depleted as they freeze out onto dust grains (di Francesco et al., 2007). N-bearing molecules such as N_2H^+ are useful tracers of the higher densities expected deeper within the cores (Crapsi et al., 2005; Lee & Myers, 2011). Crapsi et al. (2005) also noted that deuterium fractionation is a marker for high density, and we see deuterium-bearing compounds in the central region of the core.

Caselli et al. (2003) studied the central region of L1544, and found that lighter species are dominant as other neutral molecules were frozen out in the high density central region. This was also supported by the work of Walmsley, Flower, & Pineau des Forêts (2004), which suggests that complete freeze-out within a core leaves H_3^+ and deuterated species as the main core tracer. Depletion is also dependent on the current age of the prestellar core, and so one would expect less depletion for younger objects. Molecules can be divided into 'early time' molecules where they are formed before carbon creates CO and 'late time' molecules which form slowly compared to this locking of carbon into CO (di Francesco et al., 2007).

Redman et al. (2002) explored CO depletion in the prestellar core L1689B, and this is discussed in more detail in Chapter 6. There was significant depletion within the central 5000AU. They concluded that 90% of the CO was depleted within this radius. This was surmised by comparing the H_2 density obtained by dust emission measurements, using a dust to gas ratio, and comparing the value obtained by tracing H_2. The discrepancy was used to obtain a depletion value. This depletion is expected for dense cores with a temperature less than 20K as CO is frozen onto dust grains (di Francesco et al., 2007).

Walmsley et al. (2004) noted that dust grain size is important to the ambipolar lifetime. With smaller grain sizes, the timescale diminishes as the ions are more

rapidly neutralised onto grain surfaces. In the completely depleted conditions they examined they also found that where the interation is between neutral gas and charged grains, the freefall timescale is increased by one to two orders of magnitude.

1.6 The Core Mass Function

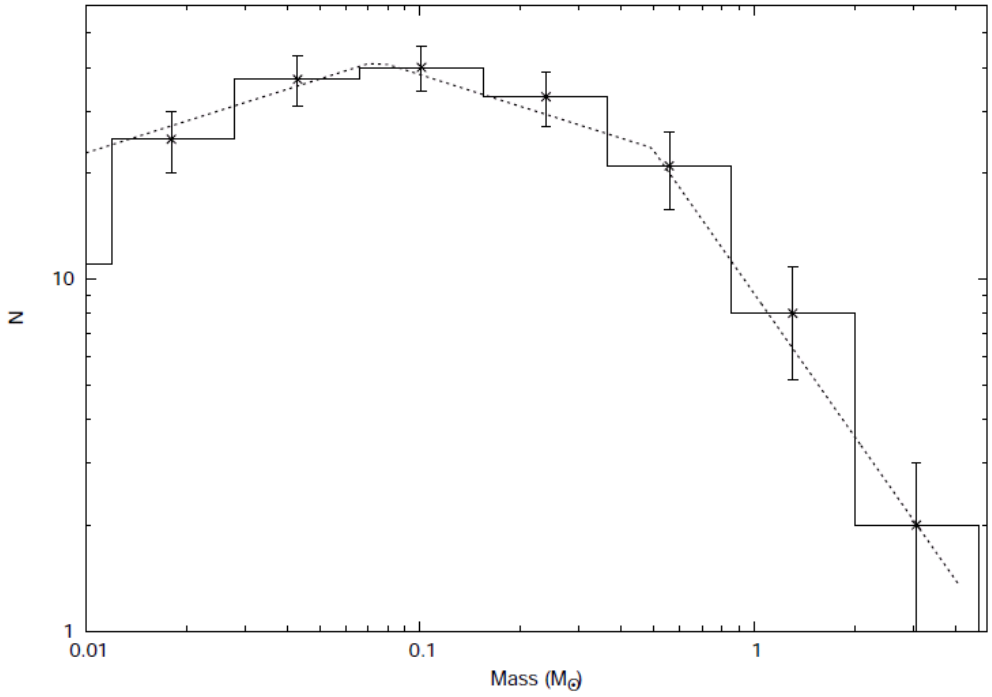

Figure 1.10: The CMF found by Simpson et al. (2008) with a three-component fit.

It is useful to consider the Core Mass Function (CMF). Analogous to the Stellar Initial Mass Function (IMF), the CMF is the function that describes the mass distribution of prestellar cores. The map of the Ophiuchus region from SCUBA archive data of Simpson et al. (2008) was used by Simpson et al. (2011) to explore the parameters of the CMF. Figure 1.10 shows their observed CMF. They fit their CMF with a power law of the form

$$M\frac{dN}{dM} \propto M^{-x}, \tag{1.25}$$

with the exponent x set to different values for different mass ranges. A full comparison with the Stellar IMF requires a correction to allow for the star formation efficiency ϵ of the cores, which they find has a best fit of $\epsilon = 0.2$ for a fxed value for each core. Allowing for a more complex model that allows for the mass of the core to determine the number of resultant stars, the fit is $\epsilon = 0.4$. This 'fully multiple star' model is outlined by Goodwin et al. (2008). In simple terms, there is defined a critical mass below which cores form double stars, and above this mass the result is a ratio of doubles and triples that is tailored for the particular model. In setting the critical mass, the star formation efficiency ϵ and the ratio of doubles to triples, one can refine the model to fit observed data (Simpson et al., 2008). Comparing the total prestellar mass in the Ophiuchus region studied, the star formation efficiency of the region to be approximately 2%.

1.7 From prestellar to protostellar cores

The formation of a star from a prestellar core requires an evolutionary path from the cold, dense gravitationally bound core towards a central hydrostatic protostellar object. Simpson et al. (2011) suggested an evolutionary path from a survey of 58 isolated cores, where infall was found by the presence of a double peaked, blue asymmetric line profile in HCO^+.

From the 58 cores surveyed, 14 showed a definite sign of infall and 10 showed possible signs of infall. The masses of these particular cores was either around, or

above, the Jeans mass required for a stable equilibrium. On their mapped data, the presence of an embedded infrared source detectable by Spitzer was taken to indicate a protostellar object rather than a prestellar core.

The evolutionary path suggested by Simpson et al. (2011) is shown in Figure 1.11, which is a radius-mass diagram of prestellar cores. The evolutionary path shown suggests that as an initial low mass, low radius prestellar core acretes mass from the local environment, its mass and radius increase. At a critical point, these values exceed the critical values for a stable prestellar core. The evolutionary line then drops showing the contraction of the core for a roughly constant mass. It is also interesting to note that all cores which show definite infall are above the critical Jeans mass. This evolutionary sequence is compelling, as it shows a definite progress from the gravitationally bound prestellar core towards the protostellar object. The diagram also shows two possible paths for 'low' or 'high' pressure for different values of the external pressure in the surrounding medium.

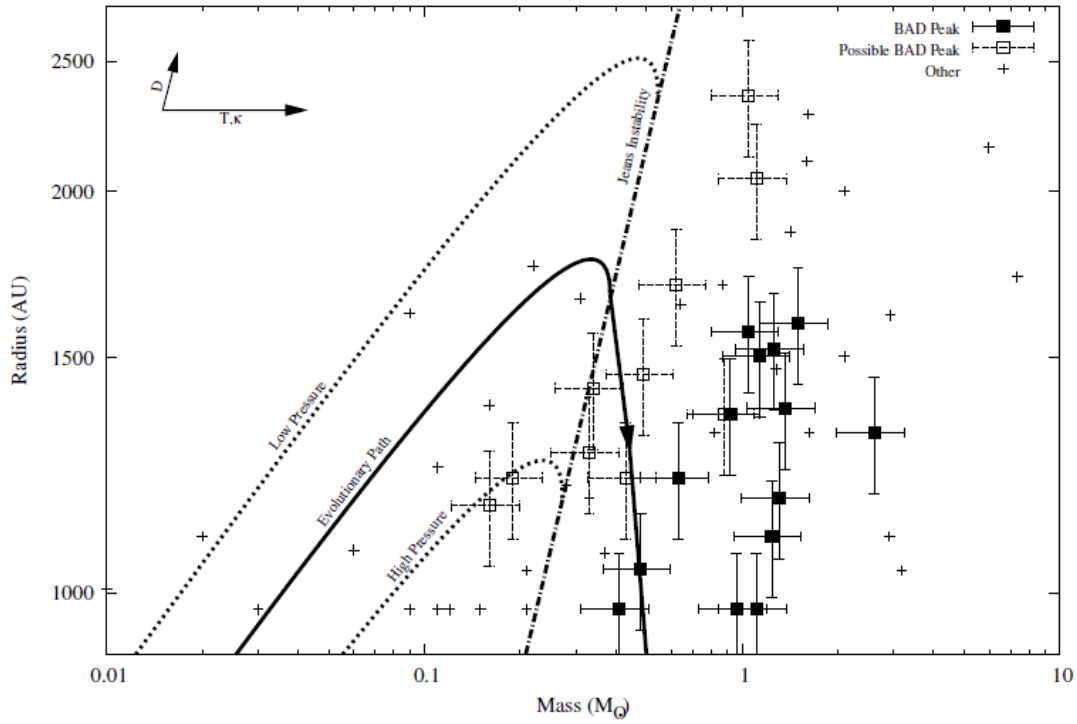

Figure 1.11: A suggested radius-mass prestellar evolutionary diagram (Simpson et al., 2011).

The transition from prestellar to protostellar is marked by the formation of a central, hydrostatic core. An acretion disk is formed and the total system of central object, envelope and acretion disk is termed a protostar.

1.8 Summary

In this chapter, we have outlined the process of low mass star formation through the characteristics of prestellar cores. From the definition of a prestellar core (Ward-Thompson et al., 1994) we have discussed core modelling for prestellar cores, through consideration of the density profile of the object. Properties such as tem-

perature, size and morphology have been examined. The mass and density of the prestellar core are determined from observations of the dust mass emission. Absorption studies have also allow study of density profile in some cases. The lifetime of a prestellar core has been deduced from statistical studies and various lifetimes have been considered in the literature.

We have also discussed the dynamical properties of starless cores, such as rotation and turbulence, as well as the effects of ambient magnetic fields on the evolution of the core. We have also considered the chemical composition of the core. From this we have seen how low mass star formation is affected by ambient conditions, and how these conditions are reflected in the composition and evolution of the core.

Chapter 2

Fourier Transform Spectroscopy

2.1 Introduction

2.1.1 Technical Discussion of FT Spectroscopy

The design of the Fourier Transform Spectrometer (FTS) is similar to that of the Michelson interferometer. It is called a Fourier Transform spectrometer as the output of the instrument is the Fourier Transform of the input spectrum $G(\sigma)$. The design of a basic FTS is shown in Figure 2.1. This is an idealised FTS as it has an infinitely thin beamsplitter. For a real FTS the beamsplitter surface is semi-reflective on one side. This means that one of the light paths will travel through more material then the other. To compensate for this, a piece of similar material without the semi-reflective surface is placed in the other optical path.

As shown in Figure 2.1 there are two mirrors within the system. One is fixed in position, the other is movable along a specifically defined range. Incident radiation

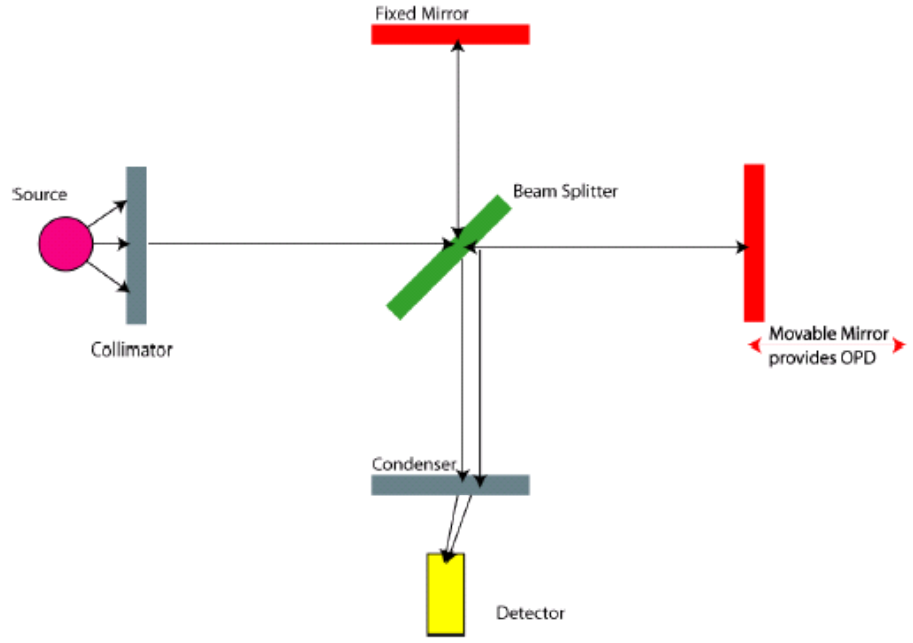

Figure 2.1: An idealised FTS.

falls upon the beamsplitter from a source, and by changing the position of the movable mirror an optical path difference (PD) is achieved, which we define as z. The radiation is recombined at the beamsplitter on the return path and then falls upon the detector, at which point an interferogram $F(z)$ is obtained.

The detector in this system can be a single-pixel (where the pixel is positioned on the optical axis of the system) or a multi-pixel camera. Multi-pixel spectrometers are more complex because one must consider the effect of off-axis rays within the system, with their differing path differences. This will be considered separately.

We can define the motion z of the movable mirror in terms of two parameters – the total range of movement available to the mirror, and the granularity of the movement of the mirror. We measure the total range from the zero path difference position to a maximum available path difference Δz. The granularity of the movement yields a step size we call δz. In the mathematical treatment of the

instrument, we shall show how these two parameters are important as they define the spectral range and resolution of the interferometer.

The FTS is a more complex design than a grating or Fabry-Perot. As the FTS consists of moving parts, a space-based FTS brings challenges in instrument design. In such a hostile environment the FTS must offer unique advantages to be chosen above simpler designs. Overall an FTS instrument can offer good spectral resolution, a high wavenumber accuracy and greatly reduced stray light or unwanted flux. Accurate movement and position of the movable mirror means that the interferogram output is known precisely – giving accuracy in the wavenumber output. Unwanted features in the interferogram can often be easily identified and removed.

Dependent on the design setup, an FTS can offer a large wavenumber range per scan (with a small optical mirror step size). However one early disadvantage of the FTS was the need for a computer to transform the output – but with the advent of cheap and reliable computer technology for performing Fourier Transforms this disadvantage is no longer important.

Comparing the output of an FTS to that of other spectrometer designs (such as the simple prism, or the grating): For a prism, the resolution is linked to the prism size. For the grating, the resolution depends on the slit width s, the focal length F of the collimator and the order angle θ. The maximum resolution for the grating in this case is given by Nm where N is a measure of the grating's grooves and m is the grating order number. For the FTS, the resolution is limited by the maximum optical path difference available to the system. When examining the prism, grating and FTS we find that the number of beams resolved in the system decreases from ∞ to N to 2 respectively but the energy throughput increases. For

the prism and grating spectrometers, a long and narrow slit is required and the area of this can never compare to the total input area possible for an FTS for the same resolution. As a result, the FTS has better sensitivity.

This helps define one of the advantages of an FTS, the Jacquinot Advantage. An FTS can have a large circular source as the input port with no resultant limitations on resolution. Compare this to the grating, whose resolution is linearly dependent on the slit width - and the detected power depends on the area squared of the slits. The Jacquinot Advantage is such that the FTS can collect larger amounts of energy at high resolution.

With a grating interferometer, the detector needs to be scanned across the image to get the spectrum. As each wavelength is scanned separately, then for N elements scanned for a time t for each element, we get a total time to scan the image of Nt. As the FTS observes all wavenumbers simultaneously, its integration time is simply t, and integrating for Nt gives an average gain in S/N of $\sqrt{(N)}$ - this is called the Multiplex or Fellget Advantage (Bland-Hawthorn & Cecil, 1997). This advantage fails if the limiting noise in the system is photon noise - noise governed by statistical fluctuations in the emitted number of photons. Here the advantage is nullified by a corresponding increase in photon noise by observing all elements together.

We can compare the throughput of a grating to an FTS (Bell, 1972) , and find the throughput of the interferometer to be:

$$E_m = A_m \Omega_m = 2\pi \frac{A_m}{R_m}, \tag{2.1}$$

where A_m is the area of the collimating mirror of the FTS, Ω_m is the solid angle of the collimating mirror and R_m is the instrument's resolving power.

We can compare this to the result obtained for a grating (Bell, 1972):

$$E_g \approx \frac{l'}{F}\frac{A_g}{R_g}, \tag{2.2}$$

where l' is the slit length and F is the focal length. Bell (1972) notes that F/l' is never less then thirty and as a result with the 2π factor in this example we see approximately two hundred times more power through the FTS than the grating.

2.1.2 Fourier Transforms

For a function $f(x)$ we can define the Fourier Transform $F(s)$ using the following equations, (see also Bracewell 1965):

$$F(s) = \int_{-\infty}^{\infty} f(x) \exp -2\pi i x s dx \tag{2.3}$$

$$f(x) = \int_{-\infty}^{\infty} F(s) \exp +2\pi i x s ds, \tag{2.4}$$

we can represent the Fourier Transform acting on the function $f(x)$ by the shorthand notation $\mathcal{F}f(x)$ and the inverse Fourier Transform as $\mathcal{F}^{-1}f(x)$. In the definitions we see that $F(s) = \mathcal{F}f(x)$ and that $f(x) = \mathcal{F}^{-1}F(s)$.

A function $f(x)$ can be Fourier transformed if it can satisfy three conditions:

1. Any discontinuties are finite in size and number;
2. $\int |f(x)|$ exists between $\pm\infty$ (i.e. it is absolutely integrable);
3. $f(x)$ has a bounded variation where there is not an infinite number of maxima or minima in a finite range (i.e. the derivatives must exist except at a finite number of points);

We can define any function $f(x)$ as the sum of even and odd functions – $F(x) = E(x) + O(x)$. Here we define $E(x)$ as the mean of the original function and its reflection in the y-axis, yielding $E(x) = 1/2[f(x) + f(-x)]$. Similarly we define $O(x)$ as the mean of the original function and its negative reflection $O(x) = 1/2[f(x) - f(-x)]$.

This leads us to the full form of the Fourier transform as above where:

$$\begin{aligned} \mathcal{F}_C E(x) &= \int_0^\infty E(x) \cos 2\pi x s dx \\ \mathcal{F}_S O(x) &= \int_0^\infty O(X) \sin 2\pi x s dx \\ \therefore \mathcal{F} f(x) &= \mathcal{F}_C E(x) - i\mathcal{F}_S O(x), \end{aligned} \tag{2.5}$$

these are known as the cosine and sine transforms, respectively. Thus we arrive again at equations 2.3 and 2.4.

The full derivation of the Fourier Transform and how we relate the transform of the interferogram of the instrument to the flux density at a given wavenumber is included in Appendix A.7. For a given wavenumber, the interferogram output

$I(z)$ can be Fourier transformed to give the flux spectrum $B(\sigma)$ where σ is the wavenumber. The wavenumber as a unit is defined as $1/\lambda$ such that to get the wavenumber in units of cm^{-1} one can use the formula:

$$\sigma(cm^{-1}) = \frac{10000}{\lambda(\mu m)}. \tag{2.6}$$

2.2 The FTS in practice

2.2.1 Practical instrumental considerations

The mathematical treatment of the Fourier transform is a mathematical idealisation that relies on integrating over infinite limits, with an infinitely narrow central field of view. As a result we must consider the modifications that finite limits impose.

2.2.1.1 A finite field of view

Using an FTS to observe an extended source, then we need to sum the flux contributions from the whole source. We do this by integrating over the solid angle Ω of the input source. We also need to consider the varying optical path across the field of view. For an oblique ray, we have an optical path difference given by:

$$z_{oblique} \equiv z \cos\theta, \tag{2.7}$$

where z is the on-axis path difference and θ is the angle the off axis ray makes with the optical axis of the system. Appendix A.6 gives the derivation for the modification term needed. Since this also affects the step size of the mirror $\delta z_{oblique}$ and the maximum path difference $\Delta z_{oblique}$ the effect is to broaden the line. This broadening is between σ_0 and $\sigma_0(1-\Omega/\pi)$ with a mean position of $\bar{\sigma} = \sigma_0(1-\Omega/4\pi)$, and the interferogram is multiplied by a sinc term.

We find $\sigma\Omega/2\pi$ is the broadened line width, and is a limiting factor on spectral resolution. This line broadening puts an upper limit on the resolving power of the instrument for a instrument with an associated solid angle Ω. The resolving power R_I is given by

$$R_I = \frac{\sigma}{\delta\sigma} \equiv \frac{\sigma}{\frac{\sigma\Omega}{2\pi}} \tag{2.8}$$

$$R_I = 2\frac{\pi}{\sigma} \tag{2.9}$$

which is the numerical definition of the Jacquinot advantage, and the sinc term that we apply to the interferogram defines the *Natural Apodisation* of the instrument. As a result,the effect of a finite field of view is to broaden linewidths and decrease resolution.

2.2.1.2 A finite path difference

Our detector must sample the interferogram in steps of PD such that we satisfy Nyquist sampling. This means that we choose our stepsize such that we sample at

least twice the highest frequency we expect to see. If we fail this condition, then we see 'aliasing' where some of the higher frequencies are 'folded' into our spectrum.

If we define z as the path difference, σ_{max} as the maximum wavenumber signal present, then by applying Nyquist sampling we get:

$$\delta z = \frac{1}{2\sigma_{max}}, \tag{2.10}$$

defining δz as the step-size of the path difference. The maximum PD of our instrument Δz defines the spectral resolution of the FTS such that we can only detect frequencies where there is an integer number of cycles between 0 and Δz. This means that the smallest determinable wavenumber interval is

$$\delta\sigma_{min} = \frac{1}{2\Delta z}. \tag{2.11}$$

and visually we can represent z, δz and Δz as in Figure 2.2.

If we consider a monochromatic source of wavenumber σ_1, scanned over an optical path difference of z we can represent the input source as a delta function. Fourier transforming this yields:

$$\mathcal{F}[B(\sigma)] = 2\cos 2\pi\sigma_1 z. \tag{2.12}$$

This is a cosinusoidal interferogram, which is a real, even function. This is to be expected as the Fourier transform of a real and even function is another real and even function. The behaviour of functions of particular type (real or imaginary,

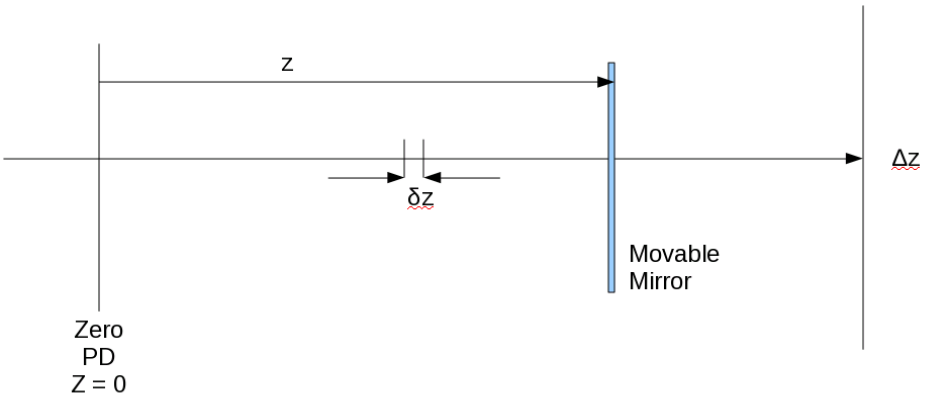

Figure 2.2: Defining terms for the movable mirror.

odd or even) is consistent upon a Fourier transform.

We can now consider the effect of finite scan length upon this. If we place an upper limit on our scan length Δz:

$$B(\sigma) = \int_{-\Delta z}^{\Delta z} 2\cos 2\pi i\sigma_1 z \exp -2\pi i\sigma \; dz, \tag{2.13}$$

this can be simplified before integration by expanding the exponential term, yielding:

$$B(\sigma) = \int_{-\Delta z}^{\Delta z} 2\cos 2\pi i(\sigma_1 + \sigma)z + \cos 2\pi(\sigma_1 - \sigma)z dz \tag{2.14}$$

$$\Rightarrow B(\sigma) = 2\Delta z \left(\frac{\sin 2\pi(\sigma_1 + \sigma)\Delta z}{2\pi(\sigma_1 + \sigma)\Delta z} + \frac{\sin 2\pi(\sigma_1 - \sigma)\Delta z}{2\pi(\sigma_1 - \sigma)\Delta z} \right). \tag{2.15}$$

We can disregard the first term in this result and obtain:

$$B(\sigma) \approx 2\Delta Z sinc\left(2\pi(\sigma_1 - \sigma)\Delta z\right). \tag{2.16}$$

This means that the interferogram from our finite FTS is in fact multiplied by a top hat function representing the truncation of the range of our interferogram. In Fourier space, this multiplication by a top hat corresponds to a convolution of the spectrum we are measuring with a sinc function.

Convolution for two functions is defined as:

$$h(x) = f(x) * g(x) \equiv \int_{-\infty}^{\infty} f(y)g(x-y)dy. \tag{2.17}$$

As this sinc function is an instrumental feature, it is termed the *Instrumental Line Shape* (ILS) function.

As a consequence of this, sidelobes are created on the spectrum of monochromatic lines as well as a widening of the spectral line. This results in a loss of resolution within the spectra, and so we seek a mathematical procedure to reduce this phenomenon.

The central peak of finite width is an adequate approximation to an infinitely narrow bandwidth line, but the sidelobes of the sinc function may appear to be false secondary sources. The line is also broadened from a delta function to a finitely wide sinc function, which gives a loss of resolution. The sidelobes are termed 'feet' and the process of reducing their effect on the output spectrum is termed 'Apodisation' – from the Greek, to remove the feet.

2.2.2 Apodisation and resolution

To minimise this effect, we multiply a taper function to the finite interferogram. This function is normalised such that at ZPD (the Zero Path Difference of the interferogram) the taper function is unity and is zero at the maximum path difference. As a result, the simplest of these taper functions is a normalised triangular function.

Multiplying the interferogram with an apodising function removes the side-

lobes in the spectrum that are a result of the limitations mentioned. We cannot completely remove either the finite width of the resultant line or the sidelobes, as they are an artefact of the limiting of theoretical infinities into experimental finite quantities. A resultant decrease in the sidelobe intensity must be compensated for by an increase in the width of the resultant line. As such, apodisation comes at the cost of decreased resolution – the full width at half maximum (FWHM) of the line is increased.

2.2.2.1 Instrumental line broadening

We have shown that the instrumental effect on a monochromatic line is to broaden it – the delta function is seen as a sinc in the output of the interferometer. This line broadening will obviously affect spectral resolution, defined in terms of the ability to separate closely-spaced features in the input spectrum. We can use the Rayleigh criterion to define spectral resolution in this case. Using this, we can define two features as being resolved such that the peak of the one feature lies at the first minimum of the second feature. This is shown graphically in Figure 2.3.

Considering the case of two monochromatic lines centred on σ_1 and σ_0, applying this criterion to the two ILS functions formed, we note that $\delta\sigma_n \equiv (\sigma_n - \sigma)$ for each line. If we define $S_n \equiv 2\pi\Delta z\delta\sigma_n$, such that the ILS for each line is proportional to $sinc\ S_n$, we find that the Rayleigh separation is defined as $\delta S = \pi$. Since $\delta S \equiv S_1 - S_0$, this gives us an equivalent resolution of:

$$\delta\sigma_{unapod} = \frac{1}{2\Delta z}. \tag{2.18}$$

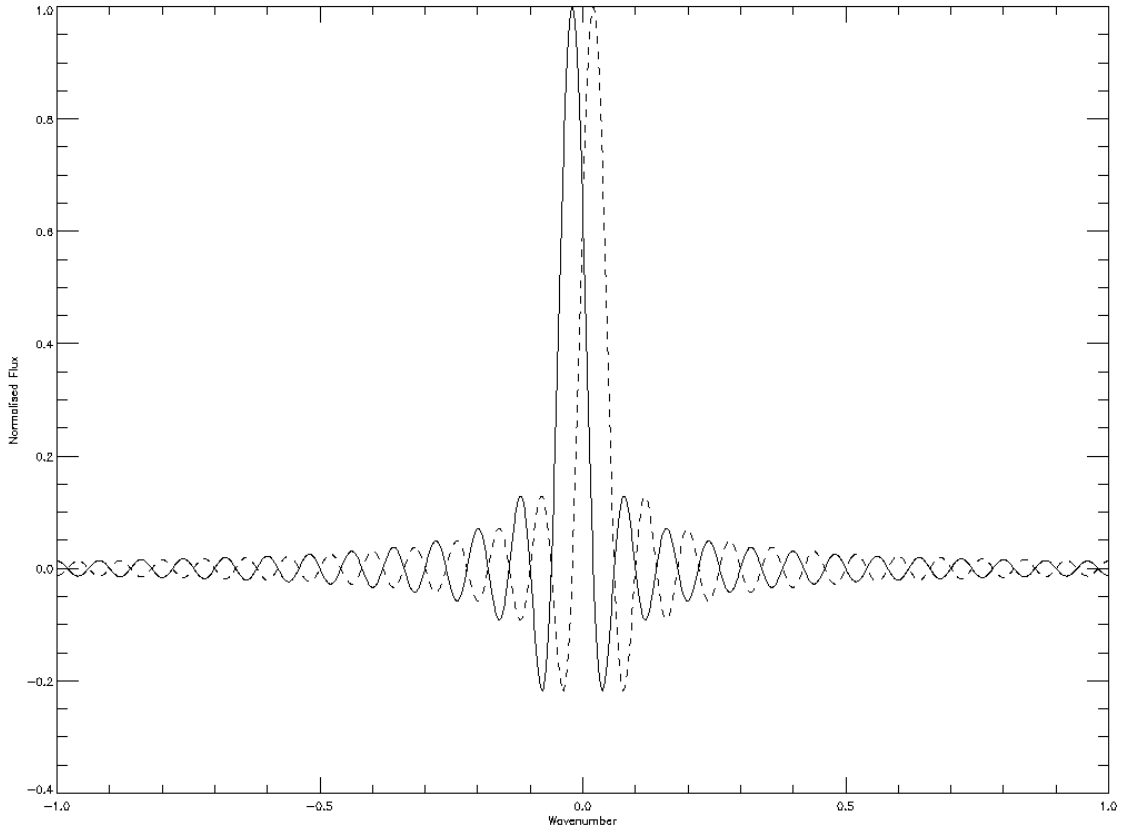

Figure 2.3: The Rayleigh criterion for two unapodised ILS functions.

As the ILS depends on the apodisation of the interferogram, applying triangular apodisation means the line profile is broadened, as this new ILS is given by $\Delta z sinc^2 \frac{S}{2}$. In this case, the Rayleigh criterion yields a separation of $\delta S = 2\pi$ and this becomes:

$$\delta\sigma_{apod} = \frac{1}{\Delta z}. \tag{2.19}$$

Sometimes the FWHM of the ILS is used to define the resolution, and for unapodised spectra this is given by:

$$FWHM_{ILS} = \frac{1.21}{2\Delta z}. \tag{2.20}$$

These measures are within 20% of each other and either can be used. For apodised spectra they agree to within 10%.

Apodising the interferogram with a normalised triangular function, yields a reduction in the sidelobes by a factor of four. The effect is to yield an output spectrum that is not defined by a *sinc* function, but rather a $sinc^2$ function. This reduction of sidelobe intensity is achieved at the expense of greater line width. Several studies (such as Filler 1964; Naylor 2007; Norton 1976) have detailed the search for what may be considered the 'optimal' apodisation formula.

2.2.2.2 Apodising functions

Filler (1964) contains a comprehensive discussion of apodisation and the initial searches for an optimal function. This paper compares some initial apodising

functions such as the top hat, triangular and $\sin^3$ functions, and later papers by Norton (1976) and Naylor (2007) expand on this. In this context, optimisation is defined as a balancing of the sidelobe intensity against line width. Here one is balancing spectral resolution against the possible false detection of spurious lines. This is obviously a greater problem as the line density in the spectrum increases, because constructive interference between nearby lines' sidelobes can generate strong spurious lines. However, for line fitting purposes the sidelobes can help constrain the fit of a sinc or sinc-convolved gaussian line profile. This will be expanded upon in Chapter 4 with regards to SLIDE, the *SPIRE Line IDentification and Extraction* program.

Filler (1964) attempts to define the quality of the apodising functions by considering the line half-widths and the largest lobe amplitudes of the functions. In this way, the user can choose the most appropriate apodising function for the purpose at hand. Norton (1976) continues this analysis, and studies the effects of over three thousand apodising functions to catalogue the line half-widths and lobe amplitudes they produce for a fixed initial input.

Figure 2.4 is a Filler Diagram, which illustrates that there is a definite boundary relationship between these two defining quantities. If we define W as the full width at half maximum and h as the first sidelobe amplitude, then the boundary relationship is of the form:

$$\log_{10}\left(\frac{h}{h_0}\right) \approx 1.939 - 1.401\left(\frac{W}{W_0}\right) - 0.597\left(\frac{W}{W_0}\right)^2. \tag{2.21}$$

For the total power under the spectral line, which is constant, Norton (1976) suggests a relationship such that:

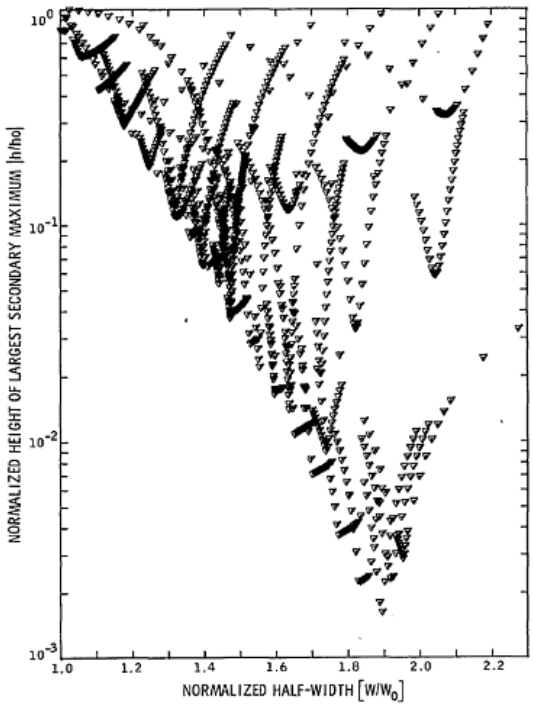

Figure 2.4: Filler diagram from Norton (1976).

$$P_{TOT} \propto \left(\frac{W}{W_0}\right)^{\alpha} \left(\frac{h}{h_0}\right)^{\beta}. \tag{2.22}$$

From this, it can be seen that apodisation involves a careful balance regarding spectral resolution and sidelobe production. A single optimum apodization function can not exist due to the fact that different FTS instruments have different limits and hence differently truncated interferograms. The use of a particular apodisation function will be dependent on the needs of the recovery process. The reduction of secondary maxima in the sidelobes is obtained at a cost of decreased spectral resolution.

2.2.3 Phase Correction

Just as the instrumental interferometer has finite limits for scan length, requiring a correction, there is also a finite scan step size. The interferogram is sampled in steps of finite size, and this may cause a problem, for example, where the point of zero path difference is between two sample steps. Without the zero path difference point sampled, we obtain an asymmetric interferogram. This results in a phase error ϕ in our generated spectrum. We can transform the inteferogram to get a complex spectrum $B_C(\sigma)$ which contains all the spectral information from our interferogram. However the offset, and resultant phase error, means that the spectral information is split amongst the real and imaginary plane of our resultant spectrum.

Should the interferogram be sampled symmetrically along this offset origin, then one can calculate $B_p(\sigma)$, the phase corrected spectrum, from $B_c(\sigma)$, the calculated spectrum, by:

$$|B_p(\sigma)| \equiv \sqrt{\Re(B_C(\sigma))^2 + \Im(B_C(\sigma))^2}. \tag{2.23}$$

Note also that $|B_c(\sigma)| \equiv |B(\sigma)|$. The magnitude of the calculated spectrum is a representation of the real spectrum, but has noise. In a general case, this phase can be corrected by taking a smaller two-sided interferogram and using that to calculate a phase correction term. This can then be applied to a larger one-sided interferogram. Using a larger one-sided interferogram gives better spectral resolution. The asymmetric interferogram has resulted in a complex spectrum, where the spectral information has been partially scattered into the complex plane. By

making an assumption that the phase error shows little variation in wavenumber, we can correct the phase error and obtain a true spectrum.

We can apply the phase correction term $\exp(-i\phi)$ to either the interferogram or the resultant spectrum. As the spectrum and interferogram are linked by a Fourier transform, the correction takes the form of either a multiplicative term or a convolution (see Griffiths 2007 for an review). This is because multiplying two functions in the frequency domain is equivalent to a convolution in the interferogram domain. This was seen in deriving the ILS of the FTS. In the interferogram domain, the ILS is the multiplcation of the infinite interferogram with a truncating top-hat function. The result in the frequency domain is the convolution of the input spectrum with a sinc function.

2.2.3.1 Forman method phase correction

The Forman method is performed in the interferogram domain, (see Sakai, Vanasse, & Forman 1968). Here we use a small, two-sided interferogram to generate the phase correction term. From this two-sided interferogram we generate a spectrum $B_2(\sigma)$ and from that we calculate:

$$\phi = \arctan\left(\frac{\Im(B_2(\sigma))}{\Re(B_2(\sigma))}\right). \tag{2.24}$$

The Inverse Fourier Transform of the phase angle term is then convolved with the one-sided interferogram $I_1(z)$ to generate a phase-corrected interferogram $I_p(z)$:

$$I_p(z) = I_1(z) * \mathcal{F}^{-1}\left(\exp(-i\phi)\right). \tag{2.25}$$

The result can then be Fourier transformed to produce a phase corrected spectrum.

2.2.3.2 Mertz method phase correction

The Mertz method is performed in the spectral domain, and is detailed by Mertz (1965). In this case, we interpolate the phase correction term $\exp(-i\phi)$ that was generated from the Fourier transform of the smaller two-sided interferogram, to the same resolution as the spectrum from the one-sided interferogram. ϕ is calculated as per the Forman method. As the phase angle is slowly varying, this interpolation does not cause a problem. We calculate the true spectrum $B_p(\sigma)$ from the one-sided spectrum $B_1(\sigma)$ by:

$$B_p(\sigma) = B_1(\sigma)\exp(-i\phi) \equiv \Re B_1(\sigma)\cos\phi + \Im B_1(\sigma)\sin\phi. \tag{2.26}$$

We keep only real terms in the expansion, as the true spectrum is real, and it can be checked during calculation that the imaginary components are zero. If not, we can deduce that there is an error in the phase calculation. The Mertz and Forman methods are mathematically identical – to see this, we note that $B_p(\sigma)$ is given by:

$$B_p(\sigma) = B_1(\sigma)\exp(-i\phi).$$

This gives us:

$$I_p(z) = I_1(z) * \mathcal{F}^{-1}\left(\exp(-i\phi)\right).$$

That is, Equation 2.26 is the inverse Fourier Transform of Equation 2.25.

2.2.3.3 Comparison of methods

It can be seen that computationally the Mertz method is simpler. Here the phase angle term is calculated from the smaller two-sided interferogram and the correction method is multiplicative. The Forman method relies on another inverse Fourier Transform and so in the past the Mertz method was seen to be preferable when computational resources were not so easily available. This is even more important in the case of multi-pixel systems where the number of calculations grows with the number of detector pixels.

Chase (1982) details a comparison of the two methods. In that paper, the author generated a synthetic spectrum with an infra-red SED with absorption bands at 1250, 1780 and $3450 cm^{-1}$, with an associated phase function. From this spectrum, an interferogram was obtained to which both phase correction methods could be applied and the recovery quality compared. This spectral form was chosen as the author explained that for a single absorption band, the dominant structure in the interferogram is due to the absorbed frequency. By choosing a multiband absorption spectrum, he ensured that he avoided the possibility of the phase errors being small at any single absorption band. The conclusion was that if errors of a few percent are acceptable, then the Mertz method is acceptable. If the interferometer can provide full double-sided scans, then the Mertz method provided excellent quality reduction.

In the next chapter we will briefly examine the use of FTS spectroscopy in astronomy, and expand upon the new challenges posed by multi-pixel systems using a small, lab-based multipixel FTS as an example. Using sample data, we

also compare the Mertz and Forman methods.

2.3 Summary

In this chapter, we have examined the basic theory behind the FTS. From a discussion of the idealised instrument, we have considered the advantages of the FTS over other instrument designs. A mathematical discussion of the instrument was presented. From the idealised instrument, we then compared how instrumental restrictions constrain the FTS. From these, we examined the observational limits such as spectral resolution and bandwidth, as well as the effects of a finite field of view on these defining characteristics.

In examining the effects of these limits, we noted that apodisation of the interferogram to remove sidelobe effects from the ILS is not a simple process. The choice of apodisation function used is important to balance the observational needs for each case.

Finally, we considered how another feature of a real FTS may cause observational issues. A finite scan size for the movable mirror need not necessarily sample the point of zero path difference. If this point is missed, then processing the resultant interferogram needs to consider the phase error introduced as a result. We considered two methods to correct for this phase error, with a short discussion on their relative merits.

In the next chapter,we expand further on the issues that affect a real FTS with the examination of a small laboratory built instrument.

Chapter 3

LoCI - a Low Cost Imaging FTS

3.1 Introduction – from single-pixel to multi-pixel spectroscopy

Herschel-SPIRE was designed to allow astronomers access to imaging spectroscopic data between $200\mu m$ and $600\mu m$, in two bands named the Spectrometer Long Wavelength (SLW) and the Spectrometer Short Wavelength (SSW) array, (Pilbratt et al. (2010) and Griffin et al. (2010)). In this range, the instrument has extragalactic and galactic applications in terms of early phases of star formation and the formation of galaxies at high z respectively.

The Herschel spacecraft was designed, built, tested, and launched under a contract to ESA managed by the Herschel/Planck Project team by an industrial consortium under the overall responsibility of the prime contractor Thales Alenia Space (Cannes), and including Astrium (Friedrichshafen) responsible for the payload module and for system testing at spacecraft level, Thales Alenia Space

(Turin) responsible for the service module, and Astrium (Toulouse) responsible for the telescope, with in excess of a hundred subcontractors.

SPIRE has been developed by a consortium of institutes led by Cardiff University (UK) and including Univ. Lethbridge (Canada); NAOC (China); CEA, LAM (France); IFSI, Univ. Padua (Italy); IAC (Spain); Stockholm Observatory (Sweden); Imperial College London, RAL, UCL-MSSL, UKATC, Univ. Sussex (UK); and Caltech, JPL, NHSC, Univ. Colorado (USA). This development has been supported by national funding agencies: CSA (Canada); NAOC (China); CEA, CNES, CNRS (France); ASI (Italy); MCINN (Spain); SNSB (Sweden); STFC (UK); and NASA (USA).

For extragalactic sources, SPIRE observes the redshifted emission from interstellar dust grains. Observations of a source's Spectral Energy Distribution (SED) allow the determination of the luminosity of the source. This provides important information in studying galactic evolution. In our own Galaxy, far-infrared and submillimetre observations of sites of star formation provide new information on the processes within. High resolution spectral maps can be used to determine bolometric luminosities of these regions, as well as to help determine the mass function. This can be supplemented by the study of faint, extended, emission regions. We can trace how the interstellar medium (ISM) is enriched by heavy elements as evolving stars eject material into the ISM (Griffin et al., 2006).

The design, construction and data analysis of a laboratory instrument is detailed below. This is a small-scale imaging laboratory FTS which gives the opportunity to examine the instrument in great detail, and to gain insight into the developmental issues of SPIRE. The instrument was named LoCI, the Low Cost Interferometer. Figure 3.1 shows a sketch of LoCI.

3.2 LoCI Development and Implementation

The development of LoCI involved an investigation of the practical and theoretical issues surrounding the construction of a $7 - 14\mu m$ laboratory-based small-scale instrument. The design of LoCI was constrained in a number of ways. A non-cryogenic detector was used to make the instrument portable. The detector also had a significant number of pixels across the array and had a fast data capture rate to help study multi-pixel spectroscopy. Components such as the beamsplitter and port windows were tuned to give optimal response over the desired frequency range. The components were procured prior to the start of this student's project (c. 2006) and so represent the best possible design choices at that date.

3.2.1 LoCI Design and Construction

The range and spectral resolution of an FTS is set by two operational parameters of the movable mirror. The total optical path difference available sets the spectral resolution. The step size of the movable mirror sets a limit on the highest frequency available (see Chapter 2). These parameters need to be matched to the detector. LoCi uses an INOCAM MIMIC II camera based upon a IRM160A bolometer, Ino (2004). This has an 160x120 pixel imaging array, optimised over the 8 - 12 μm wavelength range with a 30Hz frame capture rate. The camera has Germanium windows, and we can extend the workable range of the camera to approximately $7 - 14\mu m$ if we accept an approximate 80% transmission level at the extreme ends of the extended range. As an uncooled bolometric detector, it gives reliability and portability and this is a good match for the FTS. This wavelength range covers the black body peak wavelength for a temperature range of $200K$ to $400K$ - and

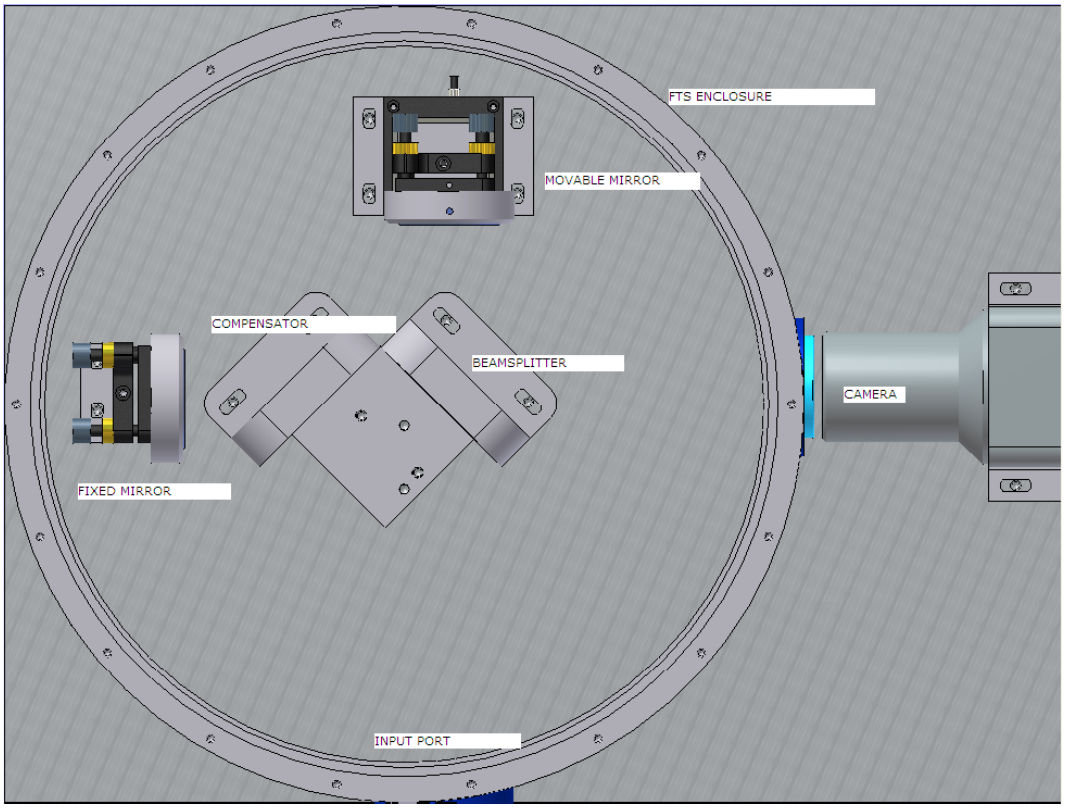

Figure 3.1: LoCI - the low-cost interferometer.

so adequately covers the room temperature range.

3.2.1.1 Mirror offset

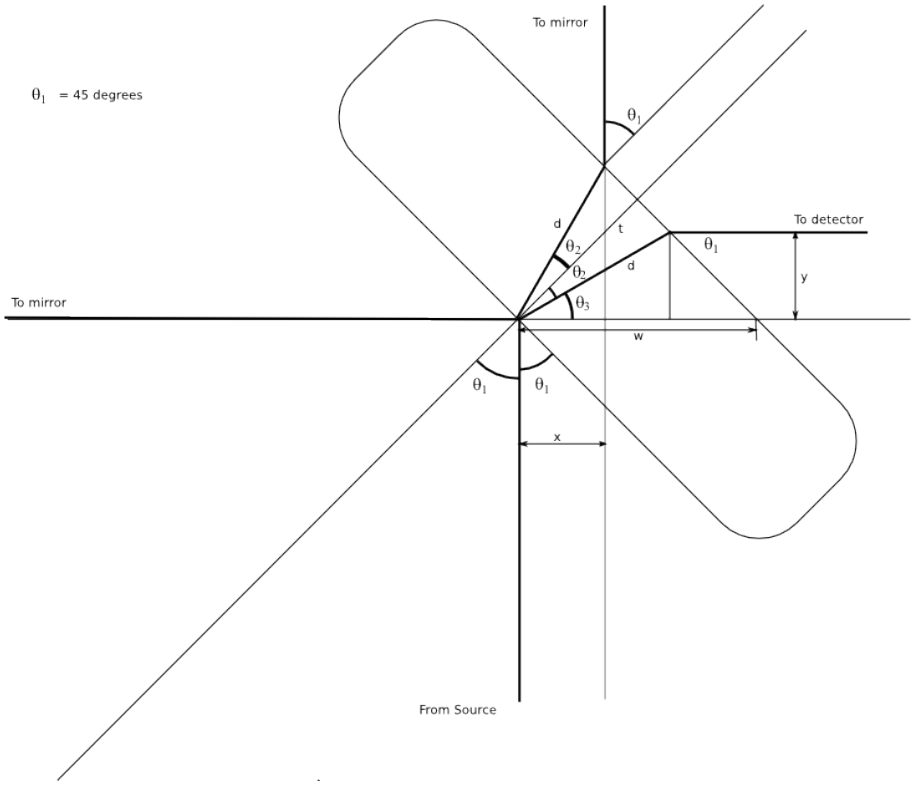

Figure 3.2: Offset correction.

The beamsplitter is not an infinitely thin element, but rather a finite thickness piece of optical glass with a semi-silvered surface. Therefore this extra optical material at 45^o to the optical path produces a slight translational offset in the rays that pass through the beamsplitter. This is shown in Figure 3.2. It is possible to trace the optical path through the system and see that, as light passes through the interferometer, one of the paths travels a lesser amount through the extra glass of the beamsplitter than the other. For the optical path which passes the fewer number of times through this glass, we can correct for this by placing a compensator at 90^o to the beamsplitter, as shown in Figure 3.1. Both compensator and

beamsplitter elements are comprised of $8mm$ thick blocks of glass, with the beamsplitter having a thin coating of Potassium Bromide (KBr) on one surface, which is the actual beamsplitting surface. This gave an acceptable degree of matching.

The addition of a compensator means that both paths will show slight translational offsets from the optical axis as light passes through the finite-width glass in the compensator and the beamsplitter elements. Rotation of the compensator by 90^0 relative to the beamsplitter element means that only the movable mirror needs to have an offset to its central position. This is because as the light travels the compensated path, the glass of the beamsplitter offsets the optical axis, then the glass of the compensator offsets it back onto the original optical axis. For this system, this means that the movable mirror only has to be offset by approximately $2.7mm$. Appendix A.8 contains a mathematical derivation of the offset.

The interferometer base and casing was machine-made in-house, and during the testing and construction phases was not fully assembled. All components were attached to the base, but the sealable enclosure was not yet mounted to the base in these phases. As the time of writing, there are still issues to surmount regarding vacuum-sealing the enclosure. Although it is possible to mount the camera outside the sealable unit, the movable mirror sits on a piezo-crystal with a thin cable carrying power and movement information. The design is still being worked on to seal this. Both input and output ports are comprised of optical port windows. Once the design is completed, these will consist of optically tuned glass. Zinc Selenide (ZnSe) offers a transmission efficiency of around 70% over our designed frequency range, as shown by the datasheet transmission diagram (Crystran Ltd 2010).The piezo controller has a range of movement of $500\mu m$ and is controlled by a PC over a serial connection to a control unit. This unit is controllable over a serial terminal via simple commands which move and interrogate the movable

mirror. It is capable of moving at a rate of approximately $60\mu ms^{-1}$

In a previous chapter we have shown that for unapodised output (see Equations 2.10 and 2.11):

$$\delta z = \frac{1}{2\sigma_{max}}, \tag{3.1}$$

and

$$\delta\sigma_{min} = \frac{1}{2\Delta z}. \tag{3.2}$$

The conversion from microns to wavenumbers is trivial, and for $\approx 7 - 14\mu m$ we find the wavenumber range is approximately $\approx 700 - 1400cm^{-1}$. Therefore:

$$\delta z = \frac{1}{2\sigma_{max}} = \frac{1}{2800}cm \approx 3\mu m. \tag{3.3}$$

This means that we need a $3 - \mu m$ step size for our movable mirror. A $500 - \mu m$ physical travel of the mirror equates to a $1 - mm$ total optical path difference yielding a possible maximum resolution for an unapodised spectrum of:

$$\delta\sigma_{min} = \frac{1}{2\Delta z} = \frac{1}{2 * 0.1cm} = 5cm^{-1}. \tag{3.4}$$

As shown above (see Section 2.2.2), apodisation will double this resolution to $10cm^{-1}$. Figure 3.3 shows the accuracy of the piezo controller movement over its full range of travel. We compared the movement that the piezo controller was

given to the reported position once the piezo had finished the movement. Figure 3.3 shows that the accuracy is better than a tenth of a micron for around 75% of the travel of the module (with the occasional error spike). The backwards scans show a slightly larger error, and overall the difference between forward and backward is shown in Figure 3.4 as about $0.01\mu m$. We require a step size of $1.5\mu m$ to give an optical step of around $3\mu m$. Hence the piezo controller movement module must be capable of an accuracy of $1\mu m$ or better. Hence, this satisfies our requirements.

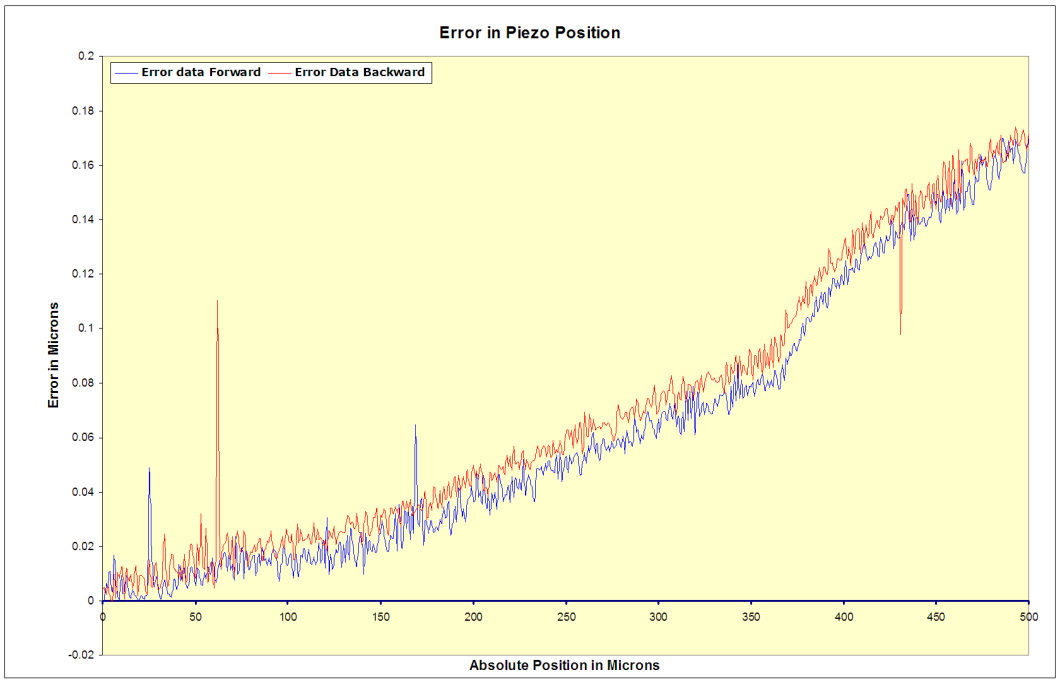

Figure 3.3: Piezo movement parameters.

In the last $\sim 25\%$ of the travel the error gradient increases sharply. If the travel range is restricted to the lower error gradient zone and we restrict our optical path difference to $0-700\mu m$ with the zero path difference position approximately $300\mu m$ then we have a usable range of optical path difference of $400\mu m$. The unapodised resolution becomes:

$$d\sigma_{min} = \frac{1}{\Delta z} \approx \frac{1}{2 * 0.04cm} \approx 12.5cm^{-1}. \tag{3.5}$$

Figure 3.4: Piezo controller movement parameters - forward and backward scans.

Apodisation doubles this to an approximate resolution of $\sim 25cm^{-1}$. This is still an acceptable resolution for this instrument, as it provides $\sim 50 - 100$ datapoints per pixel in the bandwidth of the detector. Figure 3.5 shows this graphically.

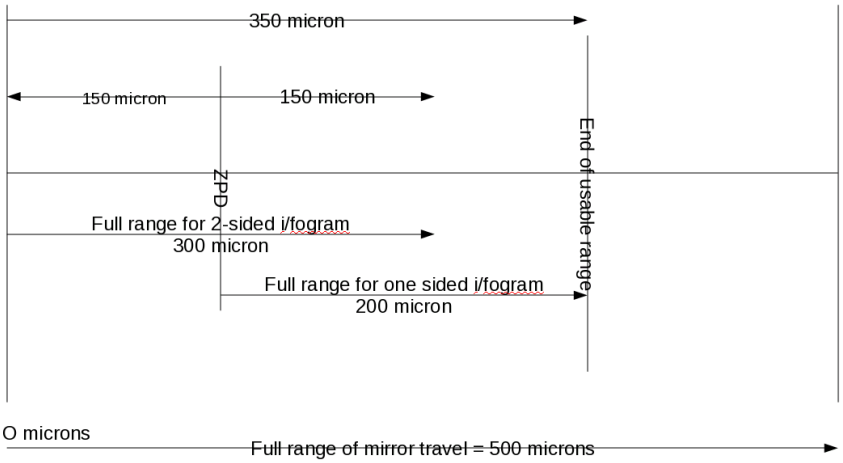

Figure 3.5: LoCI mirror travel schematic.

3.2.2 LoCI Optical Alignment

We can get a useful estimate of the zero path difference position by careful measurement of the optical path. We must allow for the fact that the beamsplitter has a finite thickness and that the beam splitting is performed on one surface only. Thanks to the compensator, both paths have equal path difference if the physical distance is the same. The compensator allows for the extra path through optical glass due to the non-zero width of the beamsplitter. The optical alignment of the FTS was performed using a standard procedure, aligning both mirrors such that a single image is present at the detector. Using a HeNe laser as an alignment source, and a screen placed in front of the detector, the mirrors can be aligned using optical fringes. The fringes produced from a HeNe laser (wavelength 632.8 nm) give us an alignment to an accuracy of around half a micron.

Alignment requires that there is zero PD between the two optical paths, and that the two images exactly coincide at the detector plane. If the mirrors are not perpendicular then the images will be inclined relative to each other. The optical axis is easily found by measuring the centre points of the mirror, and these positions were marked onto the base of the instrument. Once the optical axis is found, the laser was used to align the instrument. Examining the reflections from both mirrors, an aligned FTS will show them both recombined exactly, and only one image would be present. The mirrors were adjusted until the beams from both mirrors are superposed. If both mirrors are aligned in inclination, and there is approximately zero PD between the two then the laser should show fringes in the recombined image.

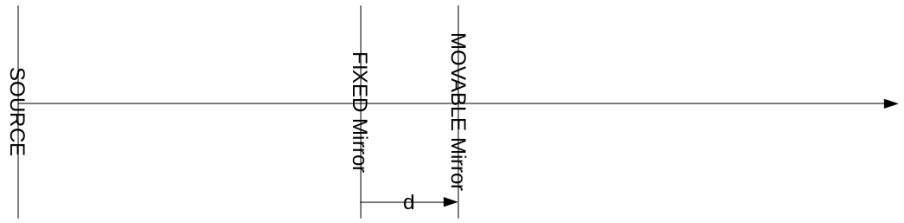

Figure 3.6: Parallel mirrors in the interferometer

Using a beam expander, we note the presence of *Haidinger fringes*, also known as *fringes of equal inclination*, (Hecht, 1998). As the incident light source is monochromatic, the Haidinger fringes are a demonstration of the interference pattern within an interferometer for parallel emergent beams. If we fold both optical paths into one dimension, then the path difference between the two mirrors is identical to a thick film of air of thickness $d = z/2$ where z is the optical path difference – see Figure 3.6. We can then consider each mirror as an optical boundary between the air film, and then Haidinger fringes are formed as a result of two-beam interference within this thickness of dielectric material. Each circular fringe is defined

by

$$2z(1 - \cos\theta_m) = m\lambda_0 \tag{3.6}$$

where λ_0 is the wavelength of the monochromatic source, m is the ring number (where zero is the centre ring) and z is the path difference. As the input light source is monochromatic, and the emergent beams are parallel, these fringes are virtual. They are formed at infinity and require a lens or beam expander to focus them into a real image. If the beam input is not monochromatic, these fringes can be seen without focussing.

Any misalignment causes the images from either mirror to be non-coincident and this produces *Fizeau Fringes* or *fringes of equal thickness*, where the inclination of the mirrors relative to each other forms a wedge-shaped air film – see Figure 3.7. These are fringes caused by interference from the images of the misaligned mirrors. These Fizeau fringes can appear as a pattern of parallel, straight fringes, although by changing the mirrors' relative inclination other shapes can appear. Careful manipulation of the mirrors' inclination allows the production of circular Haidinger fringes and shows the correct alignment of the interferometer mirrors has been achieved.

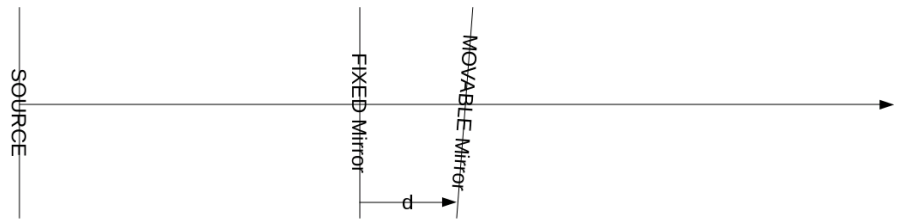

Figure 3.7: Non-parallel mirrors in the interferometer.

We then took a sample reading of a warm source that filled the field of view of the detector, in order to check that the pixel array was aligned. Figure 3.8 was generated by examining the zero PD position of every pixel over the array, where the shift of the zero PD position for each pixel (in data-points) was tabulated and plotted. There is a small angular misalignment in the camera array. The interpretational software was amended to calculate the zero PD position for each pixel as a result. This differential zero PD position for an image can be used to image the object as well as to test the effectiveness of the routine, as shown in Figures 3.9 and 3.10. Figure 3.10 clearly shows the pixel extent that is imaged through the instrument, and we expand upon this in the next section.

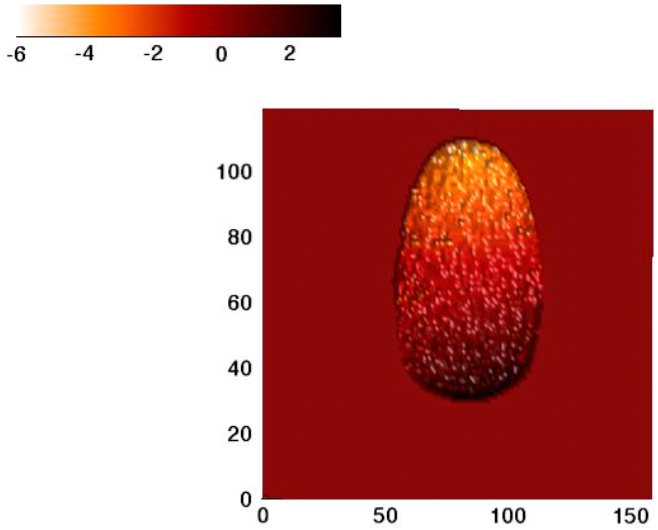

Figure 3.8: Testing the alignment of the pixel array.

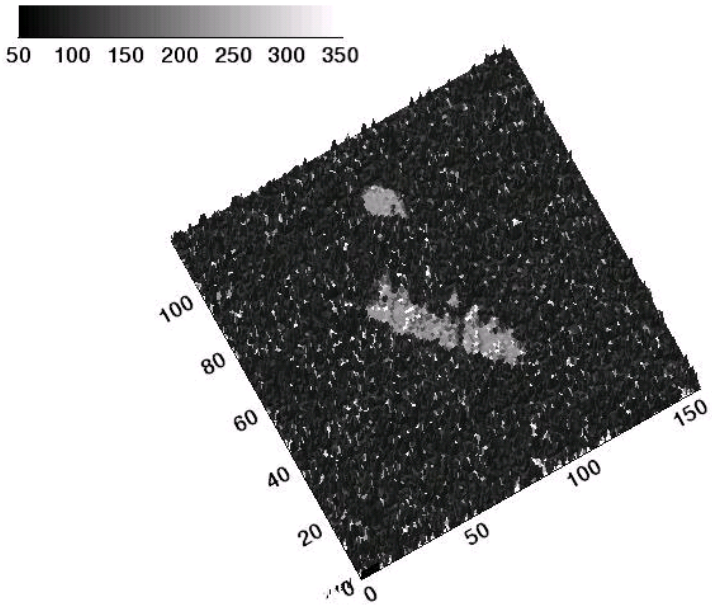

Figure 3.9: Zero PD position for a image of an object - soldering iron.

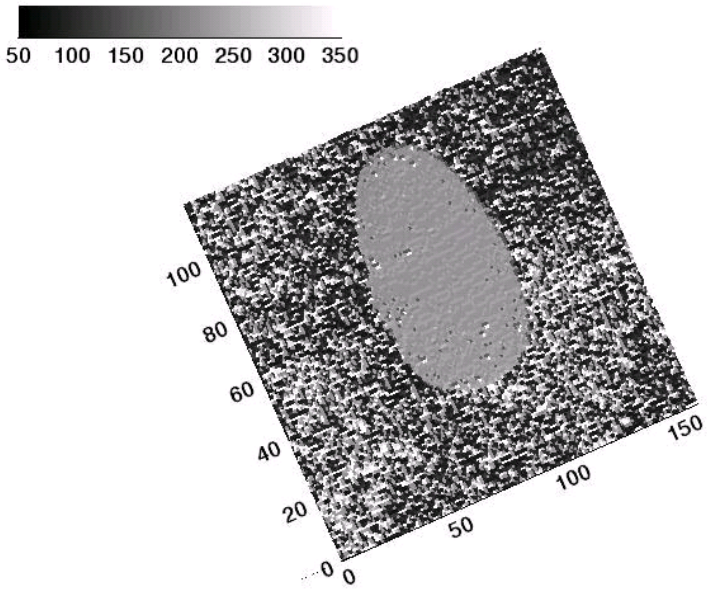

Figure 3.10: Zero PD position for a image of an object - full field hot body.

3.2.3 LoCI multi-pixel imaging

LoCi uses a multi-pixel array for an imaging device. The pixel array is 160 by 120 pixels and is centred such that the centre of the image array lies on the optical axis. As each pixel has a finite field of view, and there are pixels that lie off the optical axis we need to take into account the consequences of an extended field of view for each pixel, and off-axis pixels when examining the output from the instrument.

3.2.3.1 Rectangular pixel ILS for the detector array

Earlier we saw that for a single pixel FTS, an extended field of view for the on-axis pixel adds a sinc term to the ILS for the instrument. For off-axis pixels, we also need to consider the effect of differing path differences for the rays that are off-axis.

Using the analytic equations for the ILS for an off-axis rectangular detector (Genest & Tremblay 1999), we examined the consequence on line position and broadening for our detector array for all pixels for a sample input - in this case, a monochromatic line at 1000 cm^{-1}. Here we assume the optical axis is in the centre of the detector array, and by symmetry we can therefore calculate the ILS for an 80x60 array (one quadrant) and flip the ILS array for all four quadrants. This means the optical axis does not lie in any pixel, but rather at the intersection of four central pixels. Using these analytic forms, we see the evolution of the ILS as we move outwards towards the edge of the detector array.

Figures 3.11 and 3.12 show the ILS's evolution from a central pixel, to the furthest pixel from the optical axis. Figure 3.11 shows that for the central four pixels neighbouring the optical axis that the normalised ILS function has an amplitude of around 25%, and a maximum width of less then a ten-thousandth of a

wavenumber centred on the 1000 cm^{-1} position. Figure 3.12 shows that for the furthest edge pixel there is a reduction in the normalised amplitude to approximately 0.4%, with a maximum width of 0.25 cm^{-1} and a centroid shift position of approximately 5 cm^{-1}. Starting with a narrow, high delta we evolve towards a wide, flat trapezoid and as we can see energy is conserved as the normalised flux underneath these ILS's is conserved. We also note that, as expected, there is a shift in the line centre seen as we move outwards. The shift is a maximum at the detector edges and is around 5 wavenumbers. However at this extreme edge the normalised amplitude of the ILS is very much reduced to 0.4%.

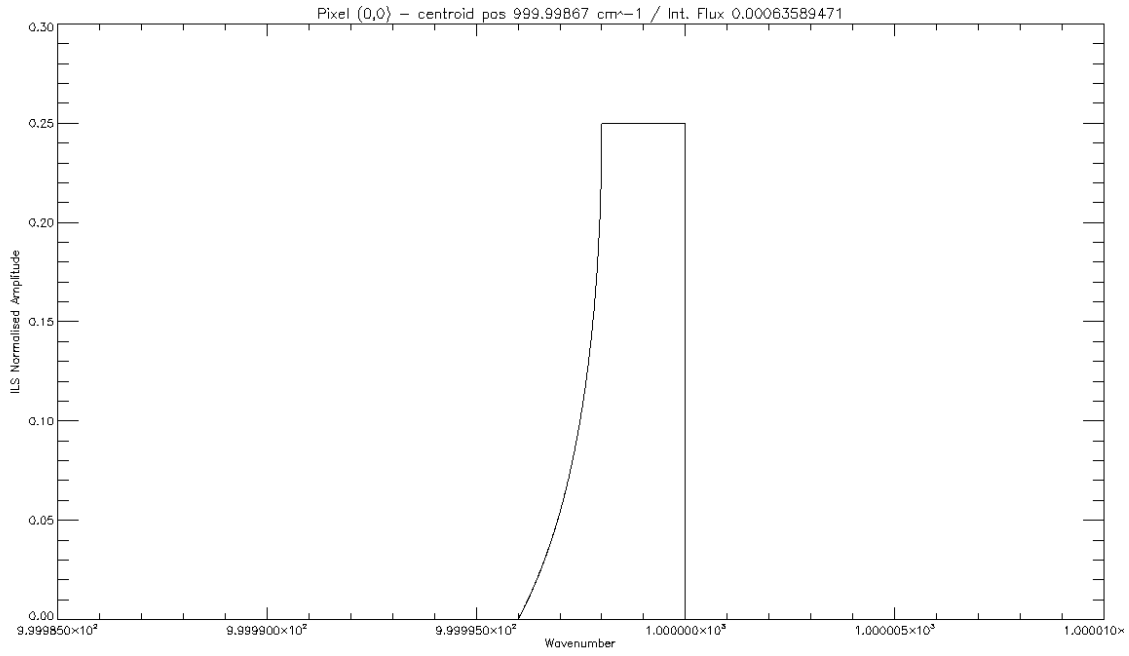

Figure 3.11: Wavenumber against normalised ILS for the 160x120 array, centre pixel

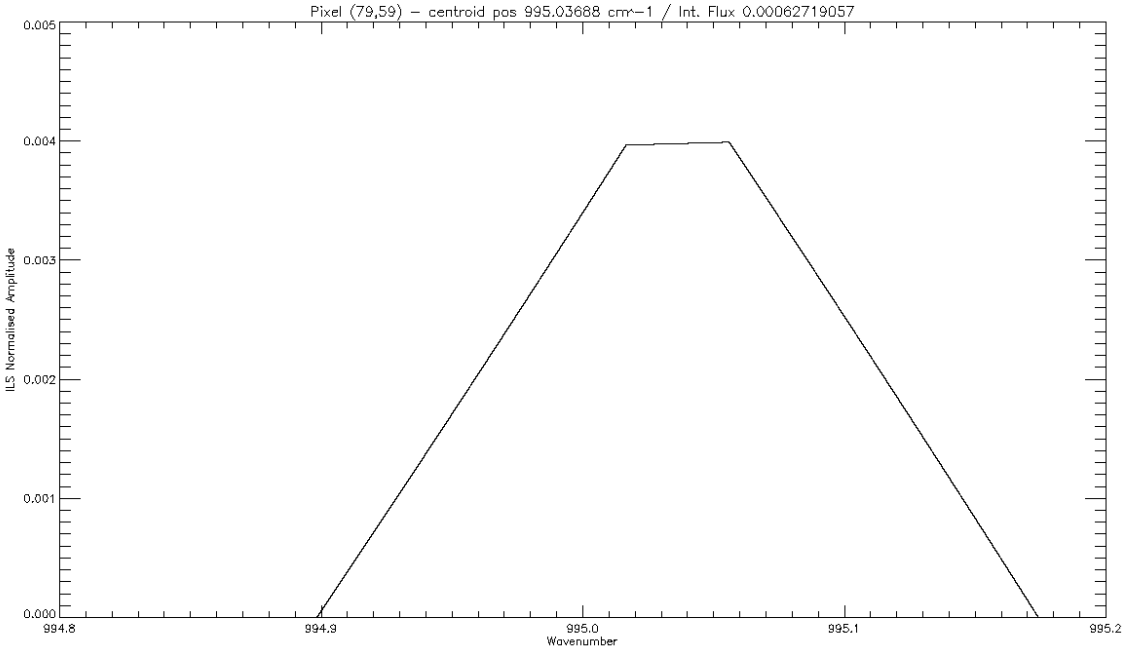

Figure 3.12: Wavenumber against normalised ILS for the 160x120 array, edge pixel.

We note that the camera field of view is larger then the output port of the interferometer and so an initial task was to identify which pixels are imaged. This task enabled us to produce a pixel mask which identifies a pixel as imaged (one) or not imaged (zero). This mask is used in the interpretational software to discard non-interferometric data. Figure 3.13 shows the pixel mask that is used. Figure 3.14 shows an image taken with the camera array, highlighting any pixels that had failed and these were also set in the pixel mask. The presence of banding on this image is a processing artefact.

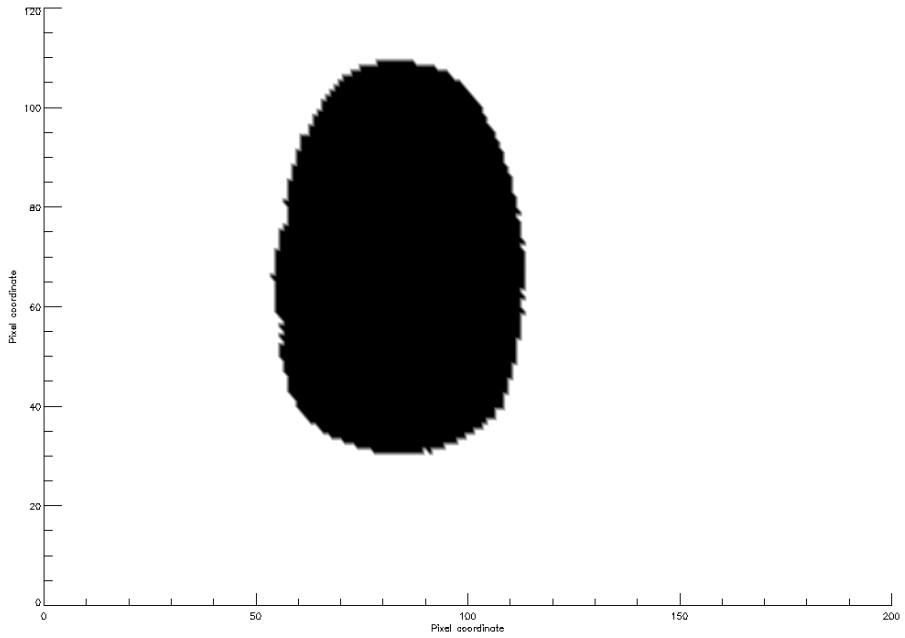

Figure 3.13: Pixels imaged through the inteferometer.

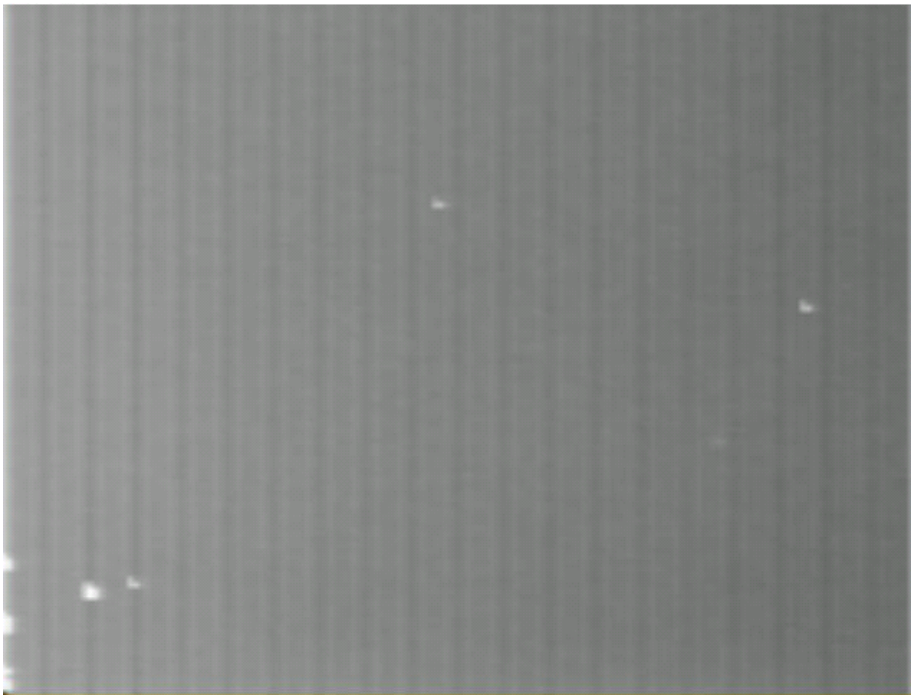

Figure 3.14: Dead pixels on the camera array.

If we only consider those pixels that are actually imaged through the instru-

ment, then the maximum line shift is only approximately $1.25cm^{-1}$ as the field of view is reduced. We also wish to avoid part of the travel range of the mirror due to larger positional errors, and allow ourselves the ability to take a two-sided interferogram for phase correction. As a result, LoCI has its zero path-difference position set to approximately $150\mu m$ into its travel range. Figure 3.5 shows the travel range for each interferogram. This gives us $150\mu m$ of mirror travel for a double-sided interferogram, which equates to $300\mu m$ of optical path difference from zero PD to maximum travel. For a single-sided interferogram, this has a mirror travel of $200\mu m$ and a maximum optical path difference $\Delta z = 400\mu m$. This means that our apodised spectral resolution now becomes:

$$\delta\sigma_{corr} = \frac{1}{\Delta z} = \frac{1}{0.004cm} = 25cm^{-1}. \tag{3.7}$$

Therefore the line shift due to obliquity is now approximately 5% of the spectral resolution of the instrument in the worst case. Figure 3.15 shows the line centroid shift for each imaged pixel and we see as expected that this extreme line shift is only seen at the furthest pixels from the optical axis. To calculate the line centre, we used the formula for the area centroid,

$$\sigma_{centre} = \frac{\int \sigma I_F(\sigma)}{\int I_F(\sigma)}, \tag{3.8}$$

where $I_F(\sigma)$ is the ILS normalised flux. Figure 3.16 shows the maximum ILS flux at this line shifted position and we see that the maximum normalised flux of the ILS is around 25% at the central pixel position, where the line shift is minimal. The normalised flux rapidly drops to approximately 1% or less at the edges where the largest line position shift is seen. We show in Figure 3.17 the width of the ILS

over the camera array, and the maximum broadening is around $0.15cm^{-1}$.

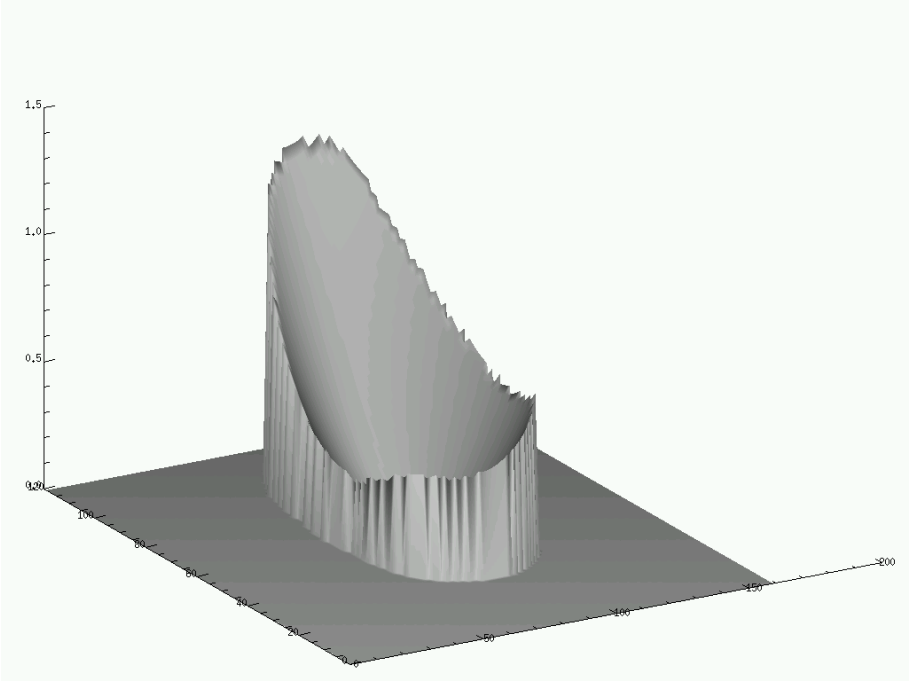

Figure 3.15: Line shift for a 1000 cm^{-1} line as imaged through LoCI.

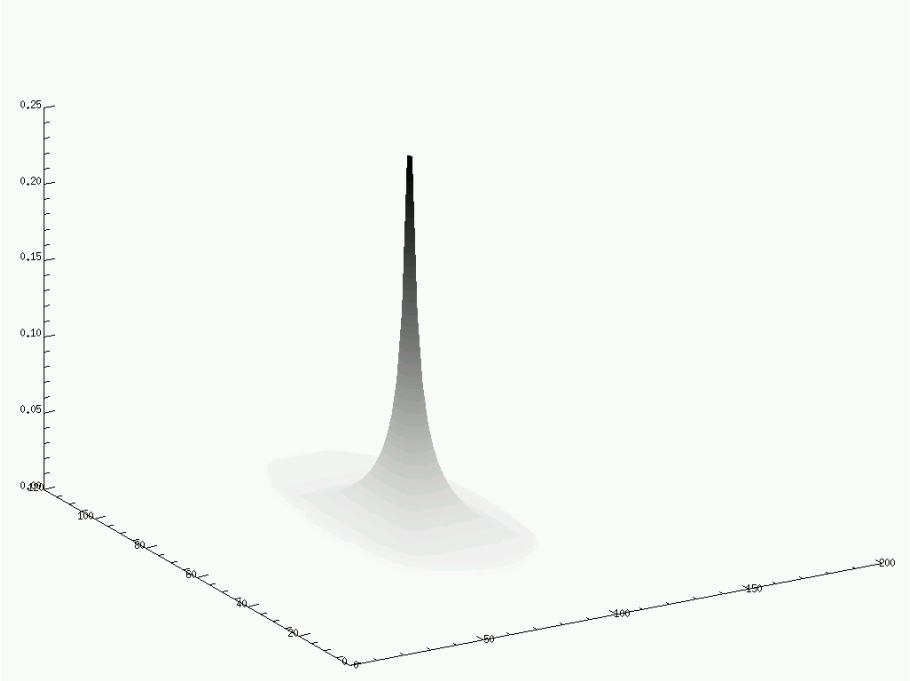

Figure 3.16: Normalised ILS flux for a 1000 cm^{-1} line as imaged through LoCI.

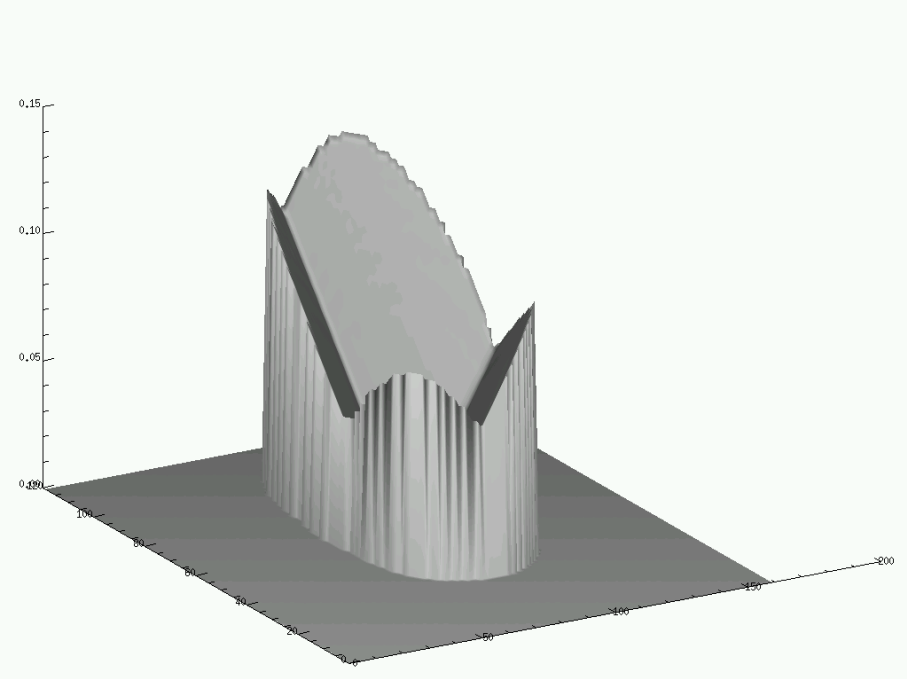

Figure 3.17: Normalised ILS maximum width for a 1000 cm^{-1} line as imaged through LoCI.

We then generated a normalised blackbody curve for $T = 500K$, whose range and limits were identical to a set of test observations taken with LoCI. We then generated a single Gaussian profile absorption line centred on $1000cm^{-1}$ of various widths and amplitudes and added this absorption line to the blackbody SED to simulate an absorption line in a SED. Water vapour in the atmosphere has many absorption lines in this infra-red region and so we aim to simulate the effects of obliquity on the observation. We do this by convolving the ILS for each pixel with the simulated SED, using the normalised feature to create a normalised SED observation. We then compared the original SED with the simulated observation for the central pixel and the furthest imaged pixel from the optical axis. We used a χ^2 metric to compare the two, and tabulated the metric for a wide range of amplitudes and widths. Figure 3.18 shows a sample SED observation from the

simulations.

These observations are binned using a sample observation to generate the wavenumber binning resolution. We compare the convolved- and non-convolved observations because we are in effect comparing two 'black-box' interferometers – one which has an obliquity correction factor applied and one that does not. Comparing the two results allows us to determine whether we need to add an obliquity correction factor into the interpretational software that we provide.

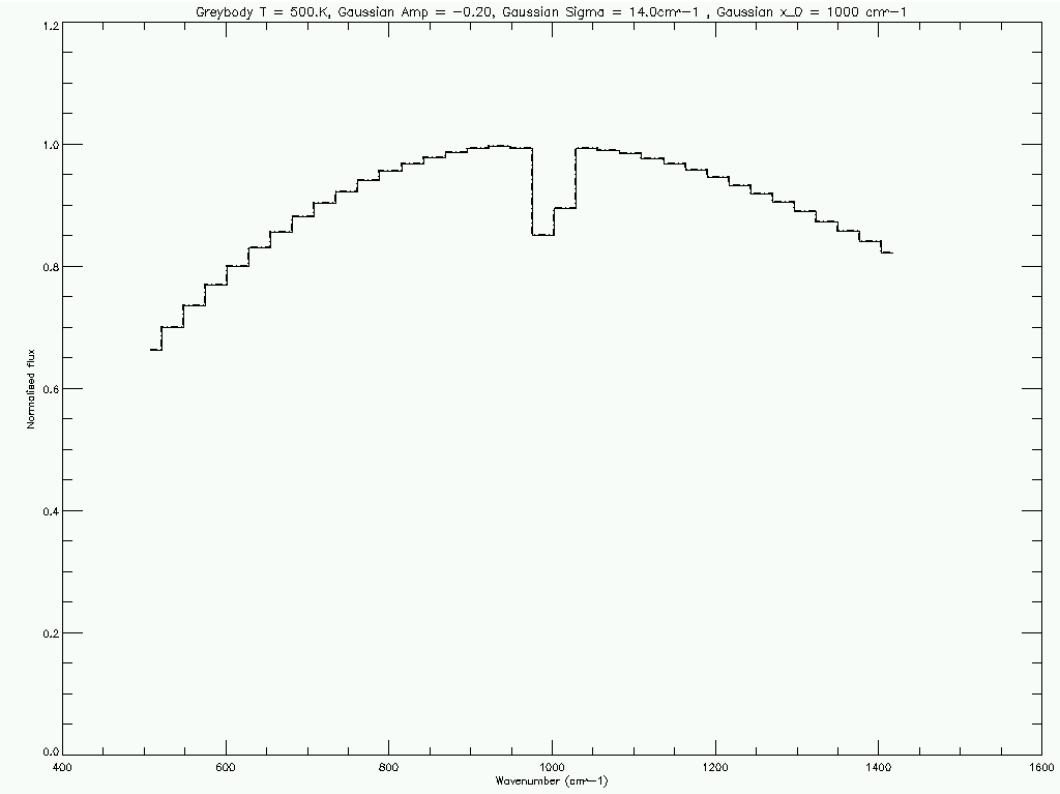

Figure 3.18: A sample observation, with a single absoprtion line.

These sampled observations had a range of normalised amplitudes from $0.05-1.00$ and a range of σ for the Gaussian profile of $1-20cm^{-1}$. Figures 3.19 to 3.24 show the resultant χ^2 metric for our simulations for the central pixel, and the furthest imaged pixel from the optical axis. The simulations were performed in

specific ranges of amplitude and σ with ten steps in each. For amplitude, we defined *narrow* as between $0.05 - 0.5$ and *wide* as between $0.10 - 1.00$. Similarly for σ, *narrow* is between $1 - 10cm^{-1}$ and *wide* is between $2 - 20cm^{-1}$. Here the *narrow* data set has a narrow step and range and the *wide* has a wider step and range of values. This enabled us to investigate a combination of depths and widths to the absorption profile. We found that, as expected, the metric increases towards the edge of the array. The metric shows that for increasing width and depth of absorption lines, we see a worsening of the fit. It is interesting to note that the maximum value is not for the widest, deepest lines but for deep lines of medium width – of order of the binning size. However, the tendency is for wider, deeper absorption lines to show a worse quality recovery than narrow, shallow lines. Over all simulations, the maximum value for the metric is of the magnitude of 10^{-6} and so even for broad deep absorption lines the obliquity of the pixel has a very weak effect on the recovery quality of the SED. Repeating these simulations with a slight offset to the line centroid, to see if the effect would become more noticeable if the centroid was offset within the bin, yielded similar results.

We should also note that the instrument returns the reading for each pixel as a 16-bit integer, and so the camera output has an inbuilt resolution of approximate 10^{-4} as a 16 bit integer varies from 0 - 65535. We can refer to this as 'digitisation noise' and is a fundamental lower limit of the noise of the device. It is therefore quite doubtful that any significant effect from obliquity on the output from LoCI will be easily detectable as it stands. Should the detector array be replaced in future, these calculations would need to be revisited.

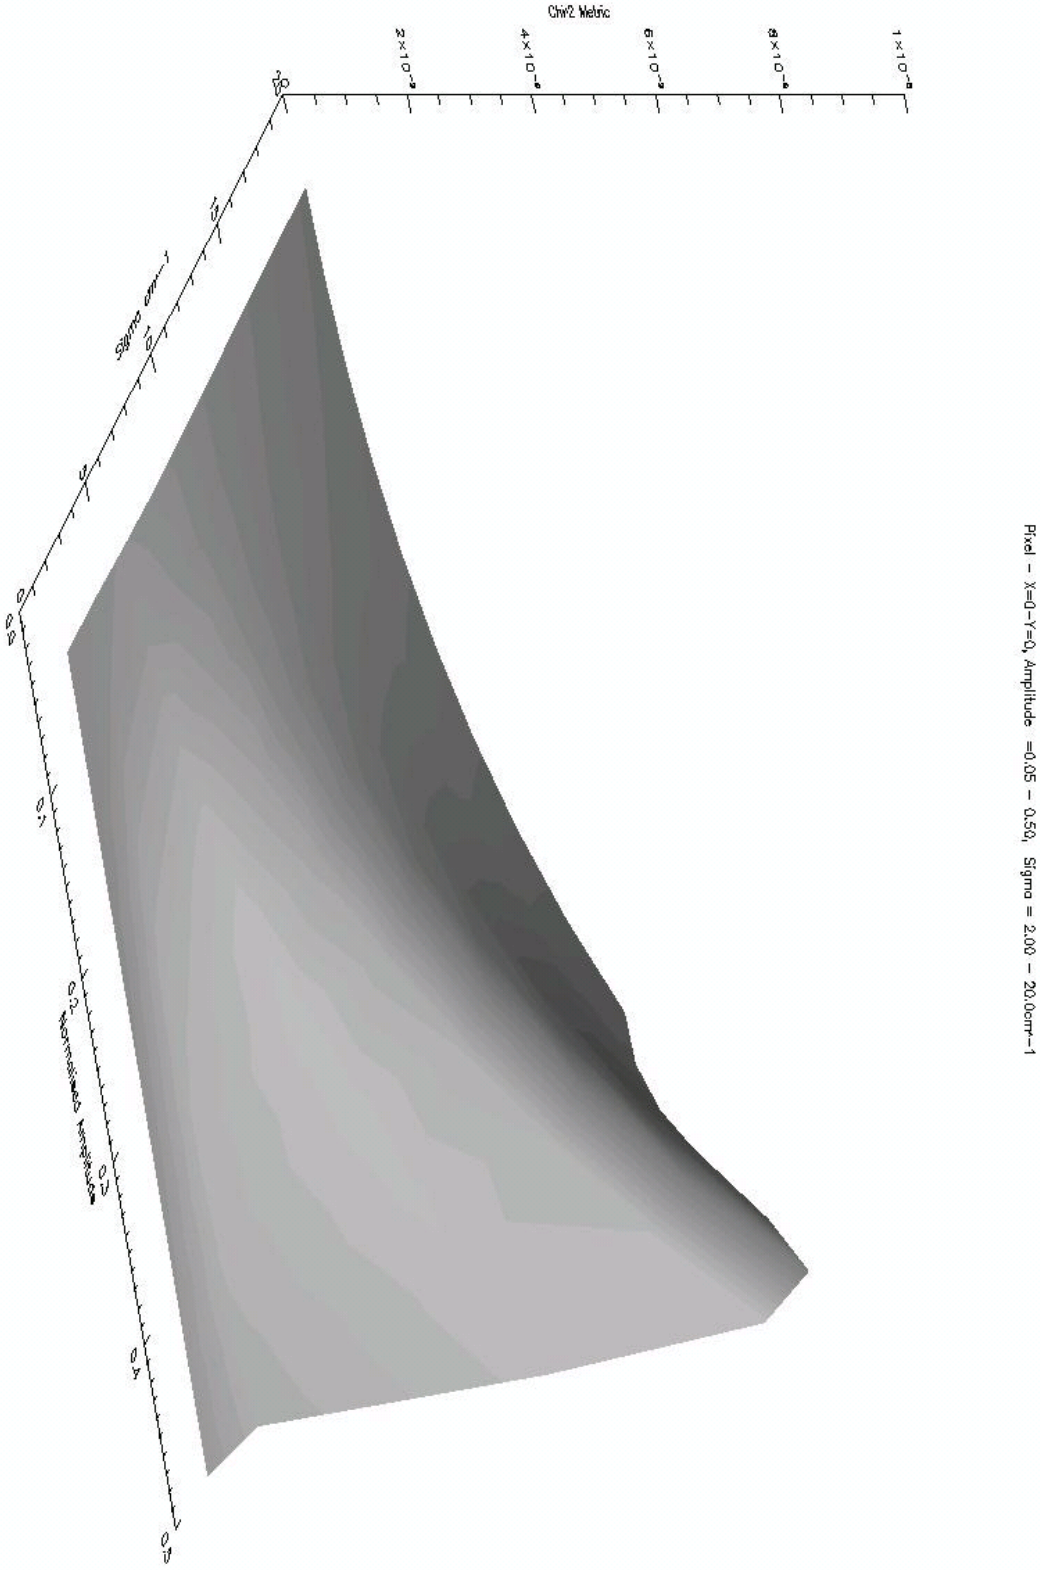

Figure 3.19: χ^2 metric for the *narrow/wide* simulations, through the central pixel.

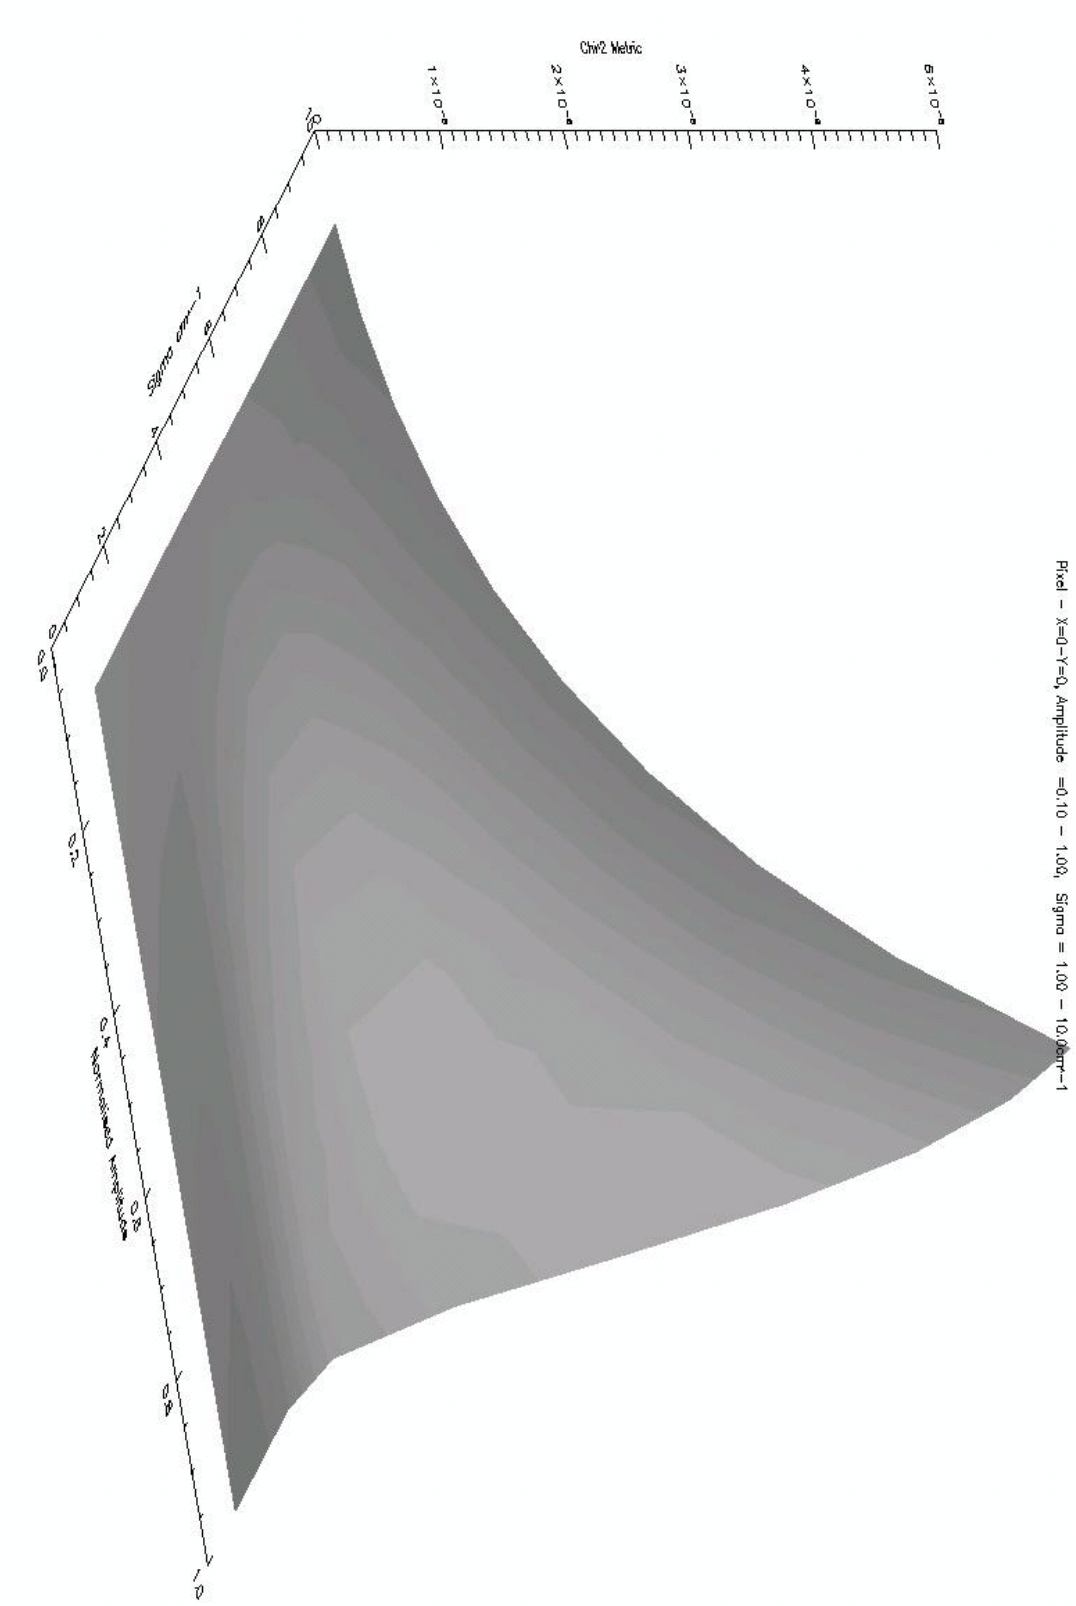

Figure 3.20: χ^2 metric for the *wide/narrow* simulations, through the central pixel.

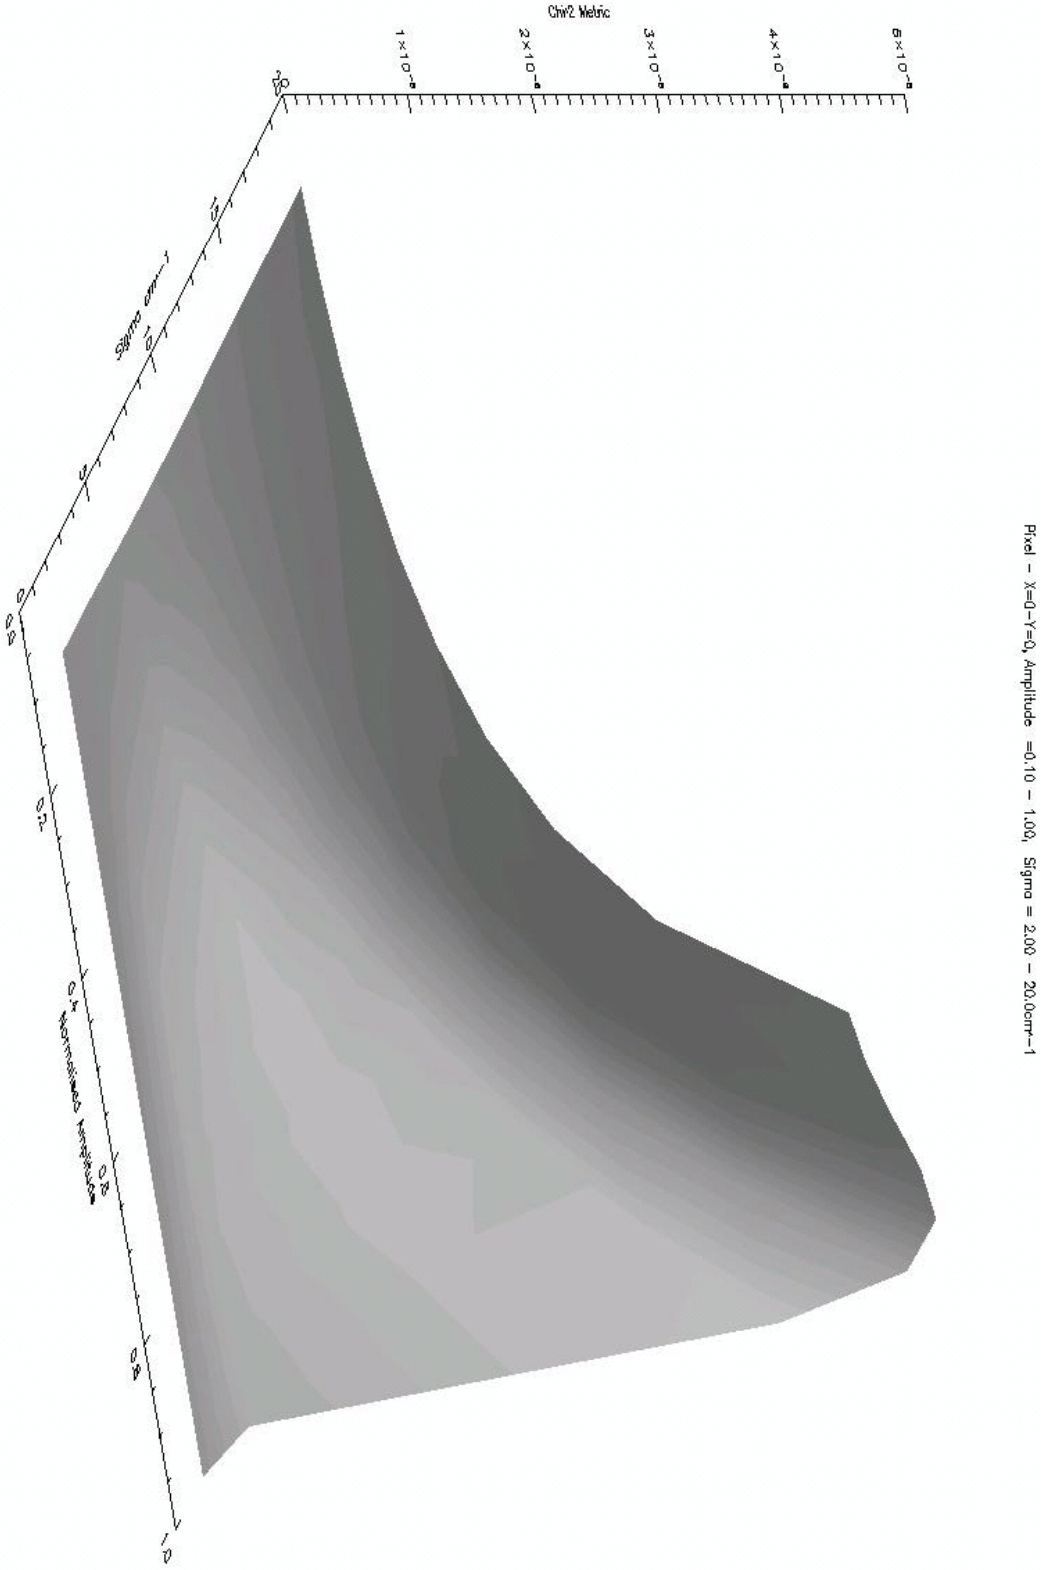

Figure 3.21: χ^2 metric for the *wide/wide* simulations, through the central pixel.

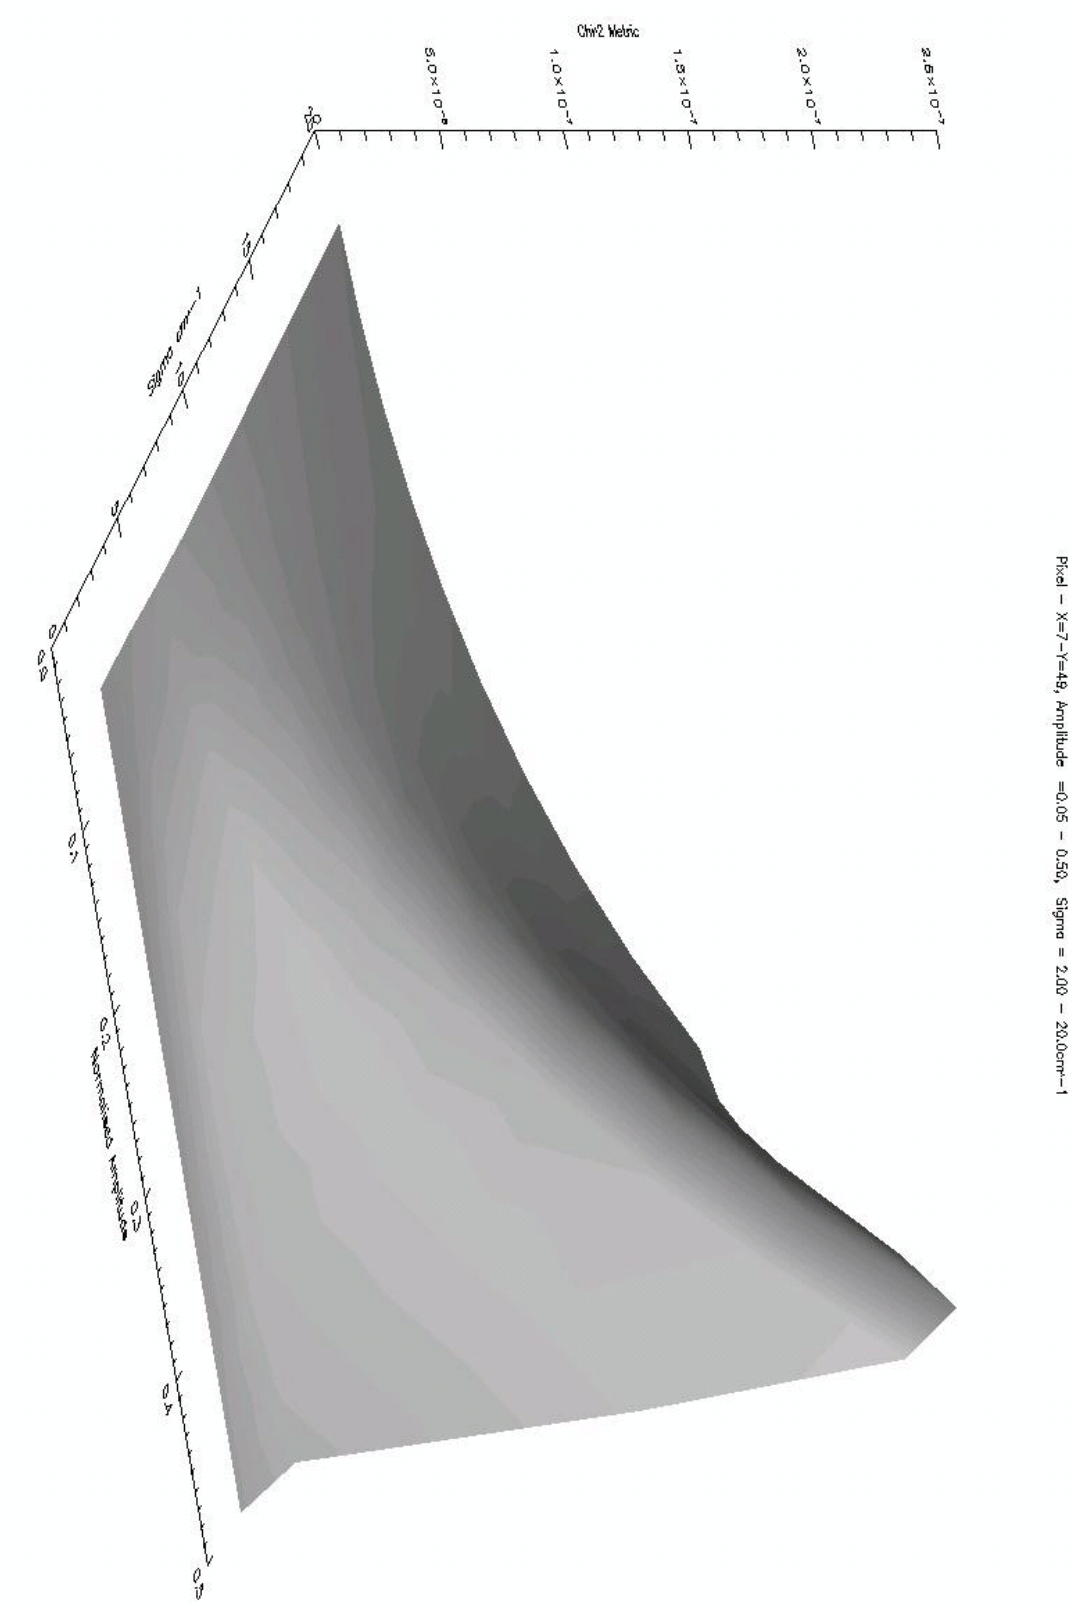

Figure 3.22: χ^2 metric for the *narrow/wide* simulations, through the edge pixel.

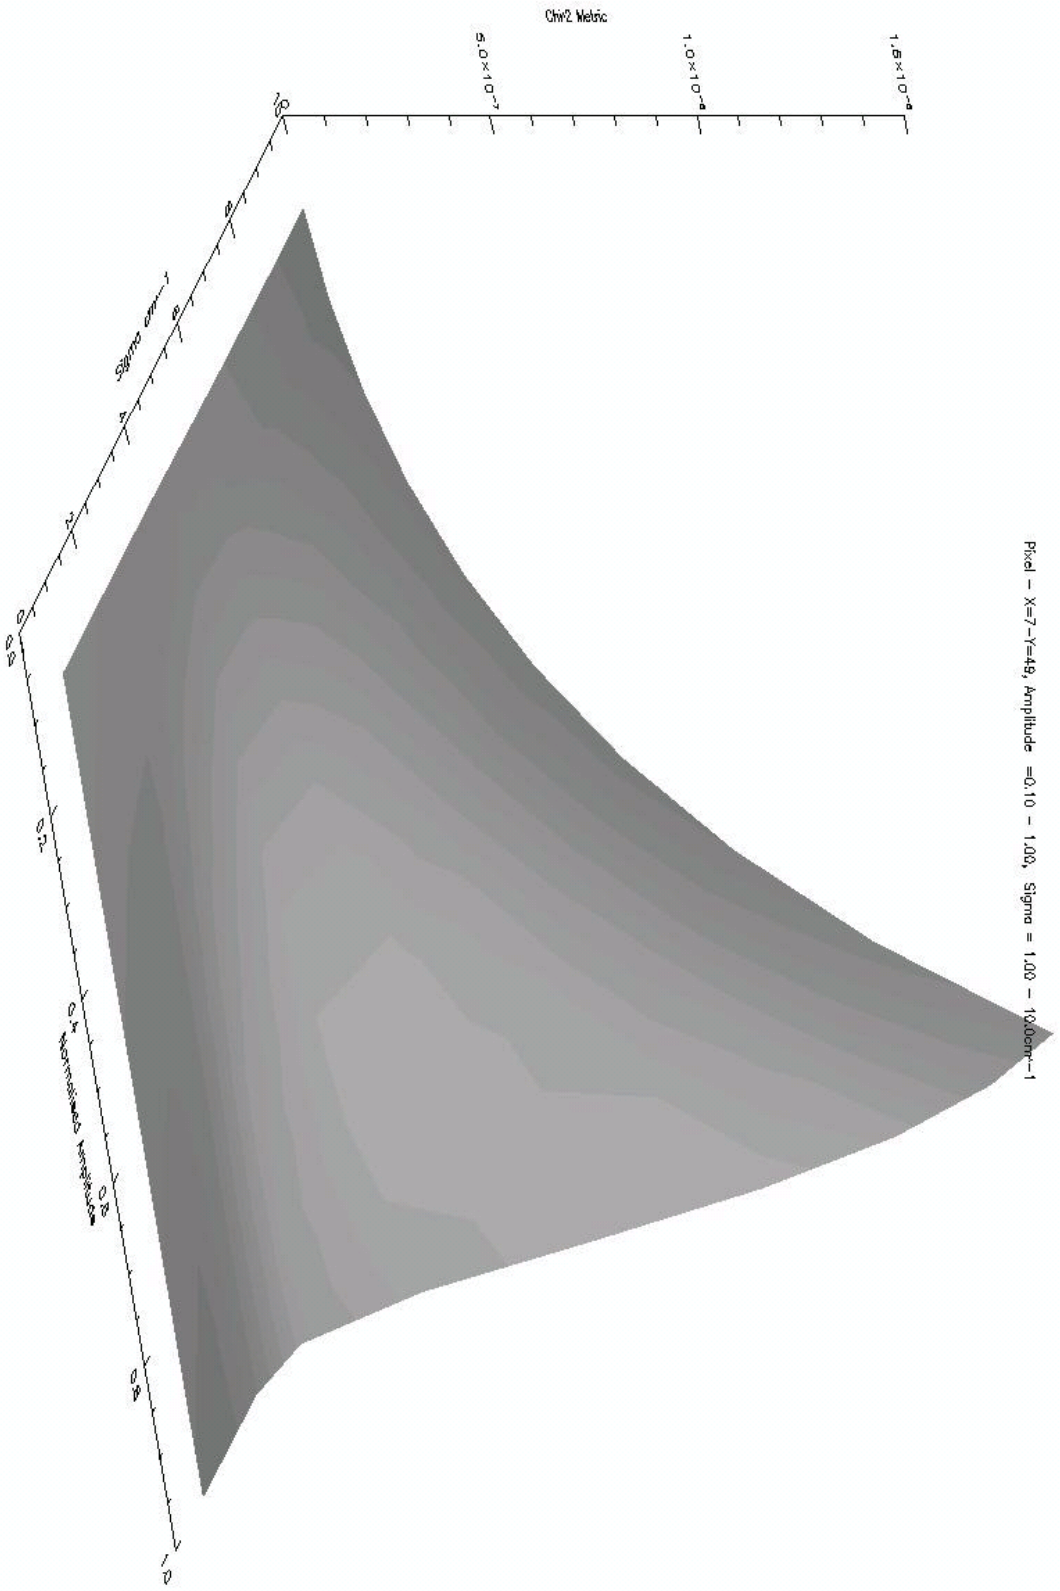

Figure 3.23: χ^2 metric for the *wide/narrow* simulations, through the edge pixel.

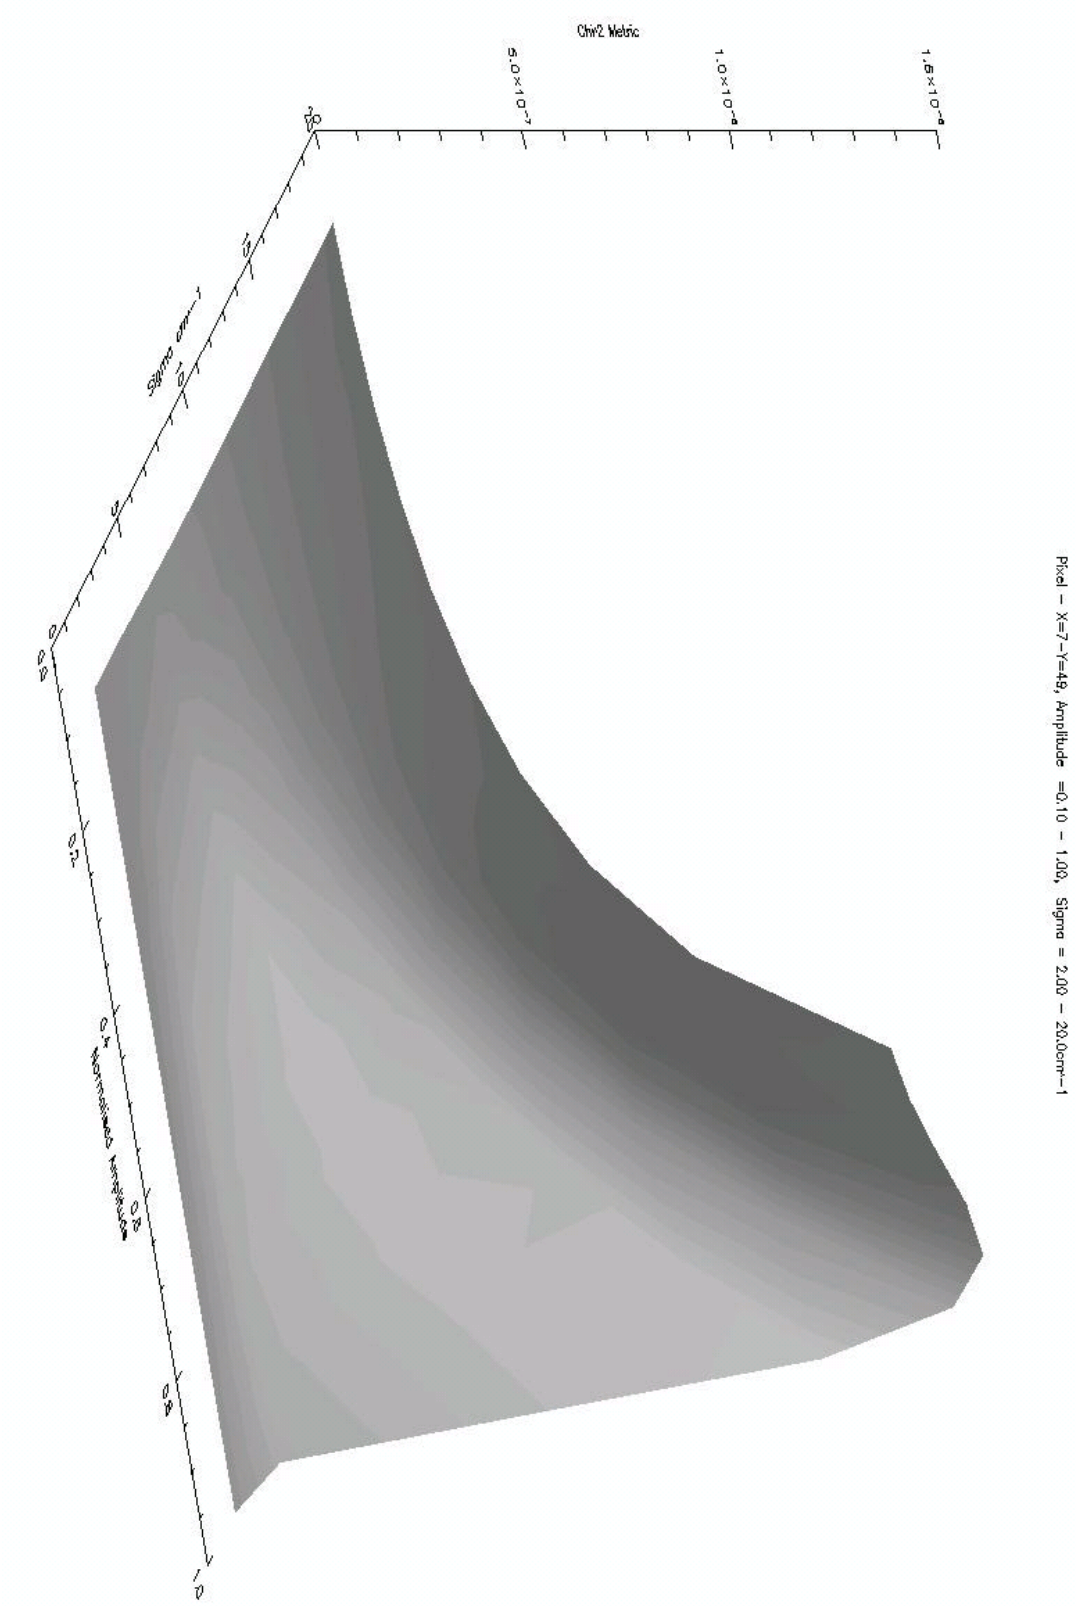

Figure 3.24: χ^2 metric for the *wide/wide* simulations, through the edge pixel.

3.2.4 Phase Correction Methods

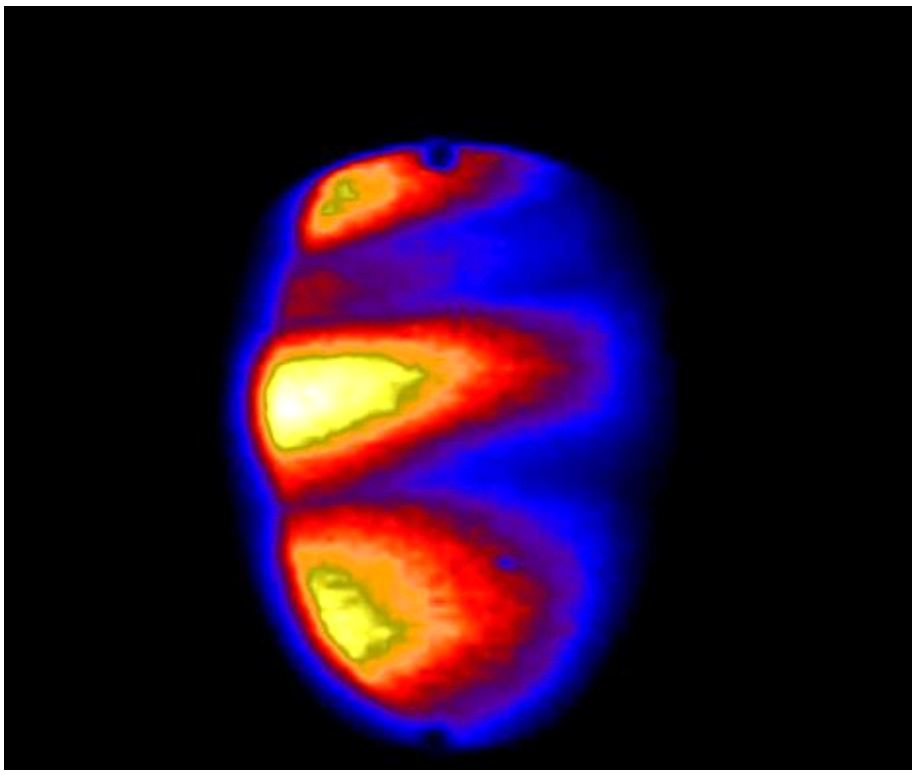

Figure 3.25: A false colour image of a hot body, imaged through LoCI - intensity.

The data reduction software provides for both Mertz and Forman methods of correcting phase errors, as outlined in Section 2.2.3. There, we noted that errors of only a few percent are likely with the Mertz method should double-sided interferograms be available. Using the data taken from a hot body, whose false colour intensity map is shown in Figure 3.25, we compared the results. The data was processed using both Mertz and Forman methods, and the central 400 pixels were compared. This was done by normalising both data sets to the Forman-corrected data. We compared data over a slightly extended wavenumber range (200 to 2000 cm^{-1}) to compare noise outside the bandwidth of the camera (700 - 1400 cm^{-1}). This enabled us to see if any effect from zeroing the phase angle in the non-usable bands in the Mertz method was visible.

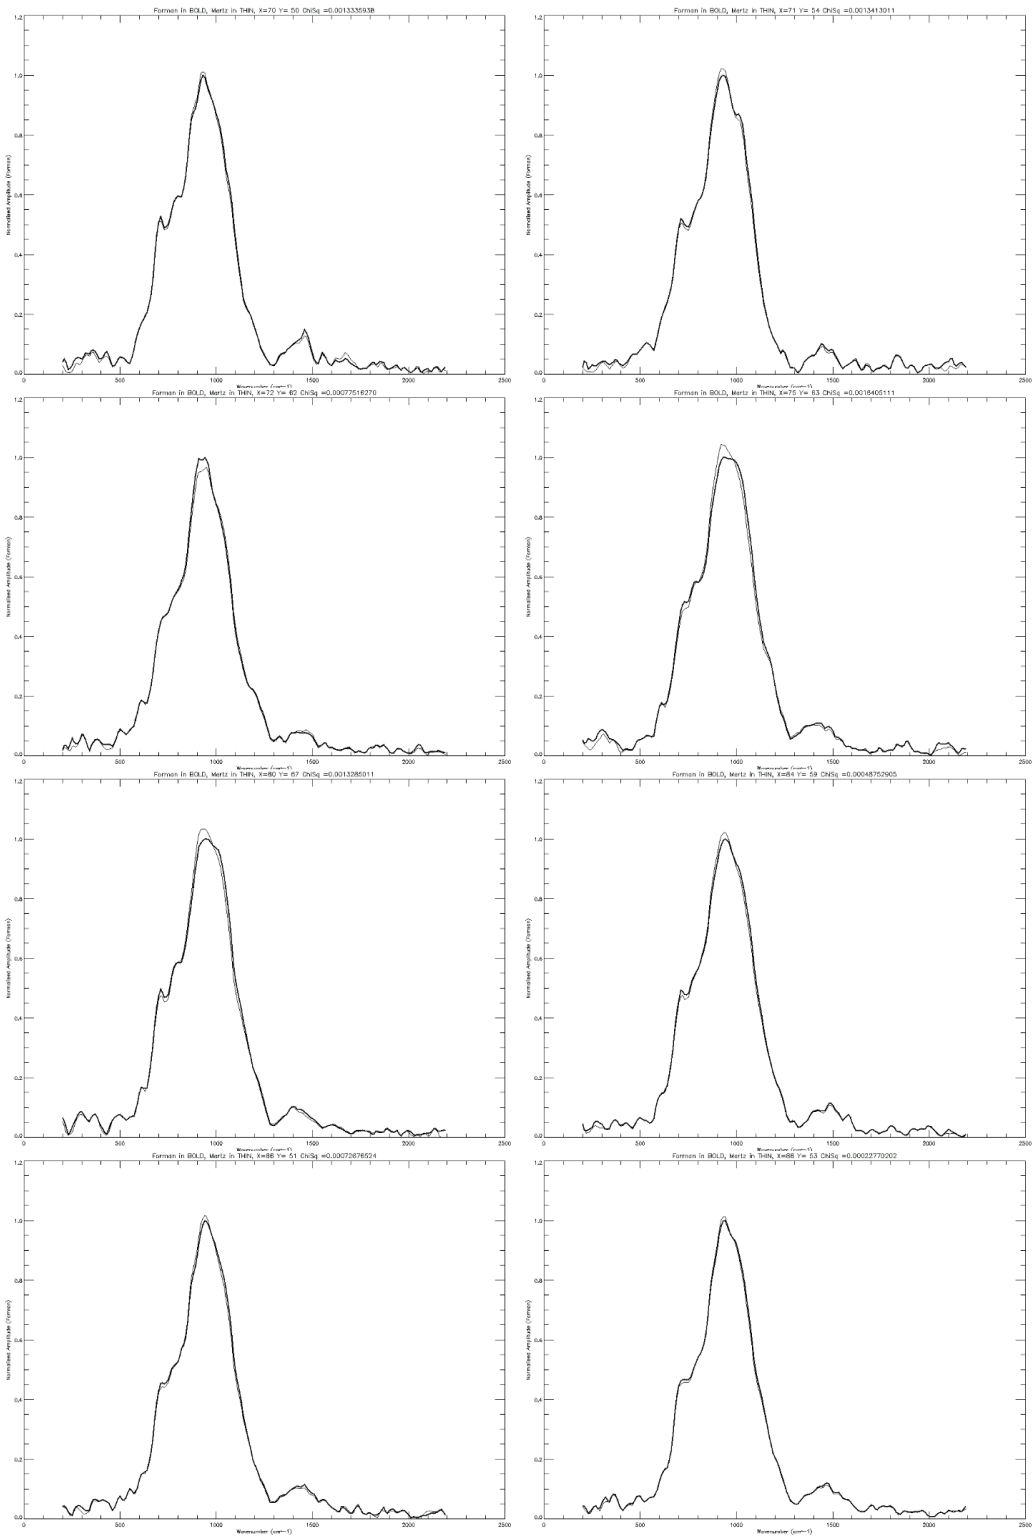

Figure 3.26: A selection of pixels comparing Mertz and Forman methods, plotting normalised amplitude versus wavenumber.

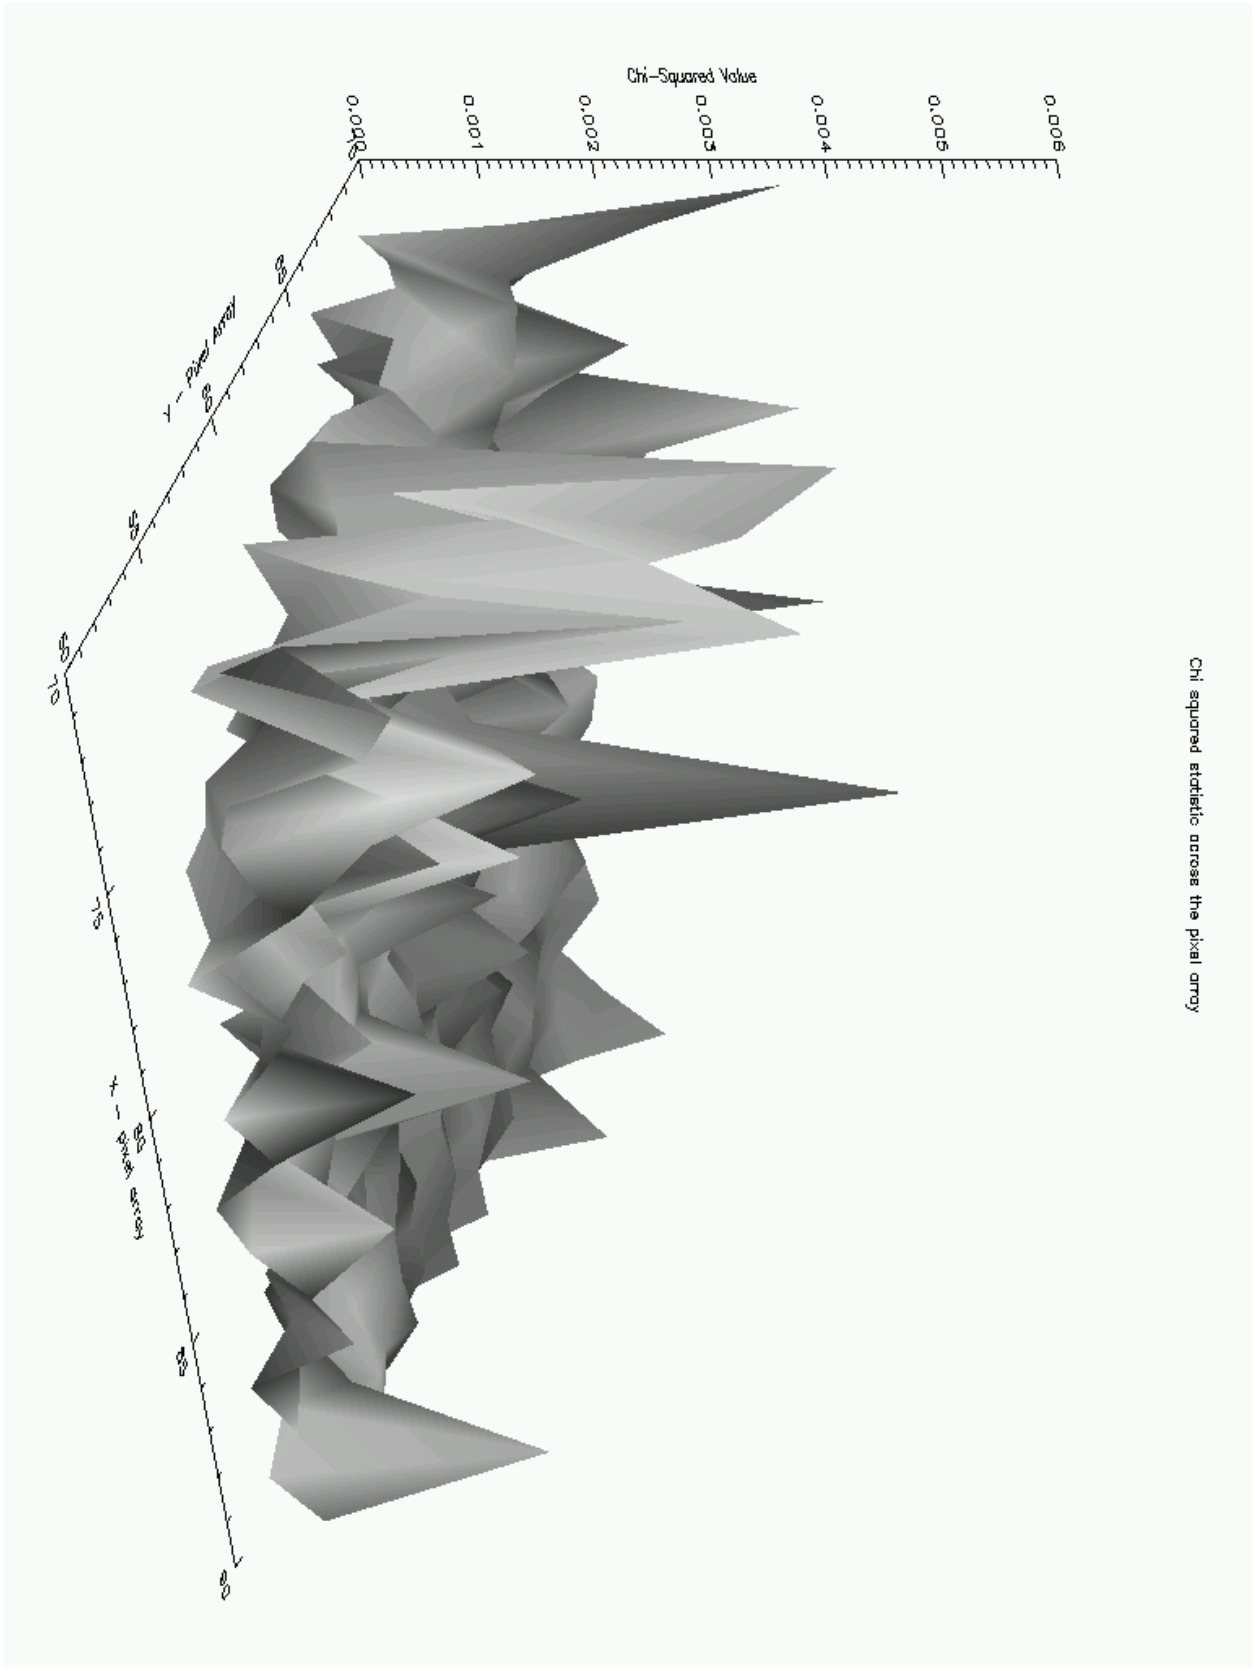

Figure 3.27: Comparing the Mertz and Forman methods - Chi-squared plot.

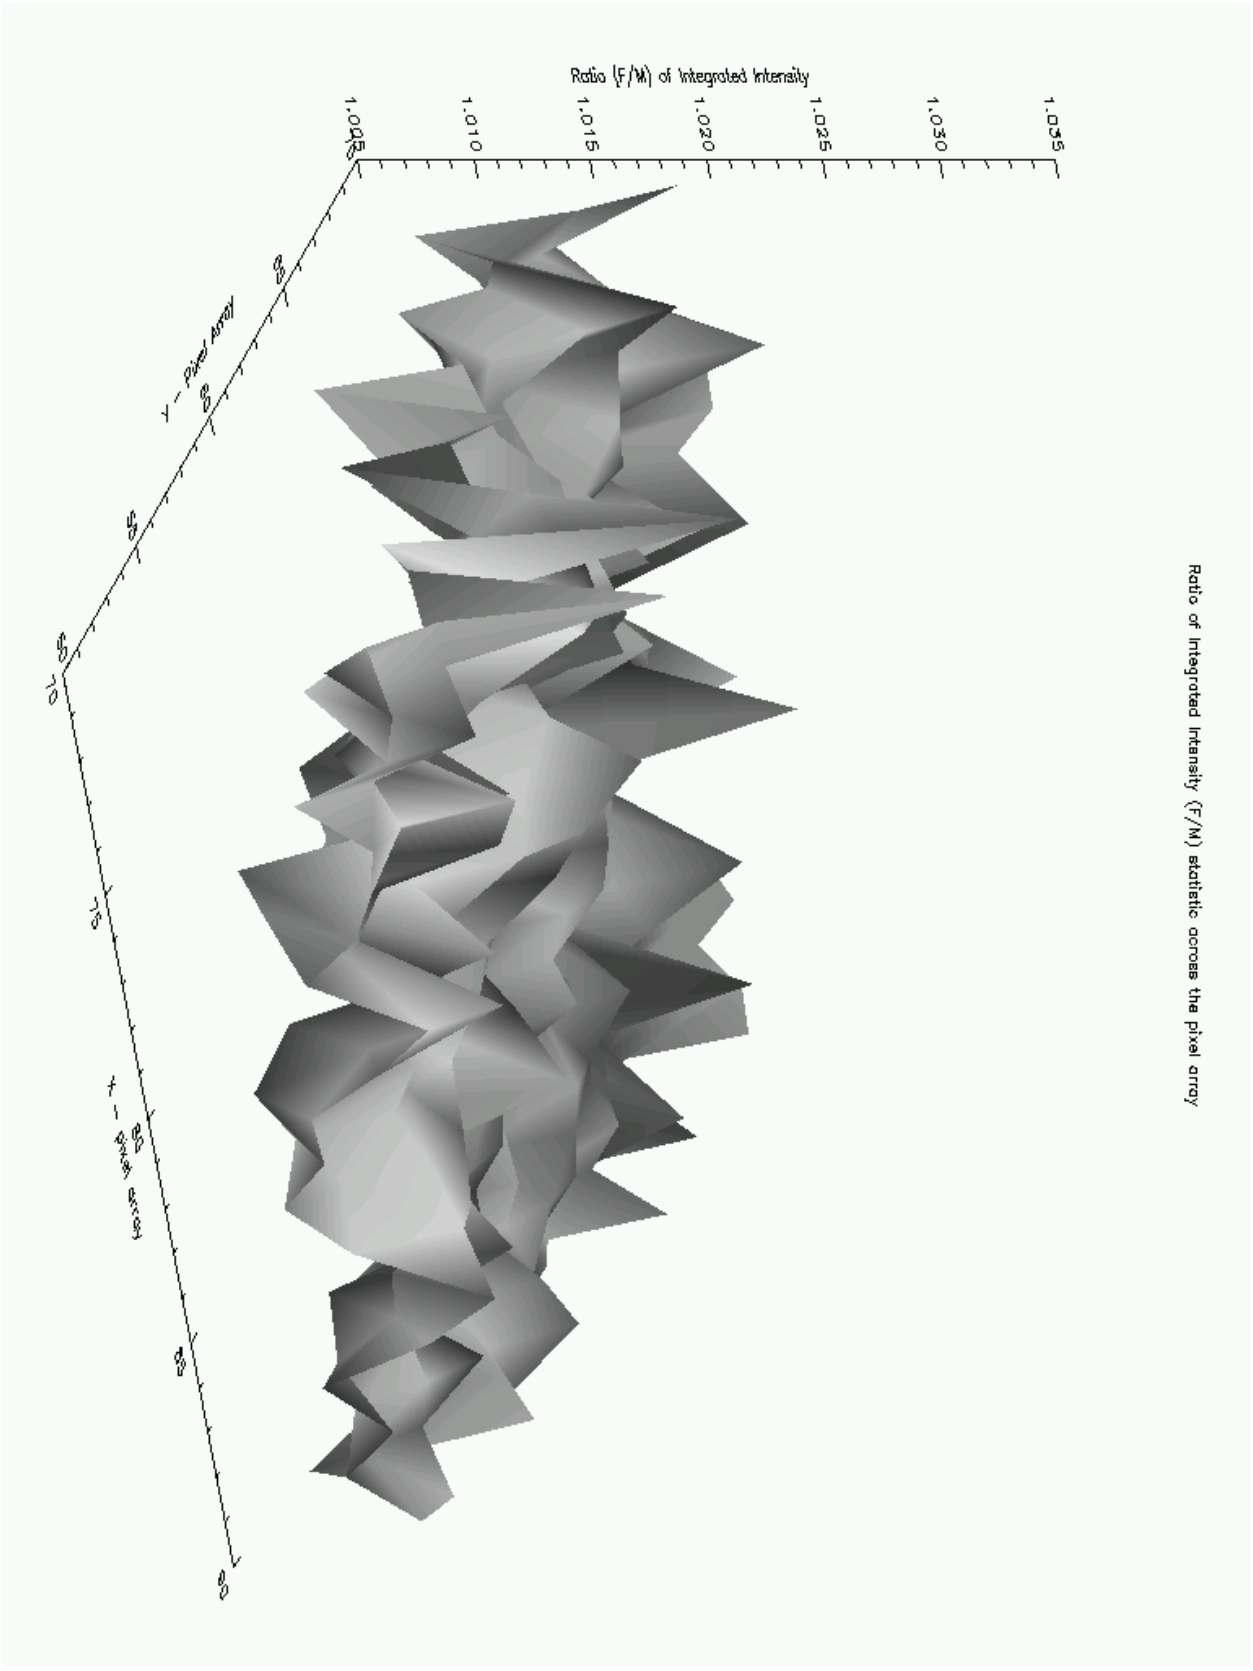

Figure 3.28: Comparing the Mertz and Forman methods - Integrated Intensity ratio.

Figure 3.26 shows a random eight pixels selected from the 400 processed. We note that overall, both methods agree very well indeed, and Figure 3.27 shows the χ^2 calculation for each pixel. We note there is no real dependence on position, save for a tendency for the χ^2 value to increase with y-position towards the end of the metric. We must point out that Figure 3.25 shows that the image is weaker across the y-range in this area, and what we are seeing is that the methods are diverging slightly for weaker spectra. However, the value for χ^2 is very small indeed across the range and both methods agree well. Future work to calibrate the instrument could shed more light on which method is more accurate, but as seen in Section 2.2.3 the difference between the two methods is, as predicted, within a few percent. Outside the bandwidth of the camera, both methods agree very well and so it seems that zeroing out the phase angle outside of the waveband of the camera has reduced noise, and as a result, both methods agree very well indeed. Figure 3.28, which shows the ratio of integrated intensity across the array, also shows excellent agreement between the two. This figure, which shows the ratio of I_F/I_M, shows that the difference in integrated intensity between the two is such that the Forman method yields a larger integrated intensity in general. However the deviation is between $1-2\%$ and again agrees with the conclusions of section 2.2.3.

3.3 LoCI Data Reduction

Data capture from the instrument takes place via a bespoke piece of software donated by the University of Lethbridge, Canada that was used for a similar instrument there, (Naylor, 2009). This software is controlled by a text configuration file that lists scan length, scan step and data output directory. The output from

this program is a standard FITS file. This file is then processed by custom IDL code that processes the interferogram and produces an output data-cube. This is a three-dimensional data structure where the x and y dimensions are the pixel coordinates and the z dimension is wavenumber. This is similar to the input data-cube, but for the input cube the z dimension is the path difference z. The camera field of view is larger than the beamsplitter size and so initial readings were taken to create a pixel mask (as mentioned in the previous section). This is a file that lists which pixels contain data that has passed through the interferometer. It also contains a list of those pixels whose response is faulty. Using this pixel mask helps reduce the processing load on the software.

A flat-field interferogram was created to help for varying pixel sensitivity across the field of view. This flat-field interferogram is used with the pixel mask to process each pixel prior to the Fourier transform. Using the scan size and range information, a routine attempts to automatically find the zero PD of the interferogram. The interferogram is interpolated to a fixed path difference grid and baseline subtracted. Triangular apodisation is performed upon the interferogram and then phase correction is applied. A choice of Mertz or Forman technique (See Section 2.2.3 for further explanation) is used. This relies on a small two-sided interferogram, and so the detection of the zero PD position is used to create this phase-correction interferogram. By automating the detection, we can take interferograms using the widest possible range of the piezo controlled movable mirror.

As the alignment routine would have given us an initial position for the zero PD, we have some freedom in choosing the range of travel. If we set the piezo-controlled mirror at approximately $150\mu m$ from the end position, and then align the instrument to this zero PD position we guarantee the position falls within the acceptable error range and sets the operational resolution as explained in the

earlier section. The software can then automatically create a smaller interferogram of the correct size to create the phase correction term by automating the zero PD position discovery. The interferogram can also be spliced at the zero PD position and the larger of the two halves used to create the spectrum by mirroring about the zero PD position to create a symmetric interferogram. A symmetric interferogram is preferable to minimise phase errors, as shown earlier in Section 2.2.3. As we cannot guarantee we sample at exactly the zero PD position, creating a symmetric interferogram in this way removes the asymmetry caused by missing the position. The resolution of the resultant spectrum is depenedent on the size of the larger, single-sided interferogram. This spliced, mirrored interferogram after apodisation is used to create the phase correction term to apply to the single-sided, larger interferogram. The phase-corrected interferogram is then Fourier Transformed to produce a phase-corrected SED for each pixel.

An IDL format .SAV file is then saved for future investigation. This contains the output data cube (x, y, z as earlier) with variables for the wavenumber range, ZPD across the pixel range and any bad pixels found.

3.3.1 Sample LoCi Output

Figures 3.29 and 3.30 show images of a soldering iron in a spring holder, as imaged by the instrument in intensity and wavenumber. A pseudo-colour image was generated by binning each pixel into six bins and assigning each bin to a RGB value. Figure 3.29 is a false colour image showing the peak intensity at each position, and Figure 3.30 is binned to show the peak wavenumber at each position. This was taken as one double -sided scan. One of the pixel observations is shown in Figure 3.34, highlighting the large range of wavelengths taken, even though only a narrow

range is actually usable. Although there is a degree of light loss within the interferometer, and a perfect inteferometer loses 50% of the light through it, this can be overcome by multiple scans and stacking the result. Figures 3.31 and 3.32 are images of a heated metal plate, again as false colour images of peak wavenumber and intensity. We see in this case some uneven heating within the plate. Figure 3.33 shows a 3D view of the LoCI observation of the heated plate.

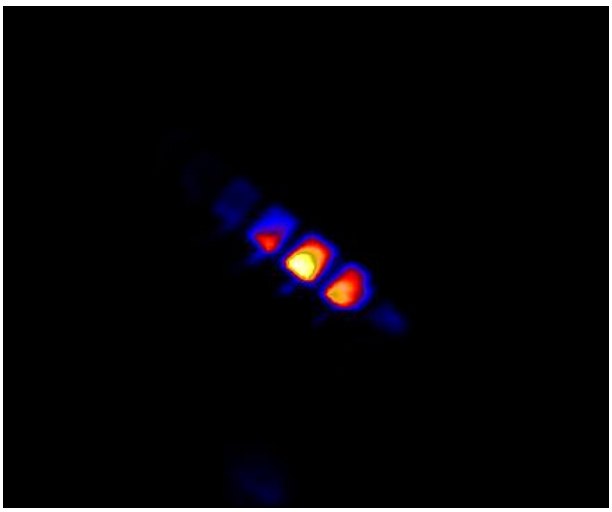

Figure 3.29: A false colour image of a soldering iron, imaged through LoCI - intensity.

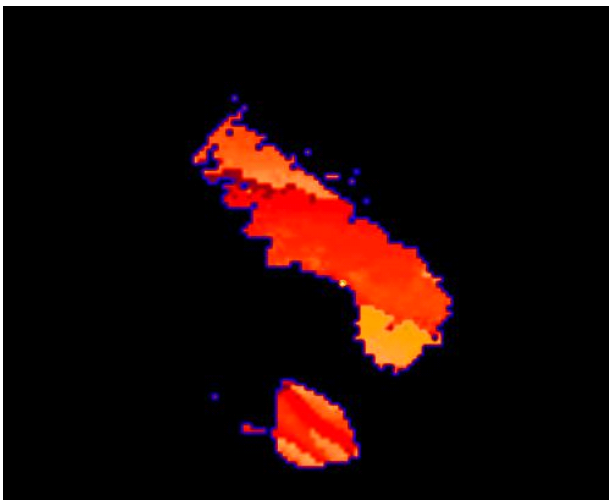

Figure 3.30: A false colour image of a soldering iron, imaged through LoCI - wavenumber distribution.

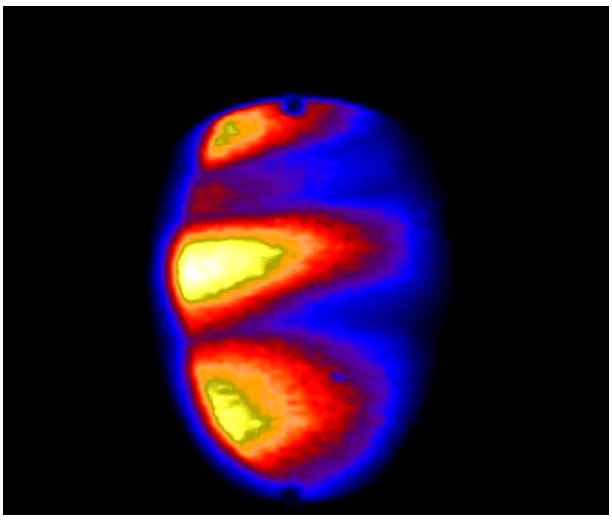

Figure 3.31: A false colour image of a hot body, imaged through LoCI - intensity.

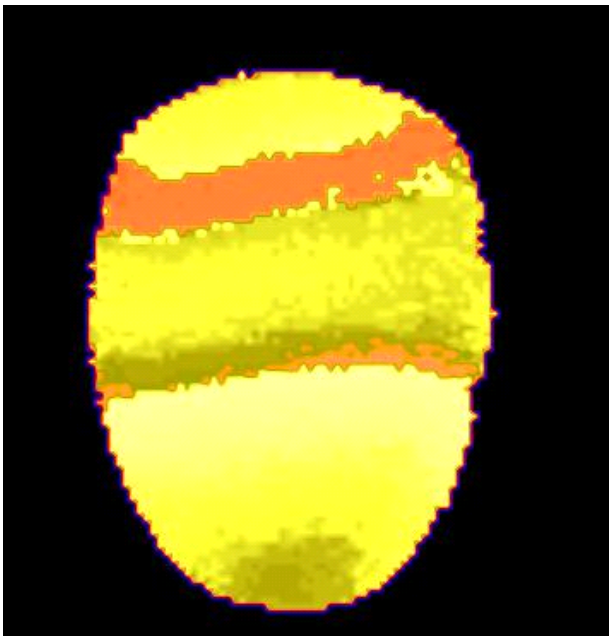

Figure 3.32: A false colour image of a hot body, imaged through LoCI - wavenumber distribution.

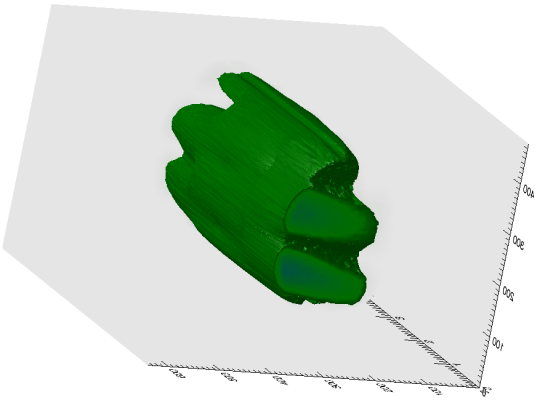

Figure 3.33: A false colour 3D image of a hot body, imaged through LoCI.

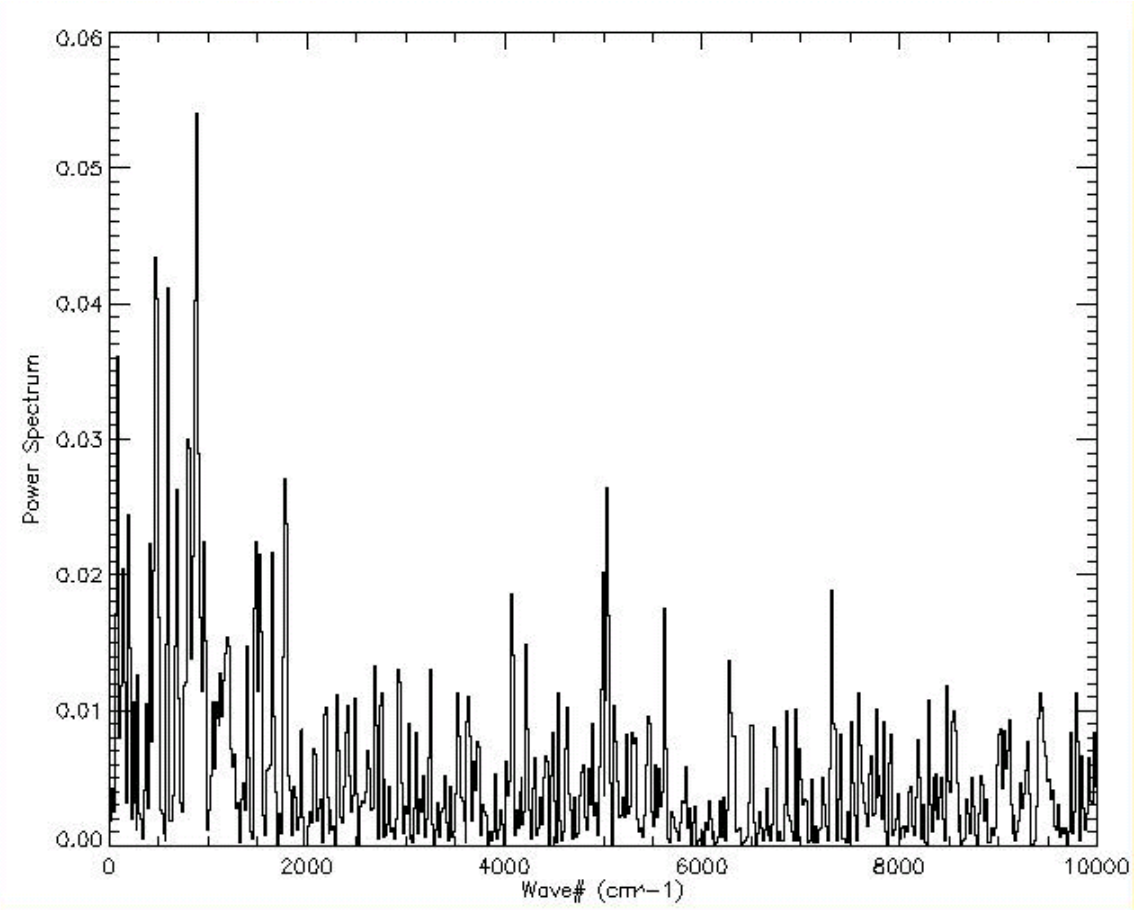

Figure 3.34: A single pixel observation through LoCI.

3.4 LoCI summary

Through the use of off-the-shelf components, we have shown it is possible to construct a portable, non-cryogenic infra-red interferometer. Using a multiple pixel detector has enabled us to examine some of the developmental issues that face the design and construction of larger instruments such as SPIRE. Although some construction issues are still ongoing, the instrument has shown that it is possible to construct a low-cost interferometer that can be used to produce useful science. The design and construction of LoCI highlights many issues that are faced in astronomical detectors, such as housing design for the instrument and calibration and testing of any interpretational software.

The components were purchased prior to the beginning of this student project, and it must be noted that in the time since this the technology has progressed. Starting a similar project at the present time would obviously allow for a more technologically advanced instrument. However, the challenges faced in design and implementation would be largely the same. It would be fair to say that the construction phase of the instrument is where the wealth of information has been gained. The alignment, testing and data analysis phases of the instrument would be largely the same with more advanced detectors.

The design of the interferometer dates to the 19th century, but its advantages over other spectroscopic designs such as prisms and gratings have meant that as technology has caught up in data interpretation it has become an important instrument for 21st century astronomy. Even at the small scale, the instrument can provide useful science.

Chapter 4

SLIDE – Spire Line IDentification and Extraction

4.1 Introduction to SLIDE

SPIRE is the imaging FTS spectrometer instrument aboard Herschel and is based upon a Mach-Zender design, (Griffin et al., 2006; Pilbratt et al., 2010; Griffin et al., 2010). Herschel is an ESA space observatory with science instruments provided by European-led Principal Investigator consortia and with important participation from NASA. This design uses folded arms and two beamsplitters. Graphically this is represented in Figure 4.1. Here M represents a mirror, BS are the beamsplitters and BCC is the back-to-back corner cube reflector. F and D are the filter and detector respectively. BB is a black-body. By folding the arms, one can change the optical path difference of both paths simultaneously with a single movement mechanism.

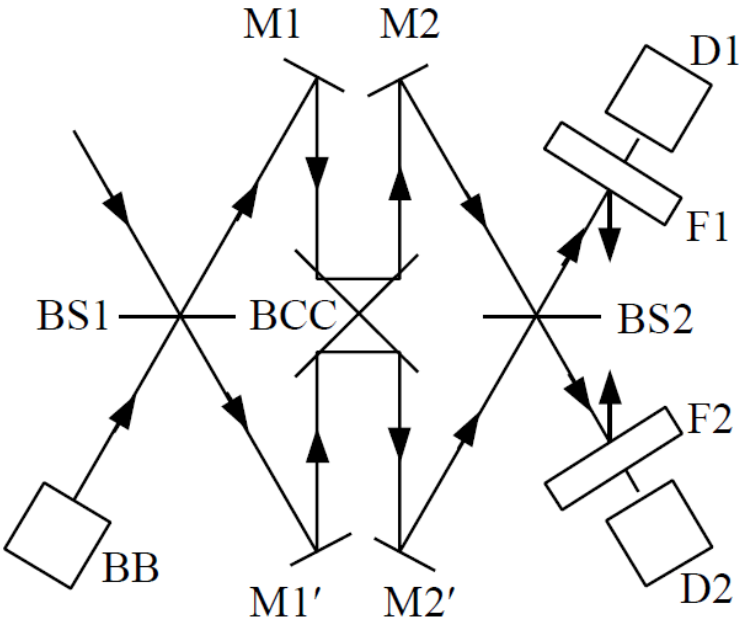

Figure 4.1: Mach Zender design for SPIRE, (Dohlen et al., 2000).

The bandwidth of SPIRE is split into two overlapping bands, SLW ($316 - 672\mu m$) and SSW ($194 - 324\mu m$). SPIRE has three resolution modes (*high, medium* and *low*) with unapodised spectral resolutions of approximately $0.04cm^{-1}$, $0.25cm^{-1}$ and $1.0cm^{-1}$ respectively. There is also a *high+low* mode where the high and low resolution observations are made simultaneously by truncating the interferogram to produce a lower resolution observation. SLW and SSW have separate arrays of detectors, packed into hexagonal arrays of 37 pixels for SSW and 19 pixels for SLW. Figure 4.2 shows the physical layout of the pixels. The coloured pixels are co-located between both arrays – and the large circles show where partial vignetting occurs due to the instrumental field of view. The dashed circles show the 2.6 arcminute unvignetted field of view for each array, and the solid outer circle shows the 2 arcminute field of view for the FTS. The beam spacing for SLW is approximately 48 arcseconds, and for SSW it is 27 arcseconds. Observations made using the central pixels *SLWC3* and *SSWD4* in Figure 4.2 basically reduce

to the simple, one-pixel FTS as their mutual centres lie on the optical axis. On this image, the two pixels D5 and F4 marked in grey are dead pixels.

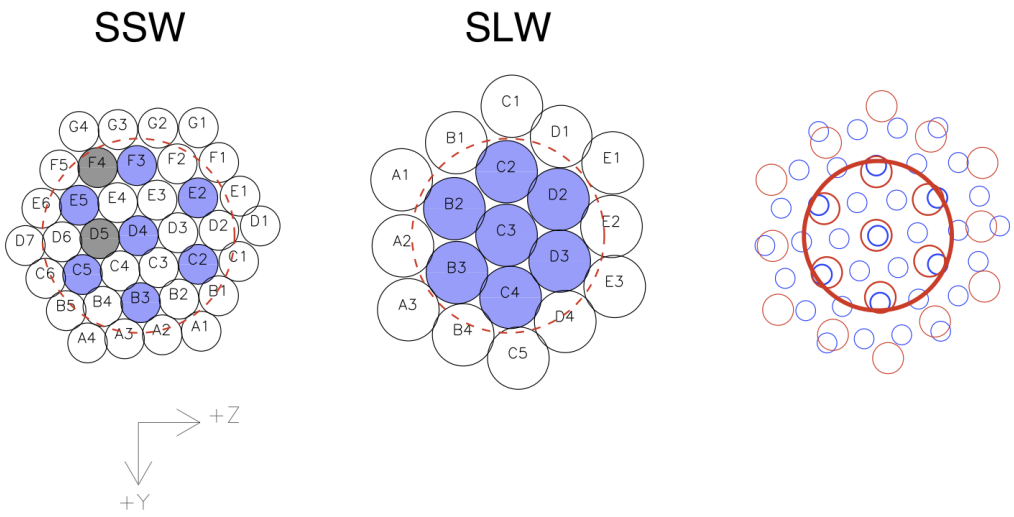

Figure 4.2: SLW,SSW Detector arrays from the SPIRE Observers Manual, European Space Agency (2010).

4.1.1 SPIRE Data analysis

The European Space Agency provides a tool for data processing called HIPE where data from SPIRE is processed to yield data products. HIPE is a joint development by the Herschel Science Ground Segment Consortium, consisting of ESA, the NASA Herschel Science Center, and the HIFI, PACS and SPIRE consortia, Ott (2010). These range in levels from raw unprocessed data (Level 0) to the final processed data (Level 2). Data can also be processed to Level 0.5 where the raw data has had the timeline processed, and Level 1 where some engineering processing has occured. Processing the SPIRE data through the pipeline results in the creation of FITS format files as final output.

It is possible to identify the data processing needing to be performed on the

output SED from HIPE for a source. Line fitting requires the identification of the continuum, as a baseline to subtract from the SED. Once we have subtracted the continuum we can then extract and identify emission and absorption lines that are present. For line identification we work with unapodised SEDs in order to have the greatest possible spectral resolution for line extraction. The sidelobes that are a result of the Instrumental Line Shape (ILS) function in the unapodised SED are useful as a fit feature when we use sinc line profiles. SLIDE was developed by myself to perform line fitting analysis on SPIRE data, see Section A.10. Currently SLIDE only processes SPIRE central pixel data, although the future plan for SLIDE is to merge it with another line fitting tool developed by the University of Lethbridge, (Jones et al., 2009; Naylor, 2010). This merged tool will be added to HIPE as part of the effort to develop a unified data analysis tool.

4.2 Continuum fitting

The continuum fitting available in SLIDE is not specifically designed for retrieval of continuum data for the source. Its primary function is to produce as close to a zero-baseline as possible for accurate recovery of line parameters. The choices for continuum fit include a greybody fit, as outlined in Colbert et al. (1999), where the form

$$F_\lambda = B_\lambda(T)(1 - \exp -\tau_{dust}) \tag{4.1}$$

is used with a Rayleigh-Jean form of the blackbody function $B_\lambda(T)$ as well as the full Planck Formula. The optical depth τ_{dust} is assumed to vary as λ^{-1}. An-

other greybody function is also available, from Priddey & McMahon (2001) for the Rayleigh Jeans form, and Kirk, Ward-Thompson, & André (2007) for the full Planck form. We define F_ν as the monochromatic flux density of the grey body at a frequency ν, $B_\nu(T)$ is the Planck function, ν_c is the frequency at which the optical depth is unity, Ω is the solid angle, f is the filling factor of the source within the observing aperture and β is the dust emissivity index. Our greybody equation in this case become

$$F_\nu = \Omega f B_\nu(T) \left(1 - \exp\left[-\left(\frac{\nu}{\nu_c}\right)^\beta\right]\right), \tag{4.2}$$

and β is a fit parameter, usually limited between unity and two in the literature – see Kirk et al. (2007) for an example. There is also the choice of a pure Planck function blackbody fit as well as a variable-degree (up to eighth-order) polynomial for cases where fringing or other artefacts distort the underlying continuum shape.

SLIDE uses a smoothing routine to generate a baseline that is used to fit the continuum, and this uses a polynomial routine. SLIDE also has the option to perform no baseline subtraction where a baseline has been externally subtracted from the data. Once a baseline -subtracted continuum is available this is passed to the line fitting routine, although this file is also written out should further processing be required. A supplementary tool with SLIDE can perform multiple polynomial subtraction from an SED for sources where the baseline is distorted – for example, where fringing dominates an SED with low signal-to-noise.

4.3 Line Fitting

The SPIRE pipeline has the capability to produce either apodised or unapodised output. Apodisation dampens the sidelobes from the spectral lines at the expense of line broadening. These sidelobes have the potential to be useful as a spectral feature for line fitting as they extend far from the line centroid. These features may be useful in generating accurate line parameters and so SLIDE has the option to fit several line types. One difficulty that has been found with fitting is the presence of fringing within the SPIRE spectra. The fringes are artefacts that are created as the result of the limitations of a real FTS compared to an ideal one.

4.3.1 Fringing and Line Fitting

We examined the effects of fringing on line recovery parameters using a function to simulate the fringing inherent in the FTS, (Savini, 2009). This involved the creation of ideal spectra based on a comb of sinc functions to simulate the effect of the SPIRE ILS on delta-line profiles, where the comb was set up such that a line was found in every $0.2cm^{-1}$ of the SPIRE Band. These idealised spectra were then multiplied by the fringing function to simulate these artefacts. Initially we added these line spectra to a greybody continuum to simulate real spectra and to examine how continuum fitting may affect the result. We further examined the SLW band by generating another matrix on a flat continuum. We then examined how well we recovered line parameters in SLIDE, using a wide range of line:continuum ratios. In this way, we attempted to define a *Threshold Detection Metric* that would characterize the quality of line recovery it is possible to obtain across the SPIRE band.

These tests revealed cases where SLIDE had issues with baseline subtraction using the original subtraction routine in cases where the lines are weak. This feedback led to improvements in the baseline fitting. In weak line cases SLIDE sees the fringing function strongly imprinted on the spectrum. The SLW band shows far stronger fringing, and for a large degree of the band a line:continuum ratio of 0.2:1 means that SLIDE does not see the lines. The SSW band shows isolated areas where fringing means large 1-σ errors (of around 25%) but fringing is not as pronounced as it was for the SLW.

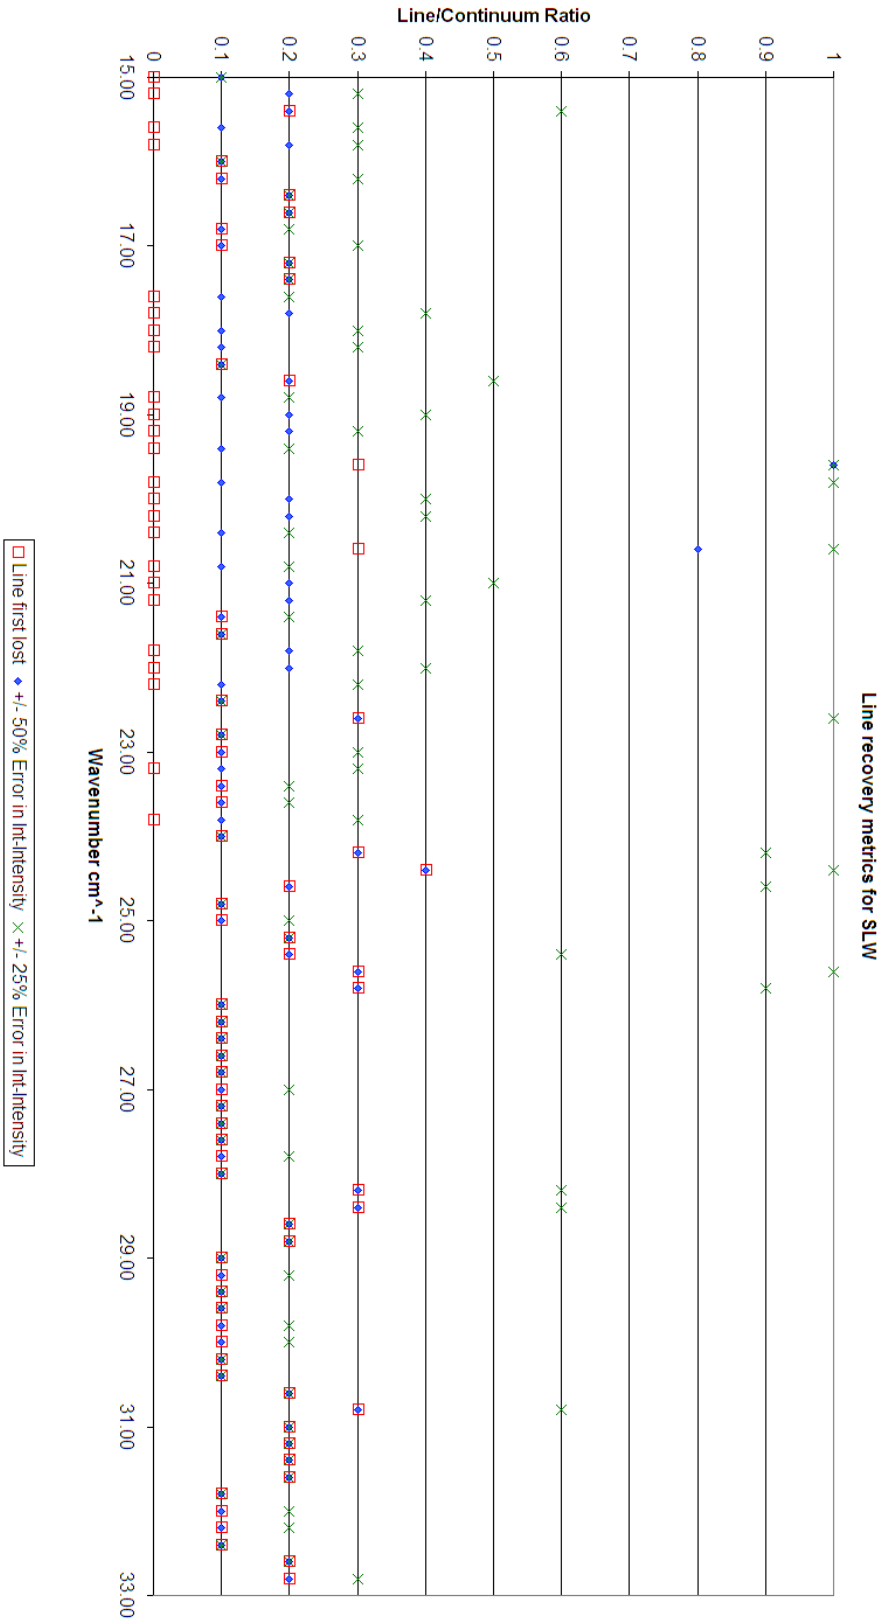

Figure 4.3: Line recovery quality metric for SLW.

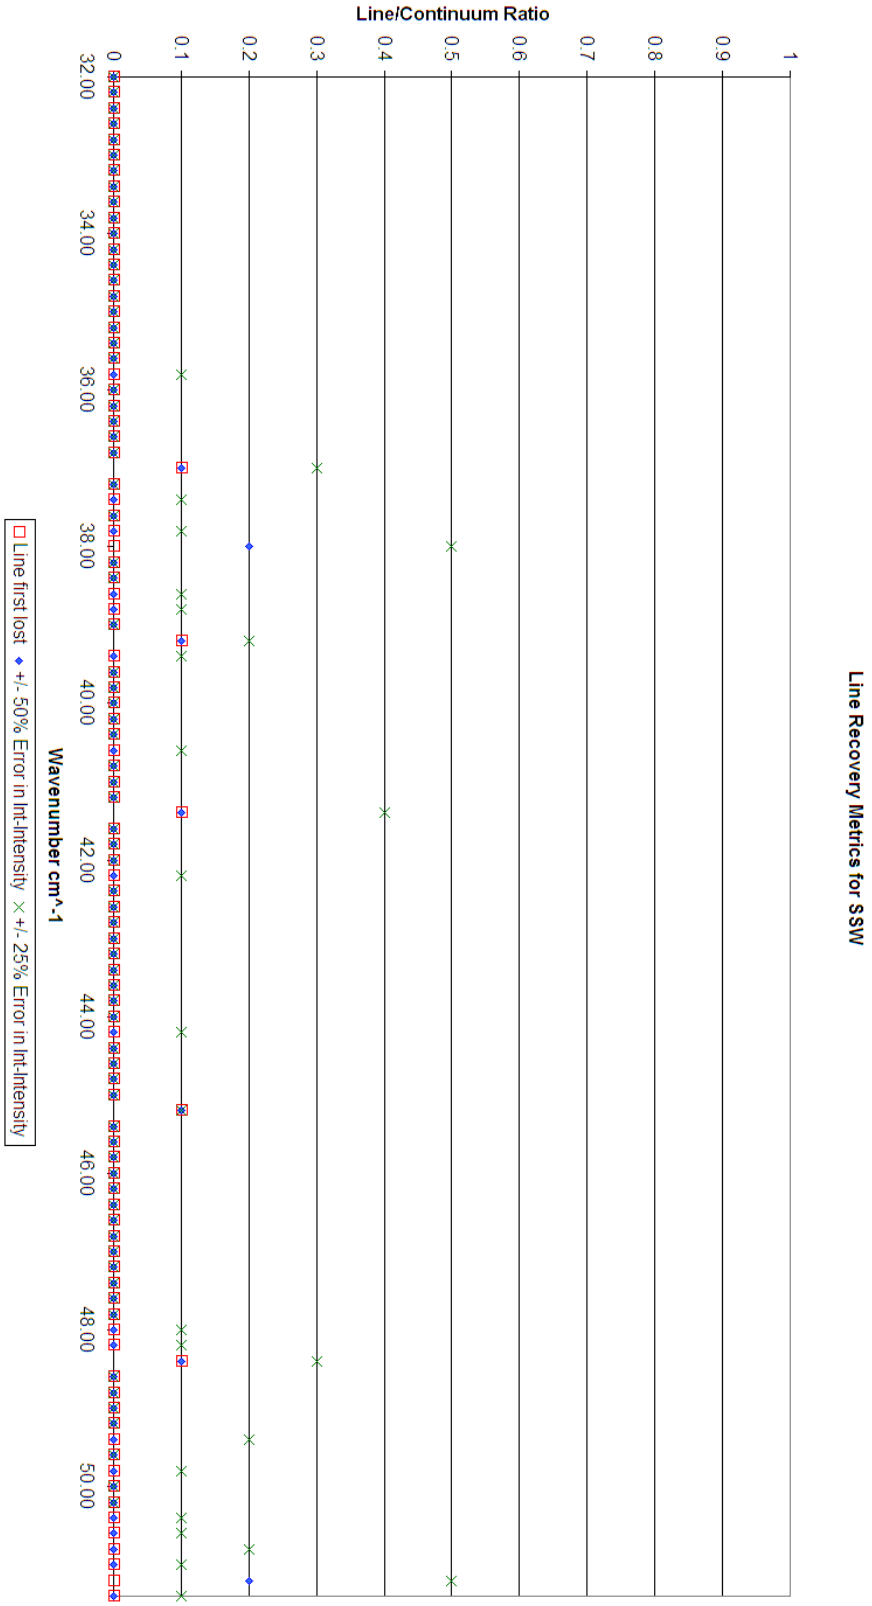

Figure 4.4: Line recovery quality metric for SSW.

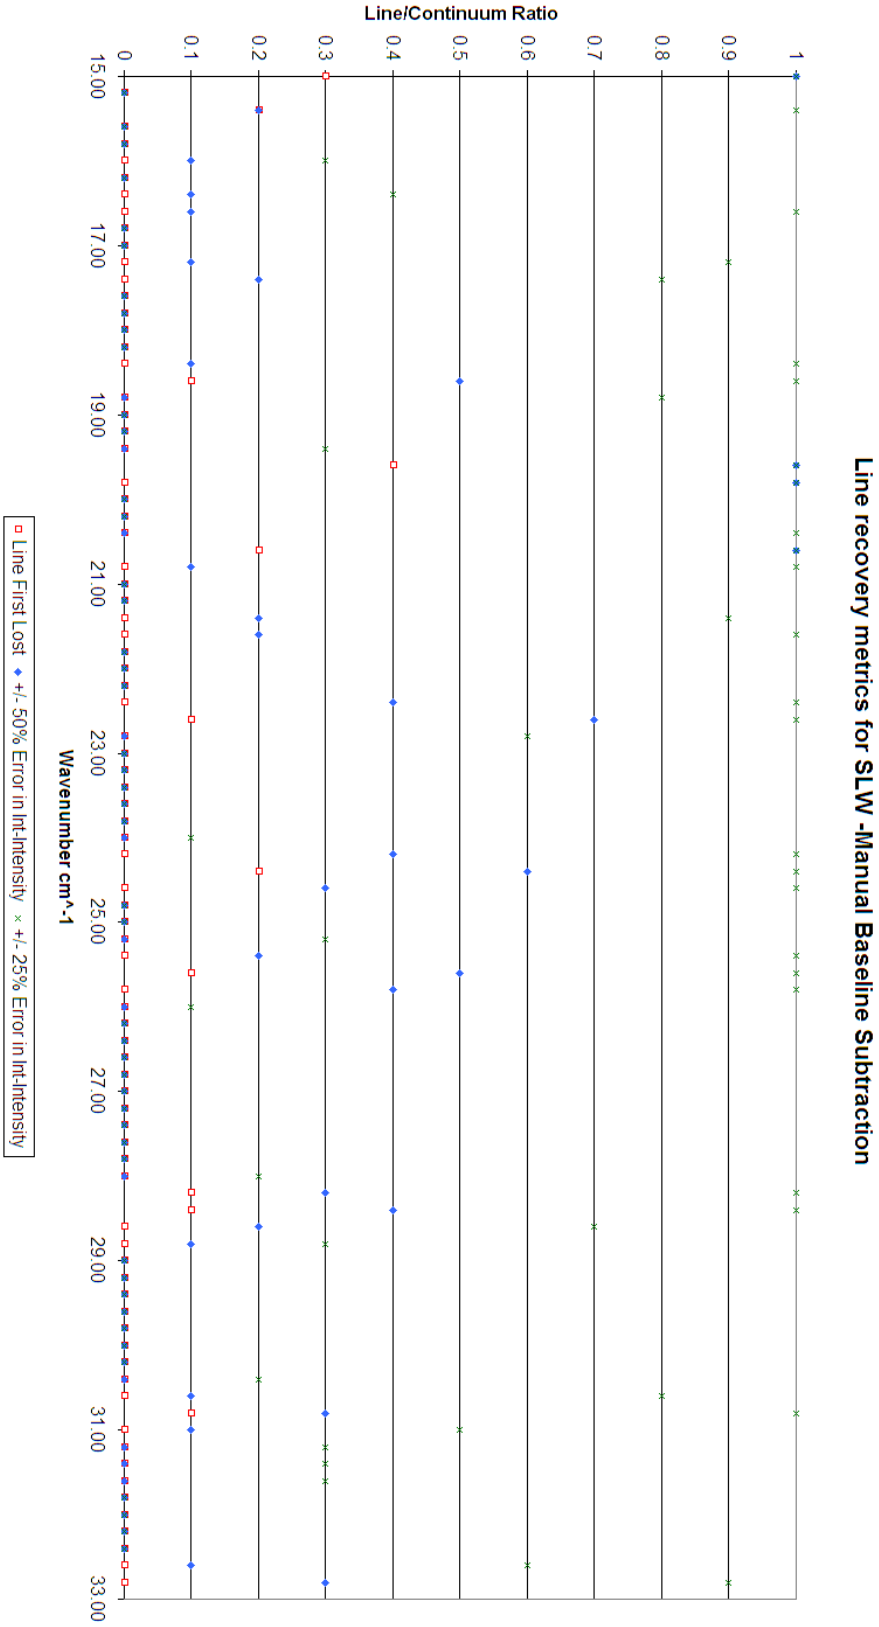

Figure 4.5: Line recovery quality metric for SLW - flat baseline.

Figures 4.3 and 4.4 show the line recovery metric graph for both bands. The x-axis lists wavenumber, and the y-axis is the line:continuum ratio at which:

Square the line is not recovered by SLIDE.

Cross shows a 25% error in flux.

circle shows a 50% error in flux.

We see from Figures 4.3 and 4.4 that it is much harder to recover accurate line parameters for the SLW band than the SSW. Below 20 wavenumbers, errors are approximately 50% for a wide range of line:continuum ratios for half the band. In comparison, SSW is far less affected by fringing and accurate line parameter recovery is possible across the band. We then examined SLW a second time, using a flat baseline to ensure that any results were not being unduly affected by baseline subtraction. Although some improvement was seen, it is still obvious that SLW is affected badly by fringing. In this simulation, there are areas where a line:continuum ratio of 1:1 show errors in flux recovery of approximately 25%.

For SLW the centroid positional error recovered by SLIDE, as shown in Figure 4.6, approaches the SPIRE nominal resolution in places of strong fringing. This becomes a zone of avoidance below 20 wavenumbers for line fitting as errors increase significantly. For SSW, these errors are around an order of magnitude less. Because SLW had more pronounced fringing issues, we examined the second case where a set of SLW data without an underlying continuum was used and examined the recovery in this case. We find that the line recovery metric is not significantly affected by baseline issues, and there are still large areas of the band with high flux errors.

Figures 4.6 and 4.7 show the absolute line centroid positional error across the band for a range of line:continuum ratios. Here we note that the wavenumber gridding of the SPIRE spectra is $0.01cm^{-1}$ and the nominal unapodised SPIRE resolution in HIGH resolution is $\approx 0.04cm^{-1}$ based on a maximum path difference of 12.645 cm. Both figures are deliberately in the same scale to highlight the differences between the bands. The area around 20 wavenumbers in the SLW band shows absolute errors spike to nearly the resolution limit, and the mean error is large for the lower part of the band. The SSW band shows a far flatter centroid error distribution with a mean of around a third of a wavenumber gridding element for the lowest line:continuum ratio of 0.2:1. Figure 4.8 shows that even with a flat baseline, fringing is still an issue for SLW and that the large errors cannot be explained solely by baseline subtraction issues.

To summarise, it is apparent that the area of the SLW band below 20 wavenumbers is prone to line recovery problems because of fringing. This must be taken into account when examining SLIDE parameters for spectra taken from SPIRE.

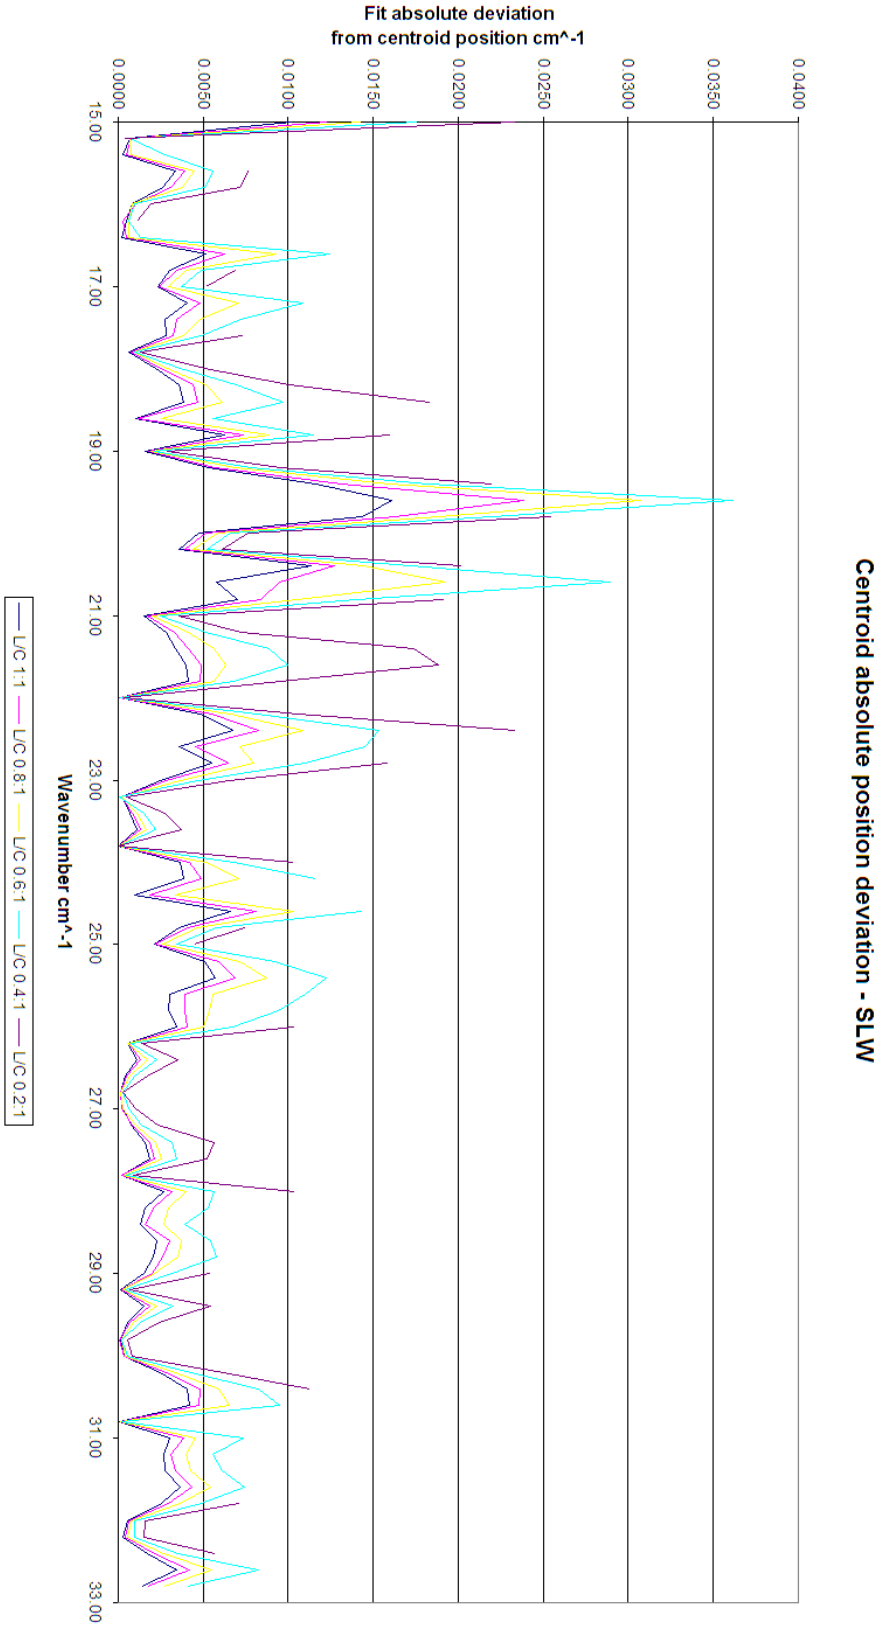

Figure 4.6: Line recovery quality metric for SLW - centroid position deviation.

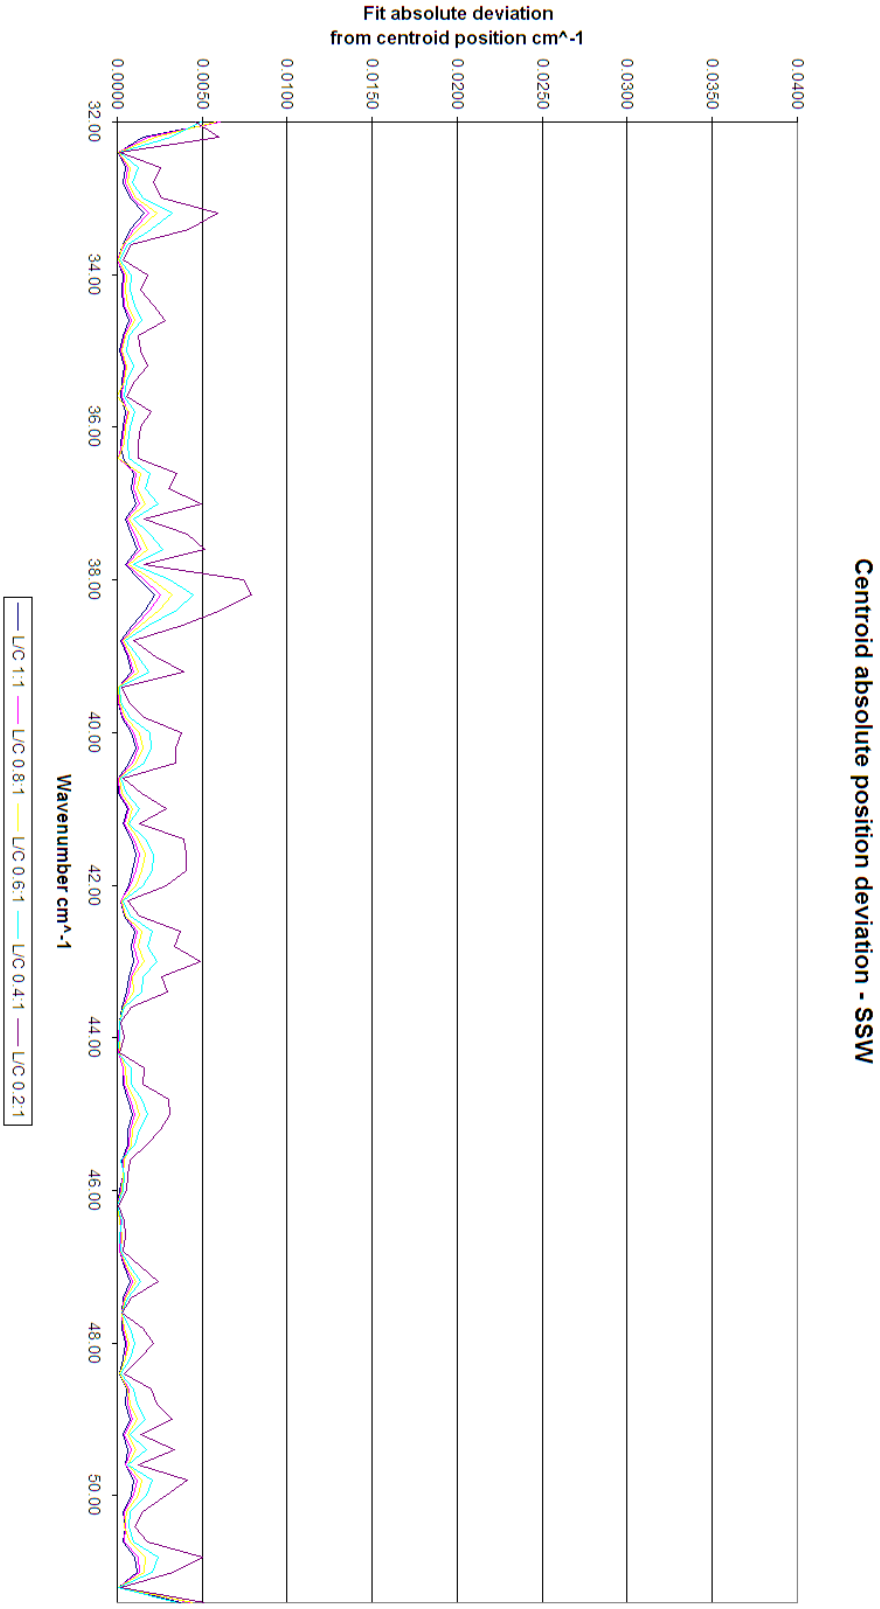

Figure 4.7: Line recovery quality metric for SSW - centroid position deviation.

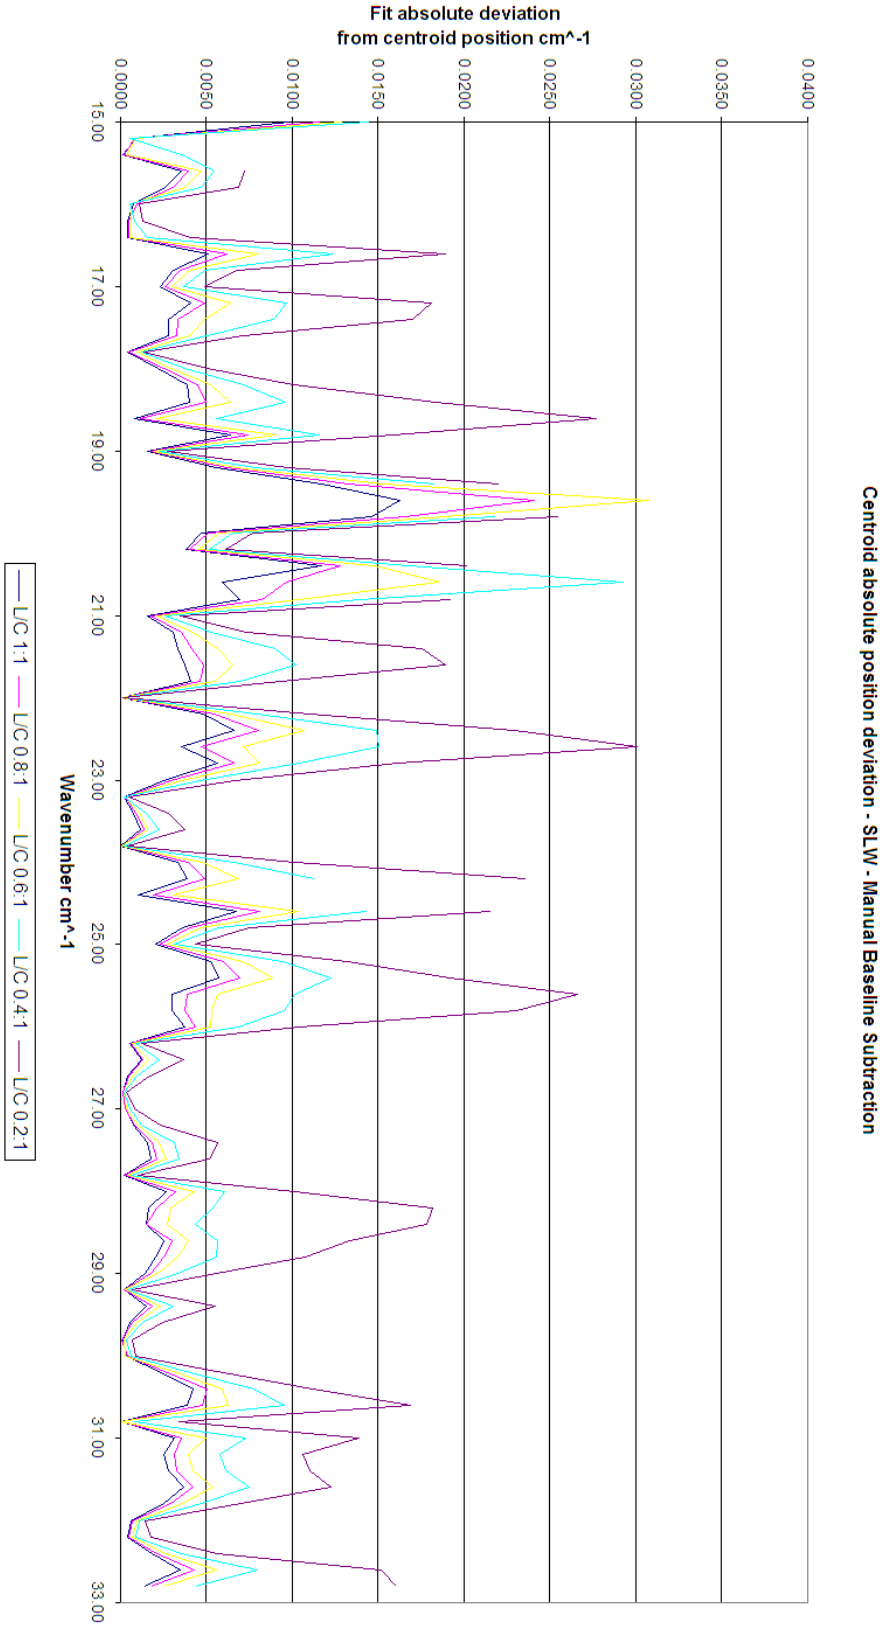

Figure 4.8: Line recovery quality metric for SLW - centroid position deviation for flat baseline.

4.3.2 Line fitting functions

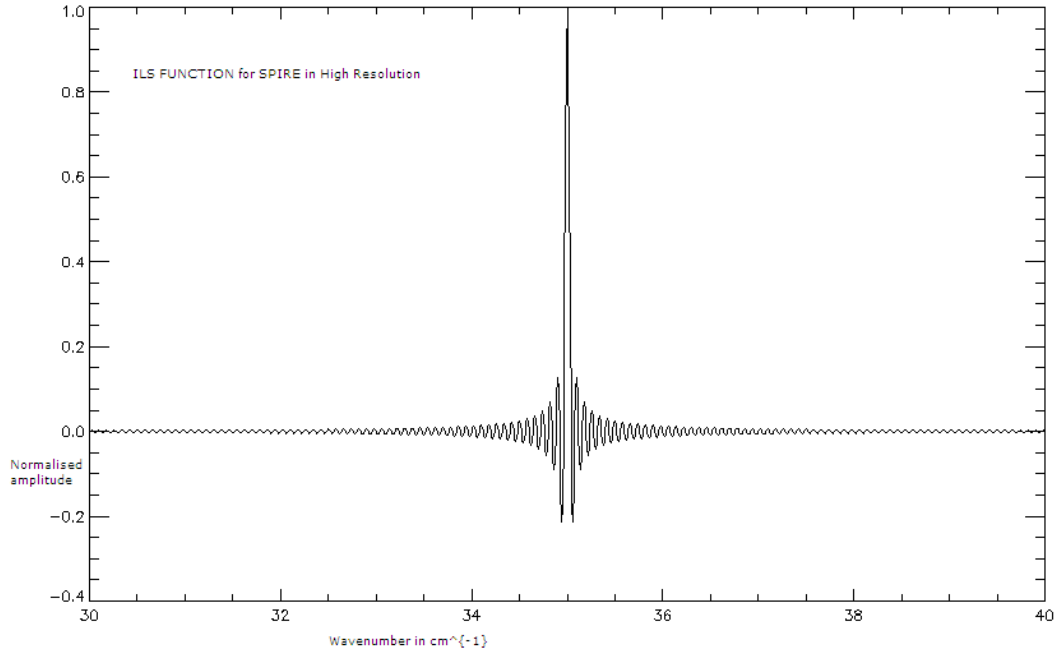

Figure 4.9: SPIRE theoretical ILS in high resolution.

For apodised spectra, the partial removal of sidelobes means that the resultant line shape is closer to a Gaussian line profile. As a result, SLIDE offers the choice of Gaussian line fitting in this case. This line shape for an apodised spectrum is still only an approximation to a Gaussian. It is necessary to take the instrumental line shape function (ILS) into account for unapodised spectra. For a finite-path-difference interferometer, the multiplication of the infinite path difference interferogram with a top-hat function is equivalent to a convolution in real space of the spectrum with a sinc function. For monochromatic lines, the resultant line profile is the convolution of a delta function with a sinc function. Figures 4.9 to 4.11 show this graphically.

For broad lines, we assume that the actual line profile is a Gaussian – and

this means that the resultant line profile is a composite profile that we term a Sinc-Convolved-Gaussian (SCG). If we compare the effect of widening the initial line profile from a delta function to a Gaussian on the resultant sinc-convolved line, we see that the output line profile evolves from a sinc to a profile that is distinct from either a Gaussian or a sinc. For an initial input of a narrow Gaussian, the output profile shows sidelobes similar to a sinc function– but the central line profile is wider than a true sinc function and cannot be accurately fitted by either a sinc or a Gaussian alone. Figures 4.12 to 4.15 show the progression in line shape from sinc-like to Gaussian-like over a range of widths of half- to twice-SPIRE resolution (in high resolution mode). In between these ranges, we get an intermediate shape which is neither a true sinc nor a true Gaussian. This is one of the challenges to be overcome in line fitting and recovering accurate line parameters.

Figures 4.16 and 4.17 show a direct comparison to the sinc-convolved Delta function and a sinc convolved Gaussian. As the full width at half maximum (FWHM) increases, the resultant line profile looks like neither a true sinc nor a true Gaussian. For larger FWHM values the profile approaches a Gaussian-like form, and conversely for smaller FWHM values it approaches a sinc-like form. Hence the provision of SCG fitting means that SLIDE can allow for the instrumental effect of the line shape function on variable-width Gaussian line profiles in its flux calculations.

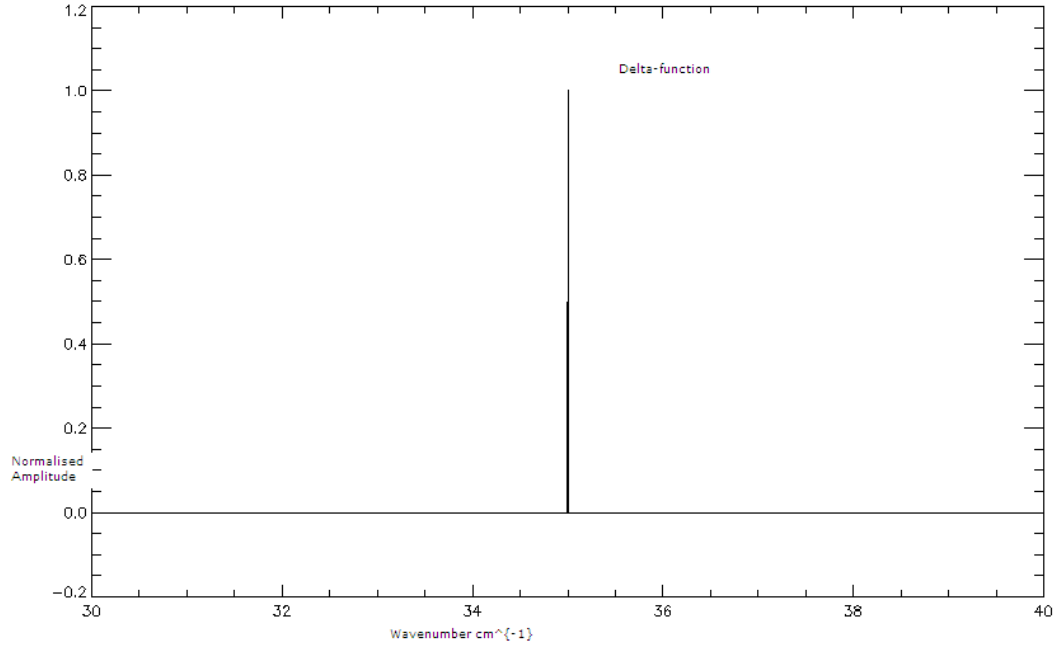

Figure 4.10: Delta function centred on 35 wavenumbers.

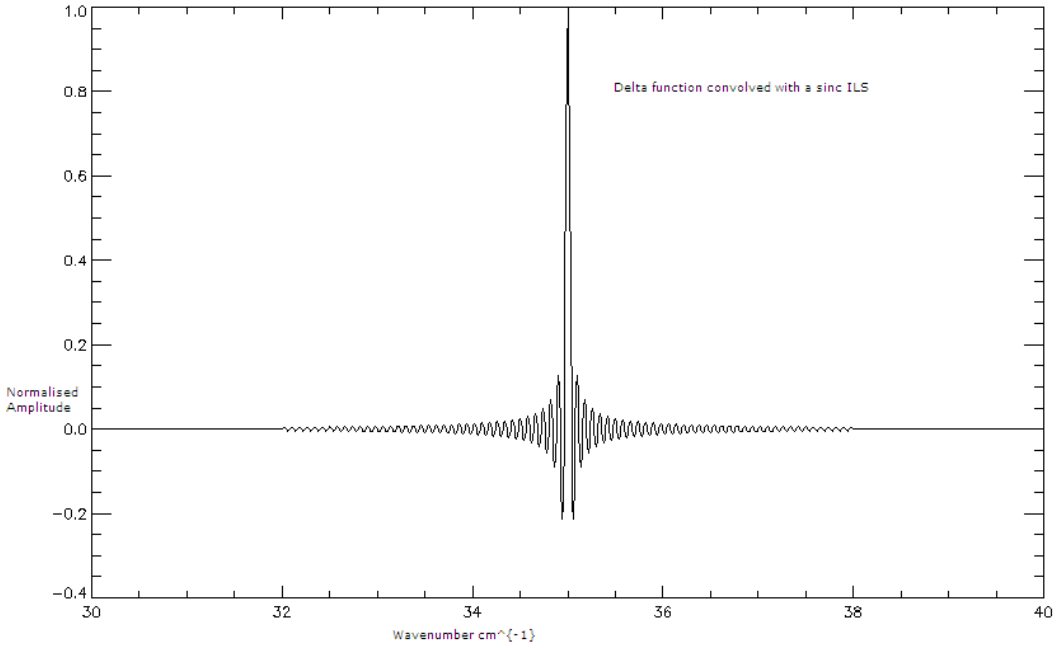

Figure 4.11: Result of convolving the ILS with a delta function.

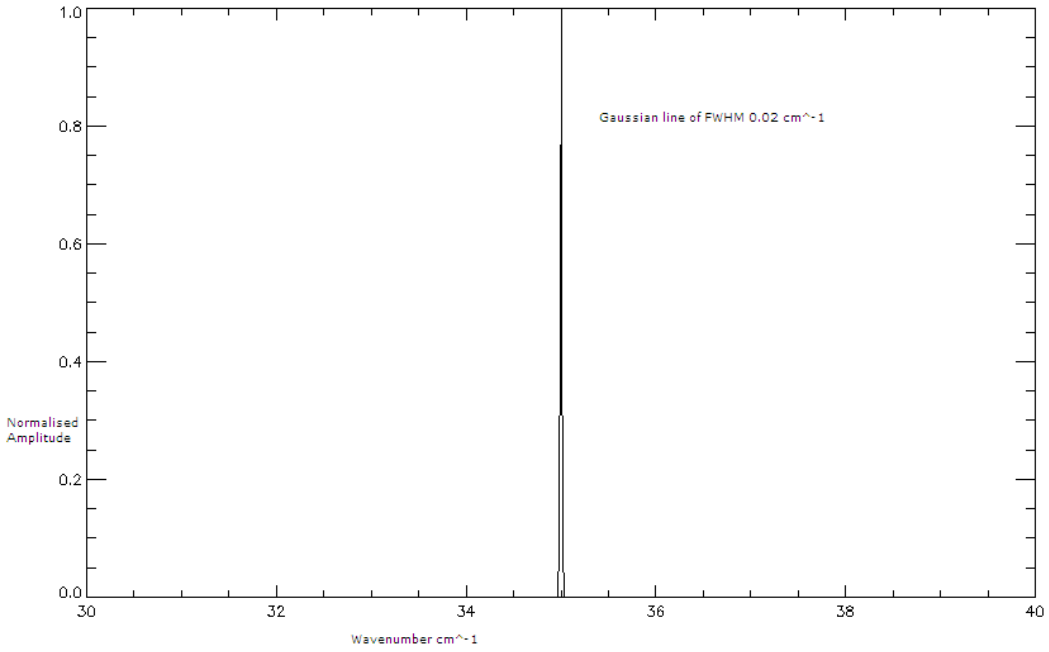

Figure 4.12: A Gaussian line, with a FWHM of half the SPIRE resolution (High).

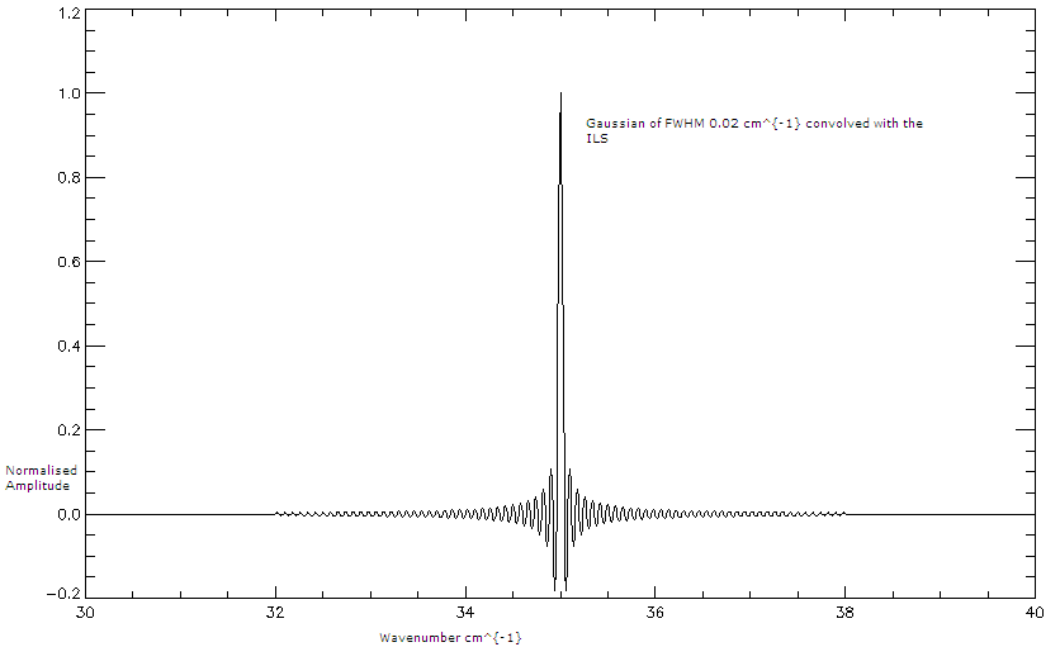

Figure 4.13: Result of convolving the ILS with the half-SPIRE resolution Gaussian.

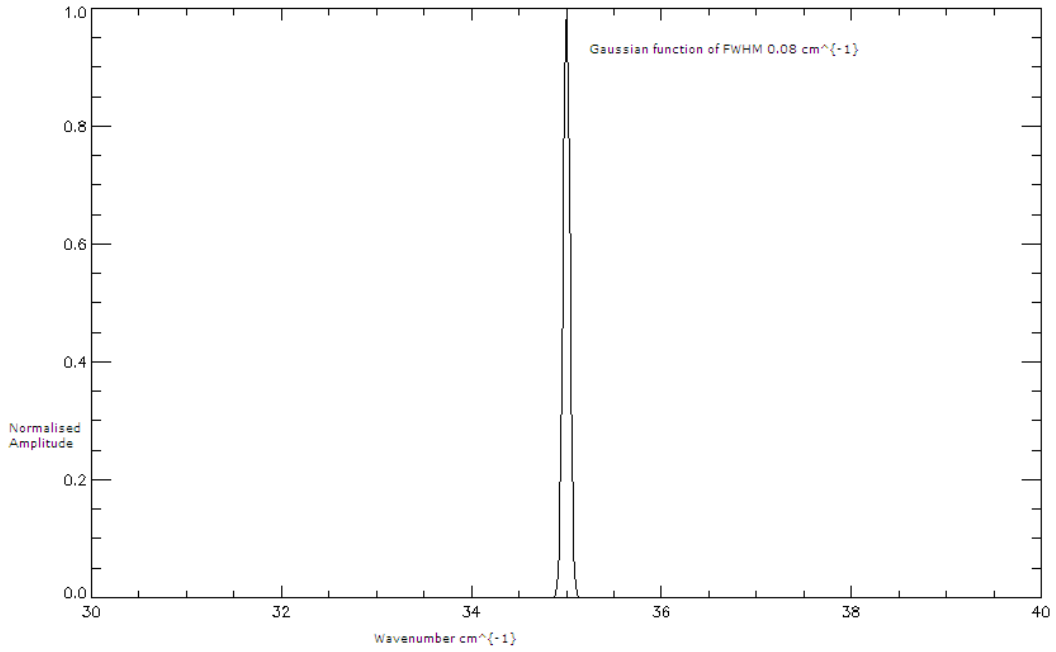

Figure 4.14: A Gaussian line, with a FWHM of twice the SPIRE resolution (High).

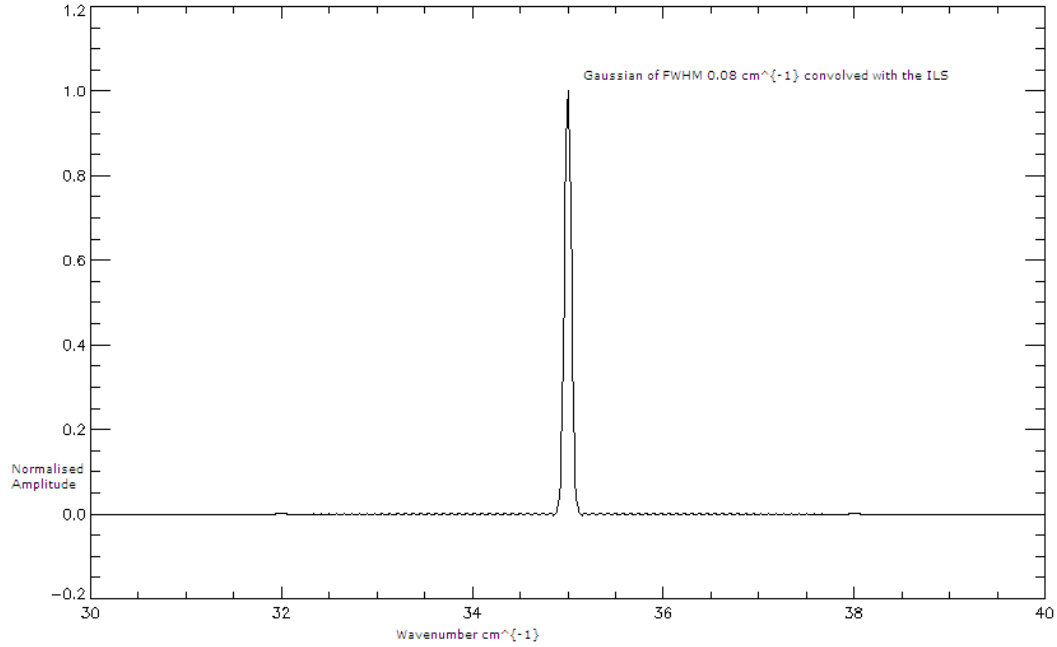

Figure 4.15: Result of convolving the ILS with the twice-SPIRE resolution Gaussian.

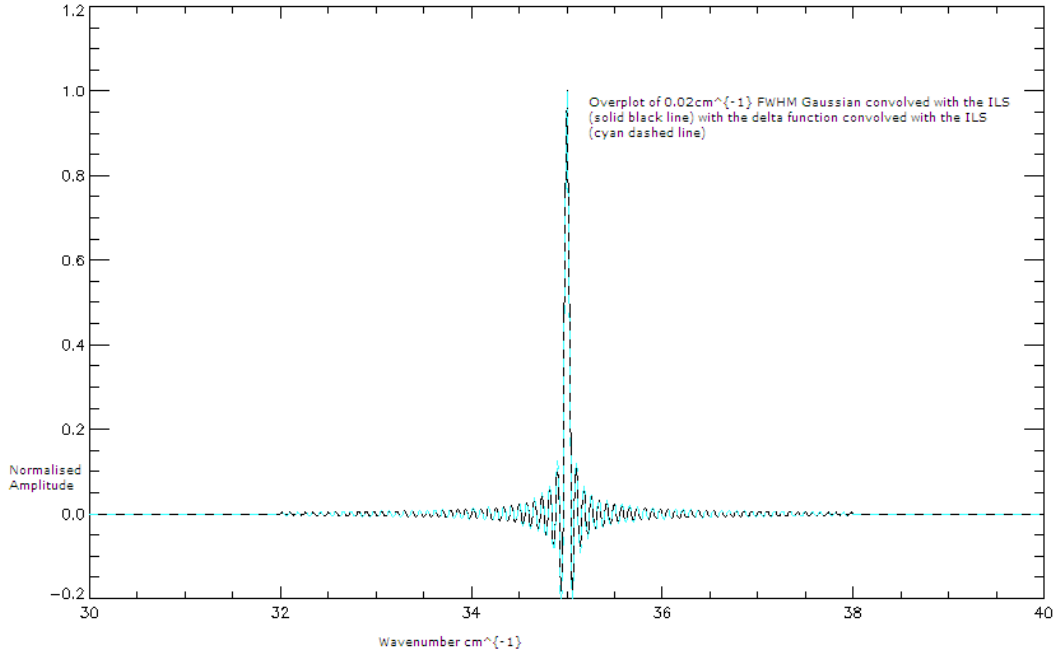

Figure 4.16: Comparing a convolved delta-function to a half-SPIRE resolution Gaussian.

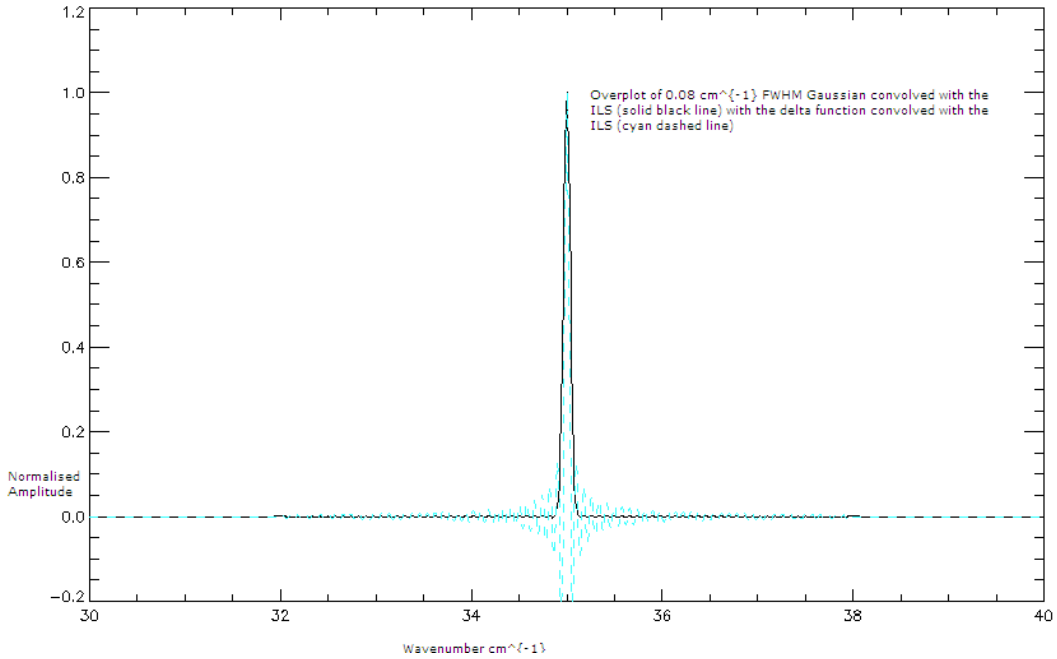

Figure 4.17: Comparing a convolved delta-function to a twice-SPIRE resolution Gaussian.

4.3.3 Binning and the SCG Profile

SLIDE can use a lower limit on the Gaussian FWHM to be $0.04cm^{-1}$ for the Gaussian component of the Sinc-Convolved Gaussian fit – that of the spectral resolution of the unapodised high resolution spectrum. This is not compulsory, but there is a risk in ignoring this limit. Any lower than this, and the Levenberg-Marquardt based MPFIT routine used (Markwardt, 2009) generates increasingly large 1–σ errors. This is largely due to the binning resolution of the spectral data.

The sinc-convolved Gaussian line profile $L(\sigma)$ is generated by creating a Gaussian line $G(\sigma, A, w, \sigma_0)$ of amplitude A, FWHM w and line centroid position σ_0. We then convolve it with a normalised sinc line profile $S(\sigma)$ representing the instrumental line shape function and this yields

$$L(\sigma) = S(\sigma) \star G(\sigma, A, w).$$

This method of line generation models the physical process that happens as a line with a Gaussian profile enters the FTS. SPIRE high resolution spectra are supersampled to $0.01cm^{-1}$ for a nominal spectral resolution of $\approx 0.04cm^{-1}$. When fitting sinc convolved Gaussian line profiles the fit routine returns significantly large errors for the parameters of the Gaussian component when the line width decreases below this supersampled resolution.. These errors are representative of the difficulty in fitting a composite function as the width decreases below the wavenumber gridding of the SPIRE data. In that case, it is impossible to fit a unique amplitude and FWHM for the underlying Gaussian, or line centroid position. This is where the Gaussian is within one bin of the data and it becomes impossible to fit to the Gaussian. As convolution is an overlap function, this means that any fitting routine will produce large errors in amplitude and FWHM as a

result.

SLIDE calculates the errors in flux as the quadrature sum of the errors in the parameters of the constituent sinc and Gaussian line profile. This can only be a first-order approximation as the calculation of true error is complicated by the presence of a convolution. For cases where the Gaussian line profile is contained largely in a single data bin, we saw there is a increasingly large parameter space for the amplitude, FWHM and centroid position of the constituent Gaussian. As the Gaussian width approaches that of the bin size, then the likelihood of large errors for the flux are a result.

We defined the SCG line profile such that it is generated by directly modelling the process by which a Gaussian line is modified by the ILS of the interferometer. It is possible to remove the large errors from the amplitude by convolving a unit Gaussian with the ILS and then rescaling it to fit the line. This removes the amplitude of the Gaussian to outside the convolution and reduces the large error bars in amplitude, and is acceptable because multiplication is commutative within convolution, i.e for a unit Gaussian g and a non-unit amplitude Gaussian G it is possible that:

$$S \star G \equiv A(S \star g).$$

Here $\star$ is the convolution operation and A is the term to rescale the output of the right-hand side to make it the same as the left. Although this method removes the large errors from the amplitude of the Gaussian, it does not remove large errors entirely. As the binning of the data results in lost information we are merely trading error bars between parameters, and the FWHM of the Gaussian using this method has much larger error bars than the FWHM of our orignal method – we

are in effect seeing a *conservation of error* type of effect. We also see a larger error in centroid position when we are within the bin because we are dealing with a Gaussian almost completely contained within one bin of the data. We are left with a loss of information about the Gaussian which makes such large errors in the fit parameters an unavoidable conclusion. Since our first-order error calculation of flux error treats convolution as just another term, then we end up with large flux errors which are not truly representative. There is no simple way to calcualate errors with the convolution process included. One way to get a representative value for the errors is to fit a sinc profile to the line, and use the sinc flux $F_s(\sigma)$ and flux error $E_s(\sigma)$ returned, with the SCG flux $F_{scg}(\sigma)$, to get a rough estimate of the SCG error $E_{scg}(\sigma)$ by:

$$E_{scg}(\sigma) \approx E_s(\sigma)\frac{F_{scg}(\sigma)}{F_s(\sigma)}. \tag{4.3}$$

When the Gaussian is fully contained within the bin, then it appears in the *binned data* as a delta-function. Convolving a delta function with a sinc should return a true sinc function. However the SPIRE observation of this particular Gaussian is not a true delta-function. The resultant line profile is a binned convolution of a 'thin' Gaussian with the ILS. If we were to improve the resolution and binning of SPIRE, then we could retrieve accurate line parameters in this case. All this would do is merely reduce the problem to a thinner Gaussian width. In general, what we are describing is a fundamental lower limit to the accurate retreival of Gaussian line parameters with SPIRE when using sinc-convolved Gaussian line profiles.

With sinc line profiles, this problem is not encountered. Although the data are binned, there may be enough information in the line profile to fit a sinc to such accuracy that one can fit the line centroid to within 10% of the resolution

element. The sidelobes of the sinc profile extend from the centroid position so as to provide additional features to tie down a fit. This is not the case with a Gaussian. The problem is encountered where the Gaussian FWHM is of order of the data bin width, although it may be seen for widths only a little greater then it. Although the fit may visually be good, the associated flux has large error bars due to the nature of the calculated error, in treating convolution as a simpler mathematical concept than it truly is. To further improve fit recovery, SLIDE also has a 'composite' fit function. Each line is fitted simultaneously by a sinc profile and a SCG profile. The choice of whether to define the line as sinc or SCG is decided by using a χ^2 metric on the residuals after subtraction of the line.

We have discussed earlier in this chapter the shape of SCGs as the Gaussian width evolves. A delta function convolved with the ILS returns a sinc line profile. As the width of the line we convolve with the ILS increases, the line profile evolves in such a way as to lose the sidelobes and become Gaussian-like. By testing both line profiles simultaneously, SLIDE can accurately predict line fluxes for both cases. There is however one caveat. In examinng finite-width lines for a narrow range of FWHM (between approximately 0.02 and 0.01 cm^{-1}) the recovery percentage drops to around 80-90% in tests with noiseless spectra and single lines. It is this region of FWHM parameter space where there is a problem, although this is easily surmounted.

By using a χ^2 selection technique, we avoid the increasing errors when trying to fit narrow-Gaussian SCG lines. These widths are below the resolution width of SPIRE and of the order of the wavenumber gridding of the SPIRE data ($0.01cm^{-1}$). What we see in this case is the difficulty in fitting the SCG in this region. The act of convolving the Gaussian with the sinc leaves a composite function whose behaviour is not easily fitted, and the MPFIT routine (Markwardt, 2009) is not

capable of accurately distinguishing these narrow-band SCG's. As a result, the 1–σ errors in the fit parameter, and therefore the line flux, are large. If the width limit for SCG's is removed, then fit recovery percentages improve – although with the risk of increasing errors in amplitude and FWHM. As a result, visually the fit is very good but we lose confidence in the recovered values of the Gaussian line amplitude and the FWHM. This becomes an issue should accurate fluxes be required, as the Gaussian amplitude and FWHM are defining characteristics of the line profile that enters the interferometer to be convolved with the ILS.

Table 4.1 shows an example set of test results from noiseless test data, using single lines and multiple lines in close proximity to examine sidelobe interaction between lines. Using the test parameters one is able to clearly define the FWHM and integrated intensity of the individual lines and then examine how well SLIDE recovers the line width and flux in each case. As line widths increase (a σ of $0.17cm^{-1}$ is equivalent to a FWHM of $0.04cm^{-1}$ – the nominal SPIRE high spectral resolution mode) we see the strength of the composite fitting routine. For delta function profile lines, which appear as sinc's when convolved with the SPIRE ILS, we see the fit routine preferentially choose sinc fits. As the line widths increase, efficiency of recovery dips to around 95% as we constrain the lower width of the SCG fit and a sinc fit is used even though the line profile is not truly a sinc profile. Adding multiple lines, in close proximity to simulate the effect of sidelobe interaction between strong lines, this recovery percentage drops to around 90% for the composite approach. However, as line widths increase, the composite fit now preferentially chooses SCG fits.

Single emission line	% age I.I (Cmp)	% age FWHM	% age I.I (Sinc)	% age FWHM	% age I.I. (SCG)	% age FWHM
Delta function	100.00	100.00	100.00	100.00	115.66	112.89
$\sigma = 0.005cm^{-1}$	97.19	97.87	97.19	97.87	112.40	110.84
$\sigma = 0.0075cm^{-1}$	94.01	94.59	94.01	94.59	108.70	107.53
$\sigma = 0.01cm^{-1}$	104.08	103.63	90.04	90.73	104.08	103.63
$\sigma = 0.02cm^{-1}$	100.00	100.00	74.08	71.98	100.00	100.00
3 close spaced emission	mean	mean	mean	mean	mean	mean
Delta Function	100.01	99.91	100.01	99.91	115.55	112.70
$\sigma = 0.005cm^{-1}$	89.05	97.79	89.05	97.79	102.83	110.66
$\sigma = 0.0075cm^{-1}$	98.97	98.74	94.09	94.49	108.67	107.35
$\sigma = 0.01cm^{-1}$	95.23	103.92	82.44	91.01	95.23	103.92
$\sigma = 0.02cm^{-1}$	100.10	100.06	71.95	70.89	100.10	100.06
4 lines (ABS,EM)	mean	mean	mean	mean	mean	mean
Delta function	100.51	100.31	100.51	100.31	115.35	113.49
$\sigma = 0.005cm^{-1}$	97.64	98.23	97.64	98.23	112.06	111.40
$\sigma = 0.0075cm^{-1}$	98.10	98.11	94.24	94.67	108.35	107.98
$\sigma = 0.01cm^{-1}$	104.17	103.95	90.22	90.72	104.17	103.95
$\sigma = 0.02cm^{-1}$	100.12	100.12	72.62	71.20	100.12	100.12

Table 4.1: SLIDE recovery effectiveness on noiseless spectra.

SLIDE performs sequential line recovery using the amplitude of the line centre to choose line recovery order. SLIDE iteratively recovers and subtracts the strongest line in the spectrum until no more lines of amplitude larger then a set threshold are to be found. The fit routine gives 1–σ errors for the fit parameters, and the integrated flux is calculated using numerical integration on the resultant line. The error in integrated flux is then calculated using a standard error formula. For the SCG fit, we make the first-order approximation that the error in the integrated flux is the quadrature sum of the error in the flux in the component Gaussian and Sinc profiles. This approximates the SCG to two Gaussian functions convolved with each other, and the cascade property in this case states the variances add in quadrature. This is because there is no simple analytic form to the area of a SCG. The code can then return line parameter information and continuum information for the input spectrum, using multiple passes if needed.

A line-fitting tool currently being developed by the University of Lethbridge, (Jones et al., 2009), was used to test SLIDE in its initial incarnation, and the information retrieved was used to improve both line-fitting routines as weaknesses in approach were identified. For example, should the code be used to dig deep into noisy spectra to extract many lines there comes a point where the noise causes a feedback-cycle effect in the fitting. The code overfits to a line and the overfit results in a negative feature, which the code overfits to to form a positive feature. This cycle repeats until the code breaks out as there is a set upper limit of lines it will recover. This only appears when the code is used to extract lines in very low signal-to-noise SEDs. One workaround is to run through the spectrum once with SLIDE, and then use the residual file that SLIDE outputs as an input to a second SLIDE run.

This does not reduce the code's efficiency or effectiveness in returning line pa-

rameters. The code has been used on a variety of papers investigating extragalactic sources such as González-Alfonso et al. (2010a) – revealing rotational lines of water emission in Mrk231, Panuzzo et al. (2010a) – used in observations of M82's prominent CO ladder in its central region and van der Werf et al. (2010a) – investigating CO and H_20 lines and exploring x-ray heating by an accreting supermassive black hole. SLIDE development and refinement is a continual process and it is released openly as a tool for the astronomical community. Section A.10 provides a fuller list of SLIDE usage as well as information for retrieving a copy of the code.

4.4 Validating SLIDE

4.4.1 Comparison with other tools

CASSIS (*http://cassis.cesr.fr/*), developed by CESR-UPS/CNRS, is a graphical line fitting program. At the time of the comparison it was in version 2.3, (Walters et al., 2005; Caux et al., 2011). We compared the behaviour of SLIDE to this tool to compare how well SLIDE recovers line parameters.

A comparison was undertaken using a SPIRE observation of NGC7027 as a control – a galactic source with unresolved lines and a clearly defined ^{12}CO ladder. With unresolved lines, we are able to test sinc-profile line recovery. The CASSIS line extraction used a fixed FWHM of $0.04725cm^{-1}$ with the other parameters free. SLIDE was run with all parameters floating, as this is the default for SLIDE. For SLIDE, the parameter L is the maximum path difference of the interferogram (therefore $L = \Delta z$) and sets the resolution of the resultant spectrum. For line fitting, we saw in Chapter 2 that this also defines the width of the sinc function for

fitting. As a result, SLIDE can set this parameter to float or be fixed. Allowing it to float by a small percentage can be useful to help constrain the fit of the line profile. Analysis of the results was done in two stages. A 'large scale' examination concentrated on global parameters such as reported positions and amplitudes, and then a 'small scale' one concentrated on fractional differences between parameters.

Figure 4.18 shows the comparison between line centroid positions found by CASSIS and SLIDE, with a line of best fit applied using a least-squares fit with a set y-intercept of zero. From this, we can see there is only a minimal gradient difference from unity, of the order of 10^{-5}. Although error bars are included on this plot, they are of such small size that they are hardly visible in this scale plot. As far as the absolute line centroid positions are concerned, SLIDE and CASSIS agree well with each other and the deviation from unit gradient slope is negligible.

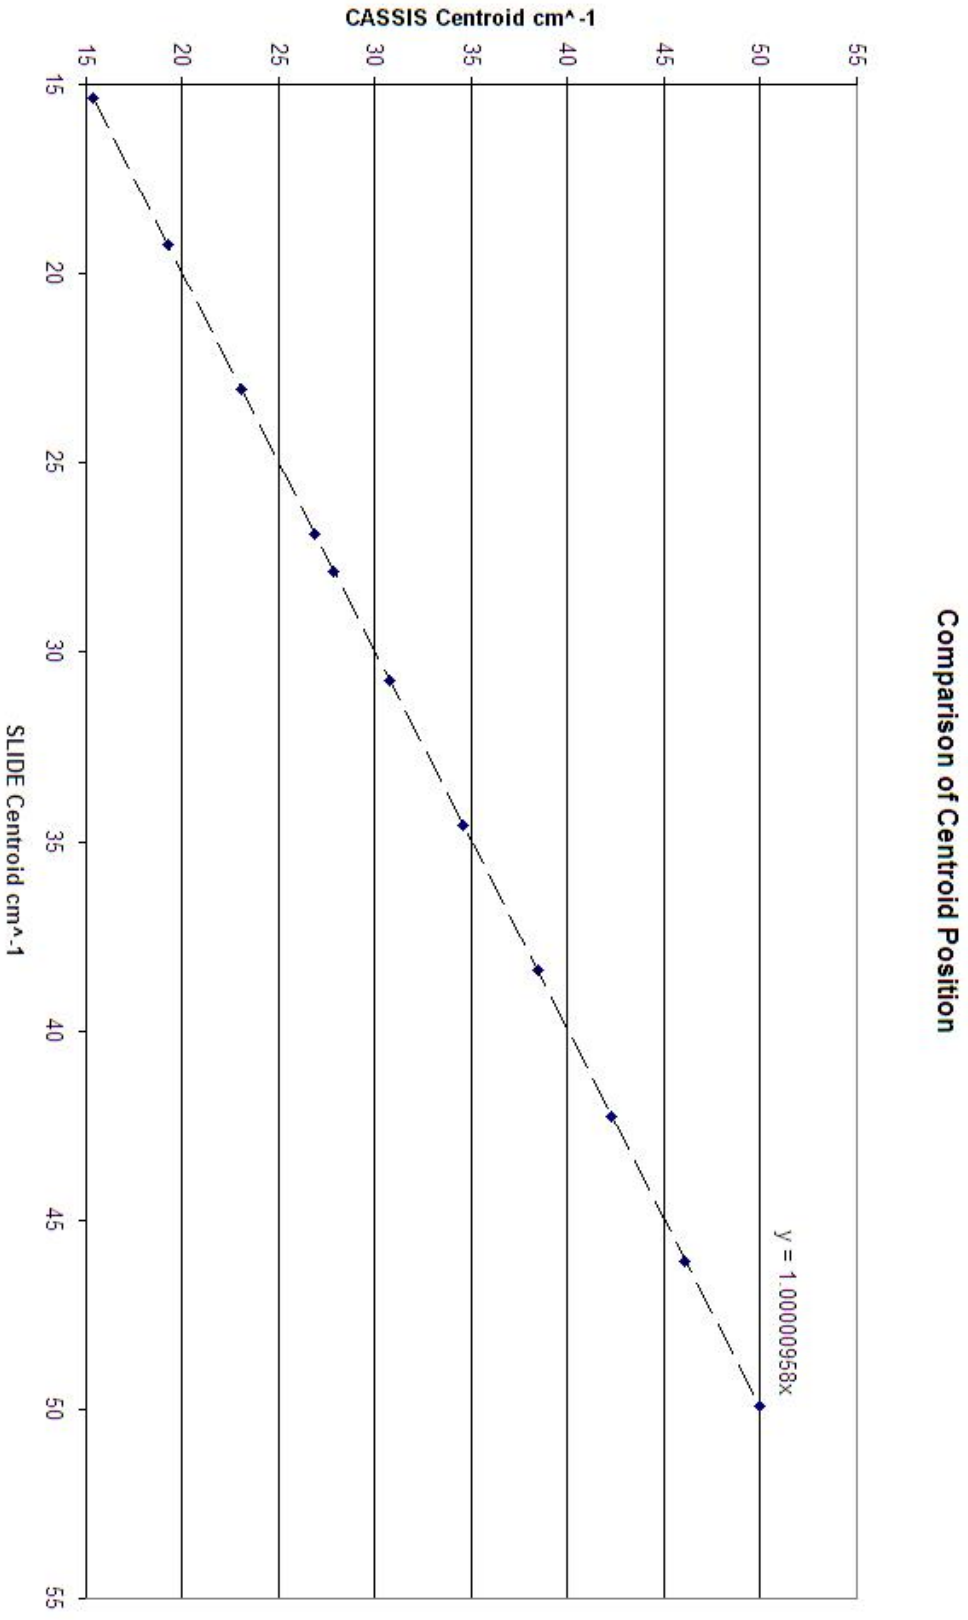

Figure 4.18: Comparison of reported line positions between SLIDE and CASSIS.

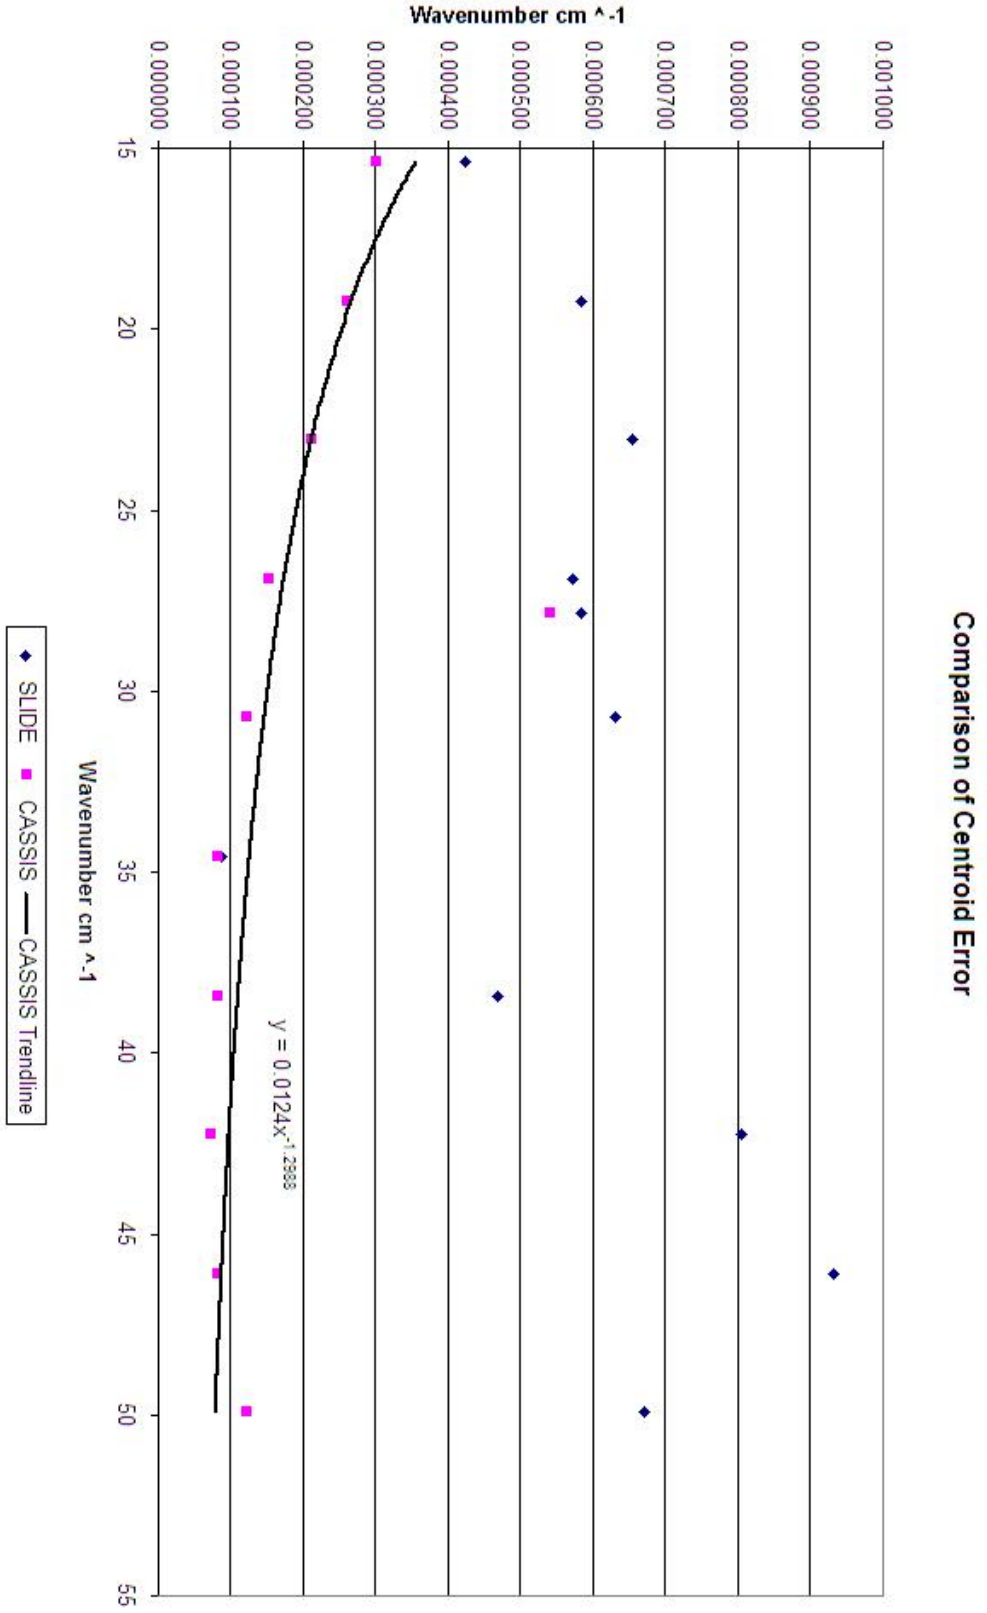

Figure 4.19: Comparison of reported line positions errors between SLIDE and CASSIS.

Figure 4.19 shows the $1-\sigma$ errors in the line centroid positions as reported by both CASSIS and SLIDE. Here we compare the error bars for the CO and CI lines in the test spectrum. Apart from the CI line at $27cm^{-1}$, there seems to be some relationship between error and centroid position for CASSIS. This seems to show a power-law type relationship with a decreasing error as centroid wavenumber increases. The deviation from this relationship for the one CI line may be due to its close proximity to a ^{12}CO line (within $0.1cm^{-1}$). This seems to show an increase in confidence in the fit as wavenumber increases for CASSIS. This may be likely due to differences in baseline subtraction and a resultant slope applied to the spectrum used for fitting. However, one must be cautious of conclusions drawn given the number of lines used in this comparison.

SLIDE shows no obvious relationship between the two with no obvious trend, although the absolute error is larger for SLIDE – but as the errors are around one thousandth of a wavenumber for SLIDE these errors are around 10% of the wavenumber gridding of the spectrum and this seems to corroborate that the sinc fit enables us to obtain an accuracy for centroid position of at least a quarter of the wavenumber gridding.

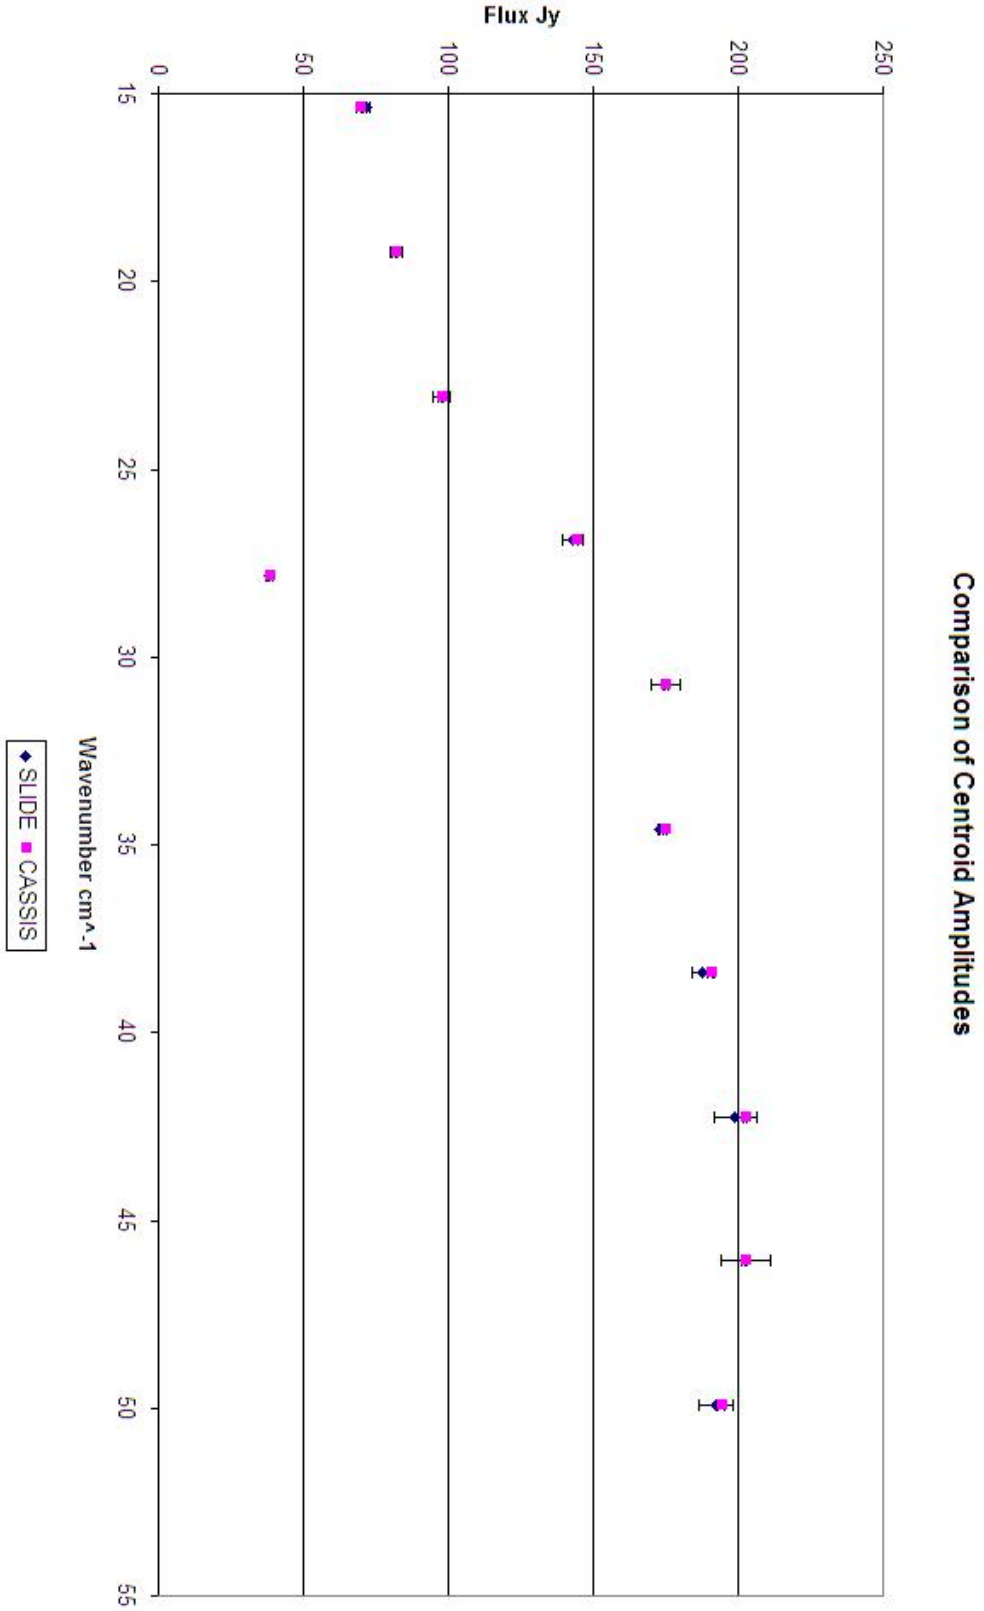

Figure 4.20: Comparison of reported line amplitudes between SLIDE and CASSIS.

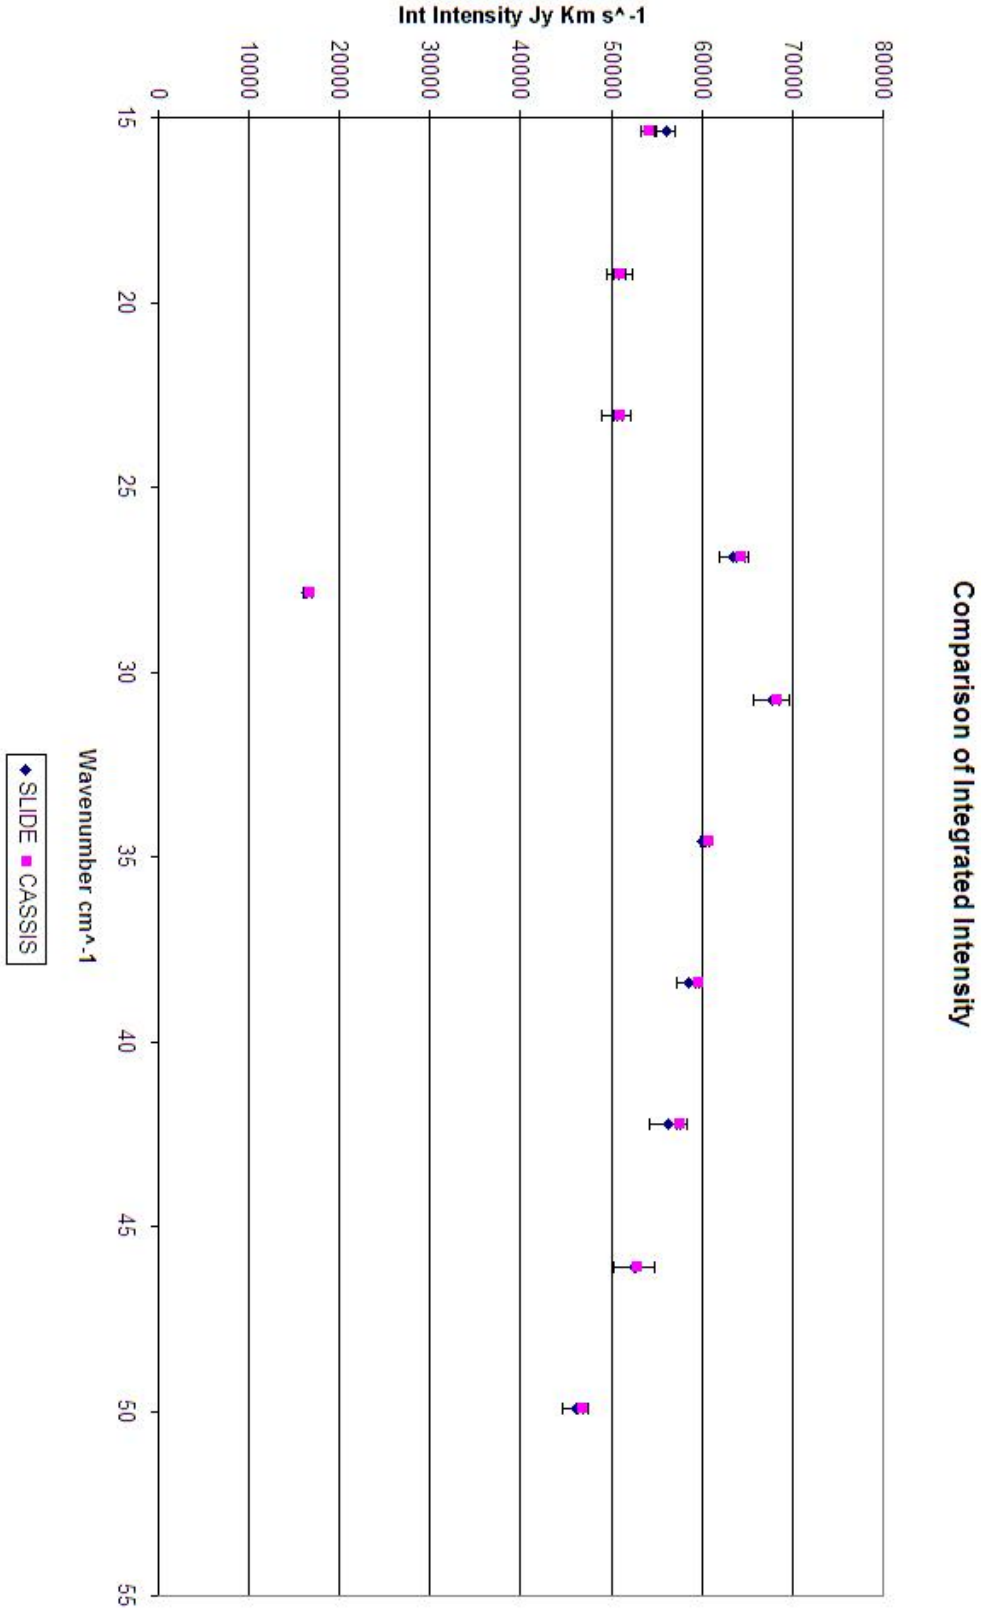

Figure 4.21: Comparison of reported line integrated intensities between SLIDE and CASSIS.

Figure 4.20 compares the line amplitudes reported by both SLIDE and CASSIS. SLIDE produces slightly larger errors in amplitude than CASSIS, but the difference is minimal and on the whole there is excellent agreement between the two. Error bars are small, and in the worst case only amount to a few percent. From this Figure, it would seem that CASSIS reports slightly higher line amplitudes than SLIDE, although we believe that it is most likely due to continuum subtraction rather then a genuine difference between the two, since the difference increases with wavenumber. This would be expected if the continuum fit is not optimal, and with many SPIRE spectra the challenge has been to adequately subtract the underlying continuum to obtain a satisfactory baseline. An imperfect baseline subtraction would result in a deviation that increases with wavenumber due to a slope being introduced into the baseline.

Figure 4.21 compares the integrated intensities reported by SLIDE and CASSIS. SLIDE shows the larger error in amplitude, and so we would expect the same with the integrated intensity (the calculation is done by quadrature sum of the component errors). There is excellent agreement between the two, again with errors of a few percent. SLIDE has larger errors for the $33cm^{-1}$ line – but this line lies on a band boundary and it is very probable that this is due to imperfect baseline fitting. CASSIS also shows slightly larger values for integrated intensity per line, but again this is likely to be a continuum fit issue.

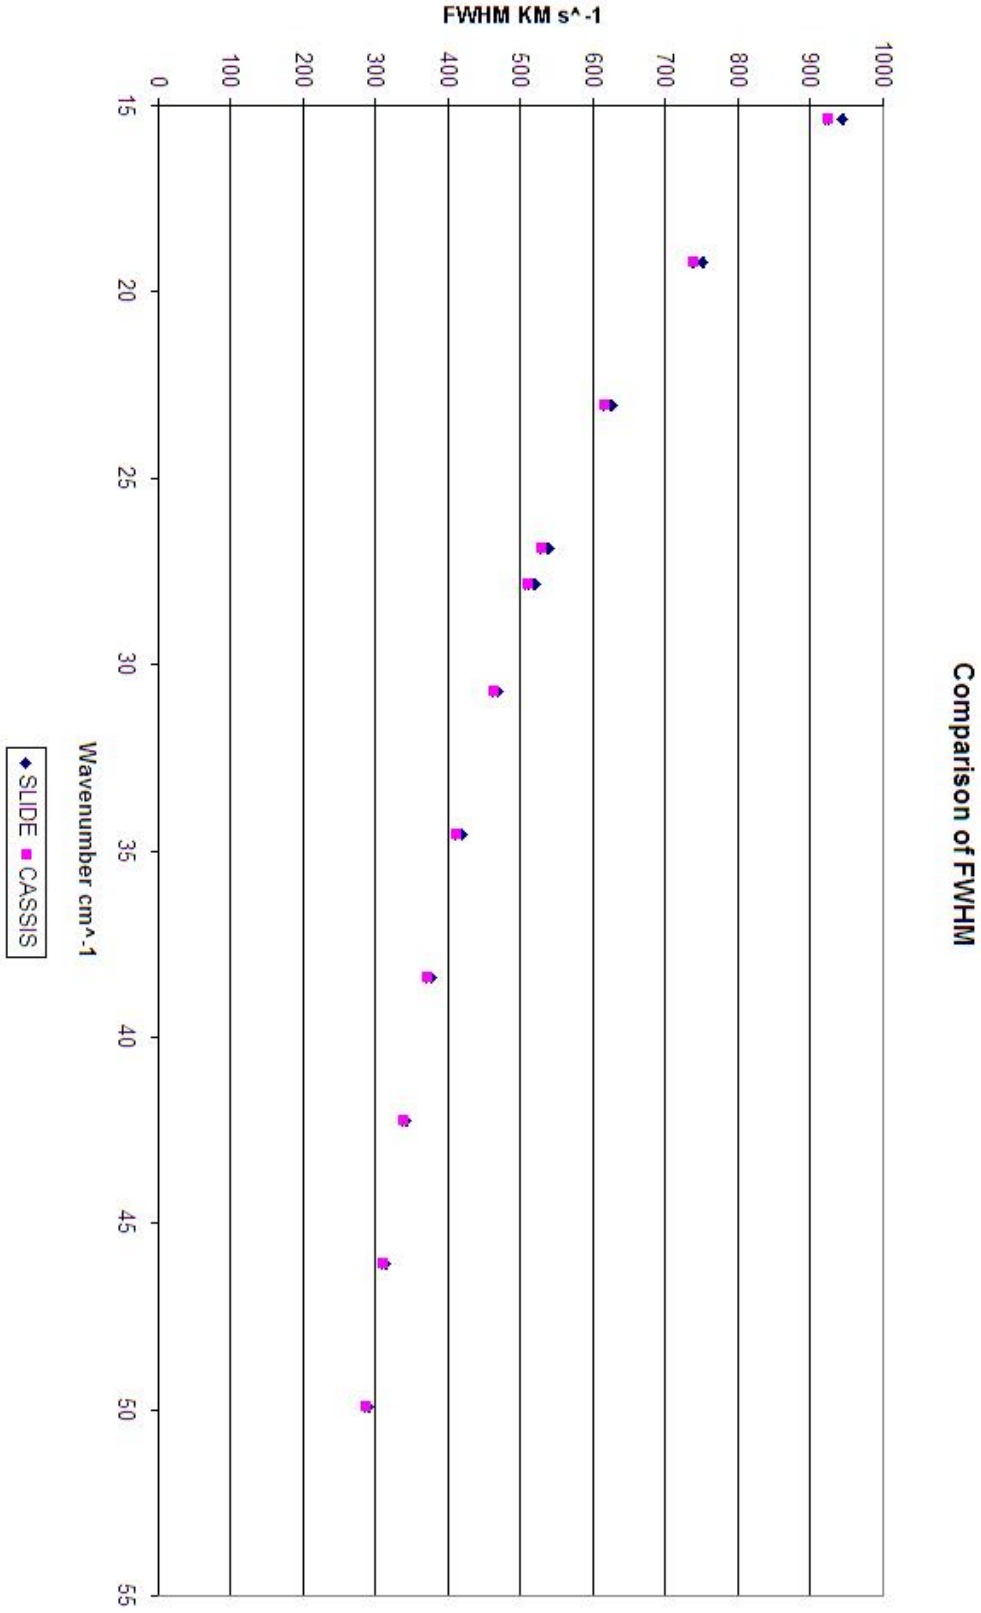

Figure 4.22: Comparison of reported line FWHM between SLIDE and CASSIS.

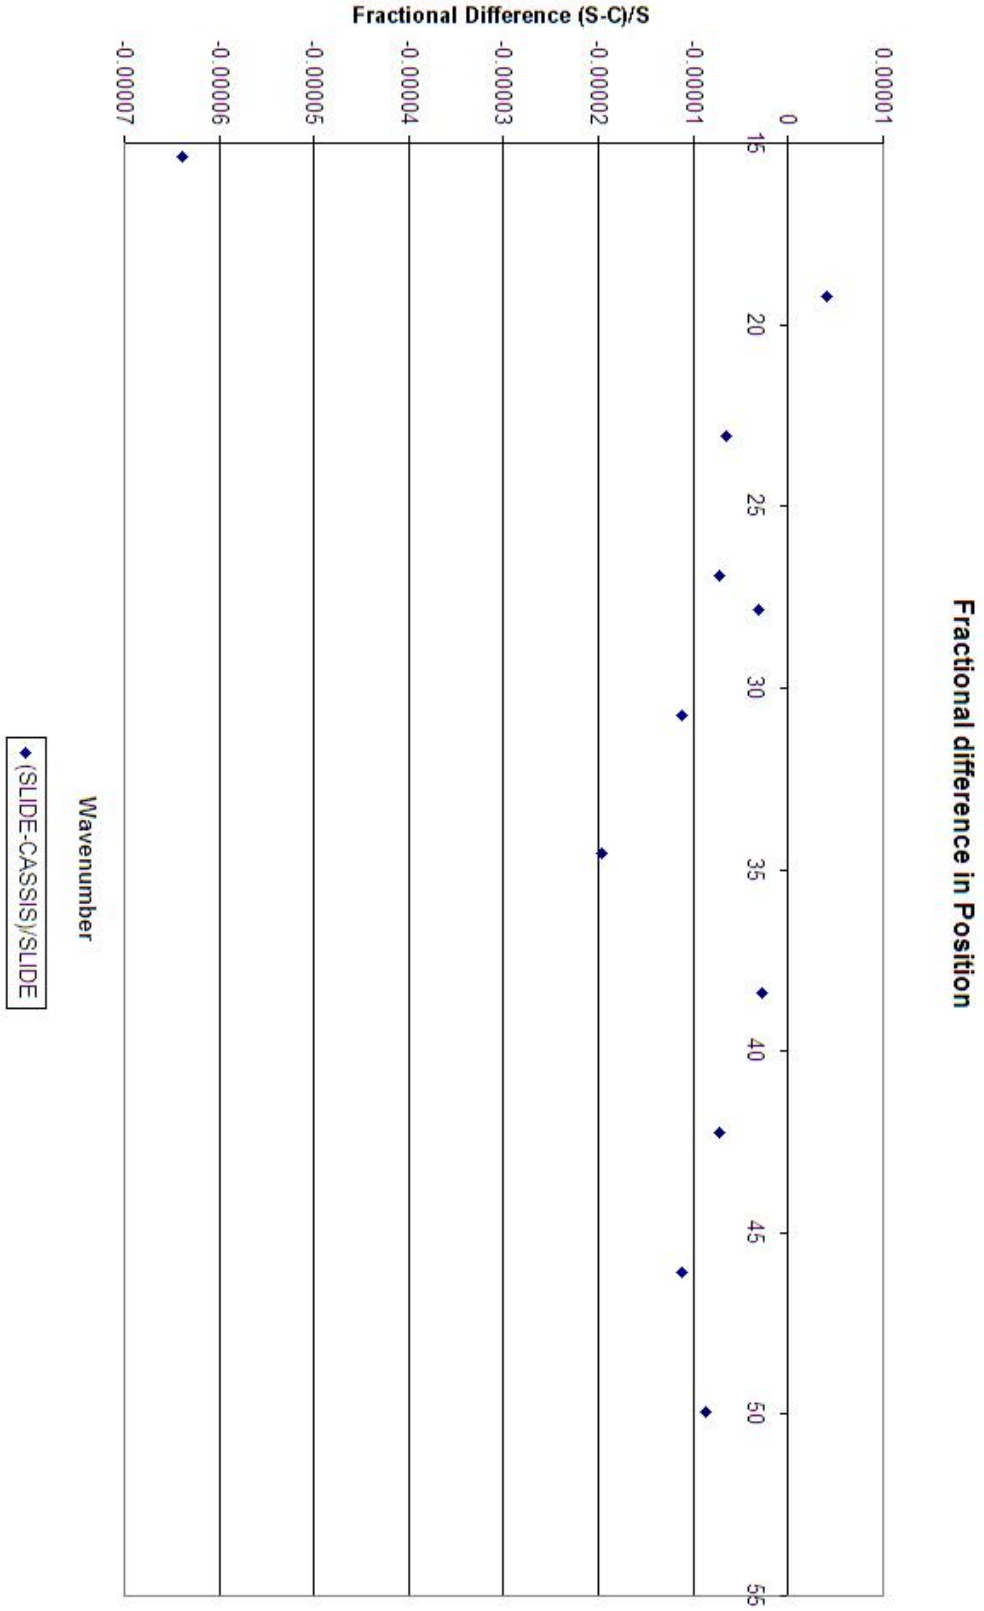

Figure 4.23: Comparison of fractional difference in centroid position between SLIDE and CASSIS, by wavenumber.

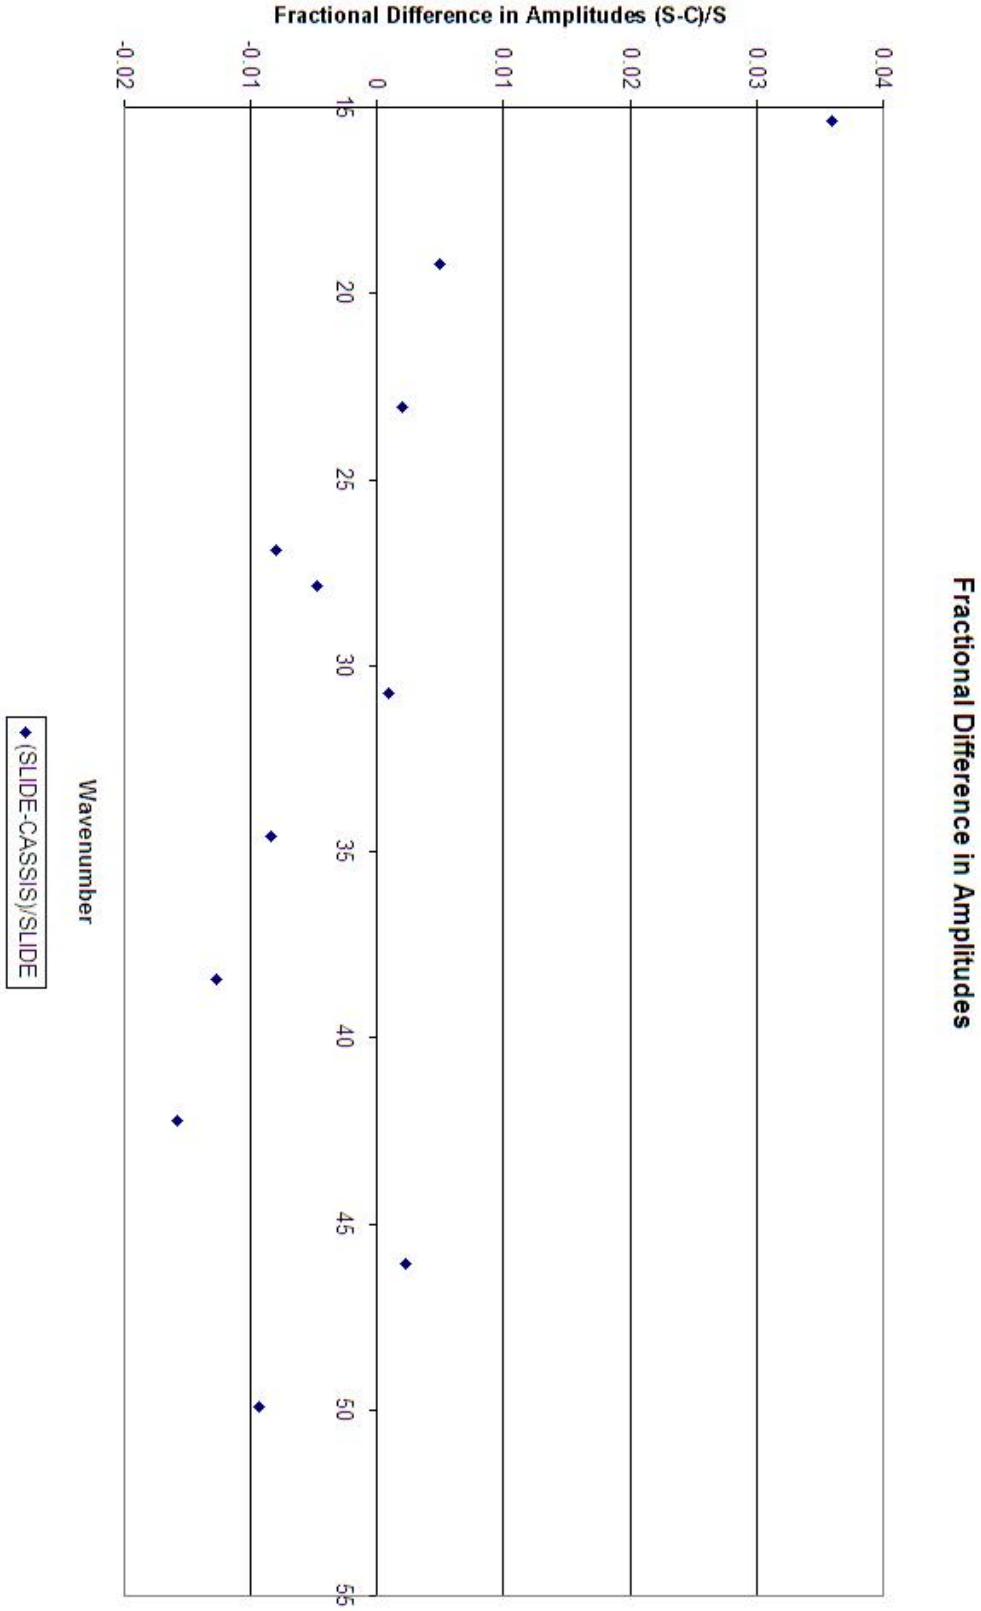

Figure 4.24: Comparison of fractional difference in line amplitude between SLIDE and CASSIS, by wavenumber.

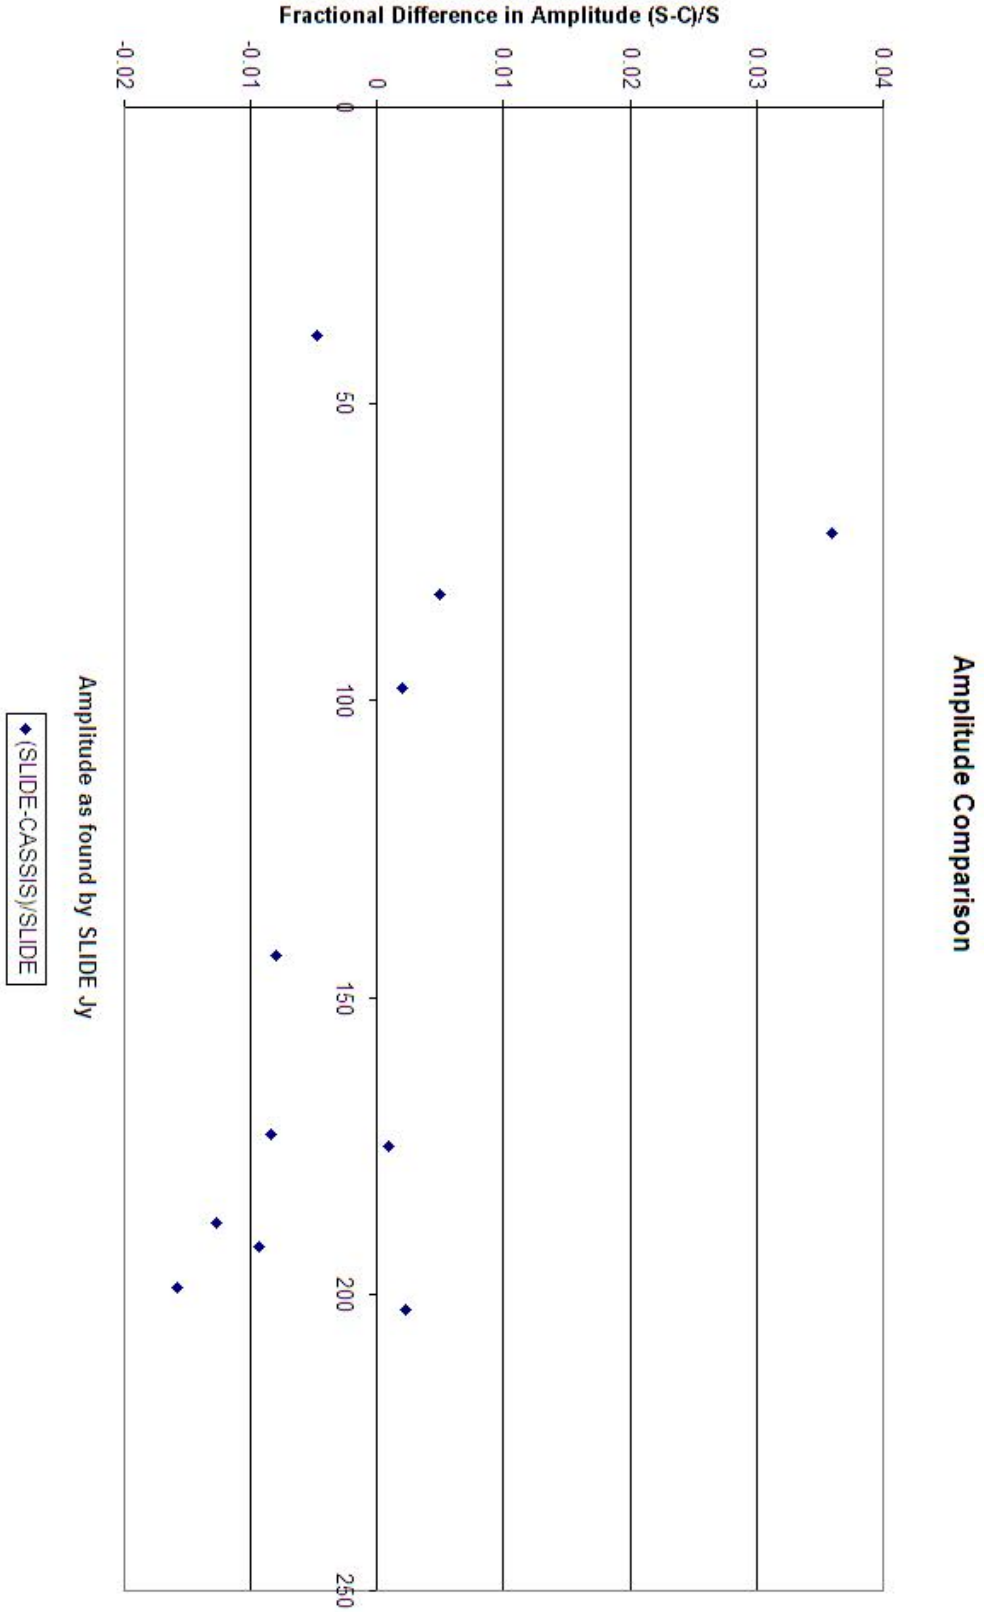

Figure 4.25: Comparison of fractional difference in line amplitude between SLIDE and CASSIS, by amplitude.

Figure 4.22 compares the FWHM of the line data returned and we see that CASSIS reports larger error bars than SLIDE. However the error values are small, with excellent agreement between the two. Overall CASSIS reports slightly lower FWHM by a few percent. Comparing the results for SLIDE and CASSIS for FWHM, intensities and amplitudes we see definite differences in the shape of the resulting fit. It is likely that this is largely due to continuum subtraction differences between the two but one must take into account the differing approach to line width used in the fit – SLIDE fits were performed with a floating maximum path difference L which defines the shape of the ILS, whereas the CASSIS fits were performed with a fixed L.

Over the SPIRE wavenumber band, there is a small fraction of a percent deviation in centroid position obtained from the fit routines. CASSIS is producing a consistently larger value for the centroid position. This is shown by Figure 4.23. One very large error is not shown. This is at the low end of the SLW band where ripple structure dominates, hence it is discarded.

Figure 4.24 shows a comparison between the fractional difference in line amplitude returned by CASSIS and SLIDE by wavenumber. We see that there is a increase in absolute amplitude deviation with wavenumber. This uses only eleven lines, hence with so few lines in the comparison one must be wary of drawing conclusions. Again, it is likely that an imperfect baseline fit has introduced a slope. The $15cm^{-1}$ line lies in an area where there is a lot of underlying structure, and so yields a larger difference. The overall deviation is around 1% which is to be expected. Differing continuum fits would yield some deviation in returned amplitude and this small deviation is therefore unlikely to be significant. Compared to Figure 4.23, there is no obvious distinction between the SLW and SSW bands. Overall, SLIDE and CASSIS agree well.

Figure 4.25 shows the fractional difference in line amplitude by amplitude of the centroid, as opposed to the absolute centroid position of the previous graph. Comparing the fractional difference in amplitude, sorted by amplitude returned, we see that it is difficult to see a relationship between amplitude and fractional difference in amplitude for SLIDE and CASSIS. The scatter on the graph (noting yet again the problematic $15cm^{-1}$ line) is minimal (around one percent) and is likely a result of differences in continuum fitting.

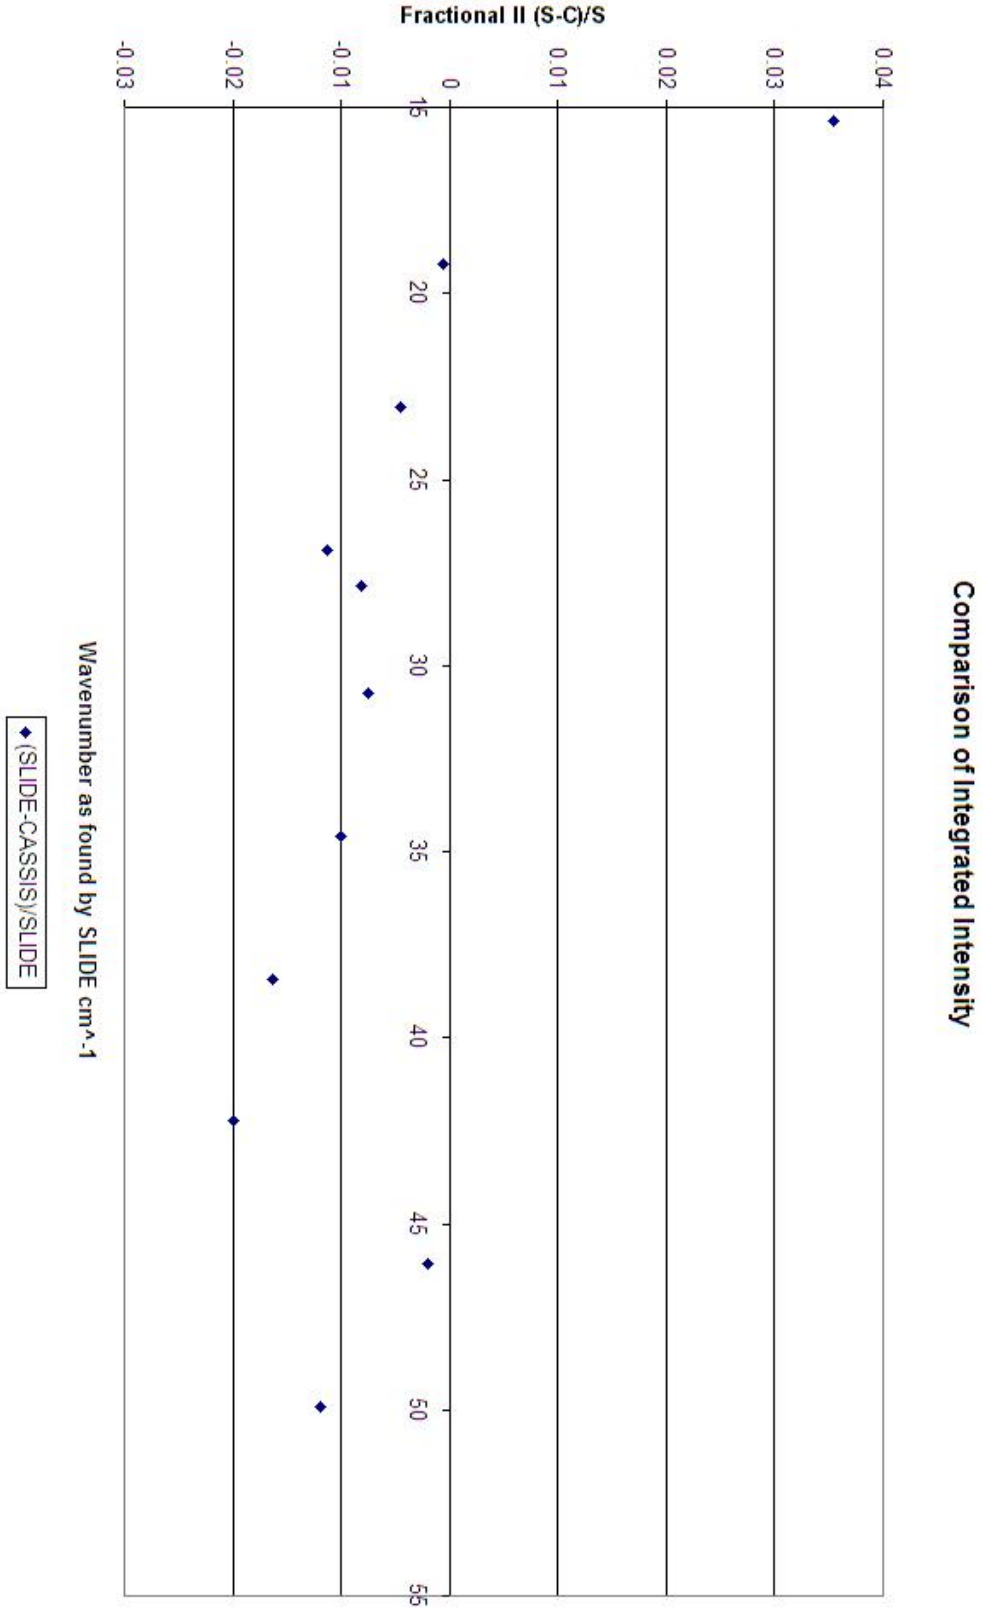

Figure 4.26: Comparison of fractional difference in integrated intensity between SLIDE and CASSIS, by wavenumber.

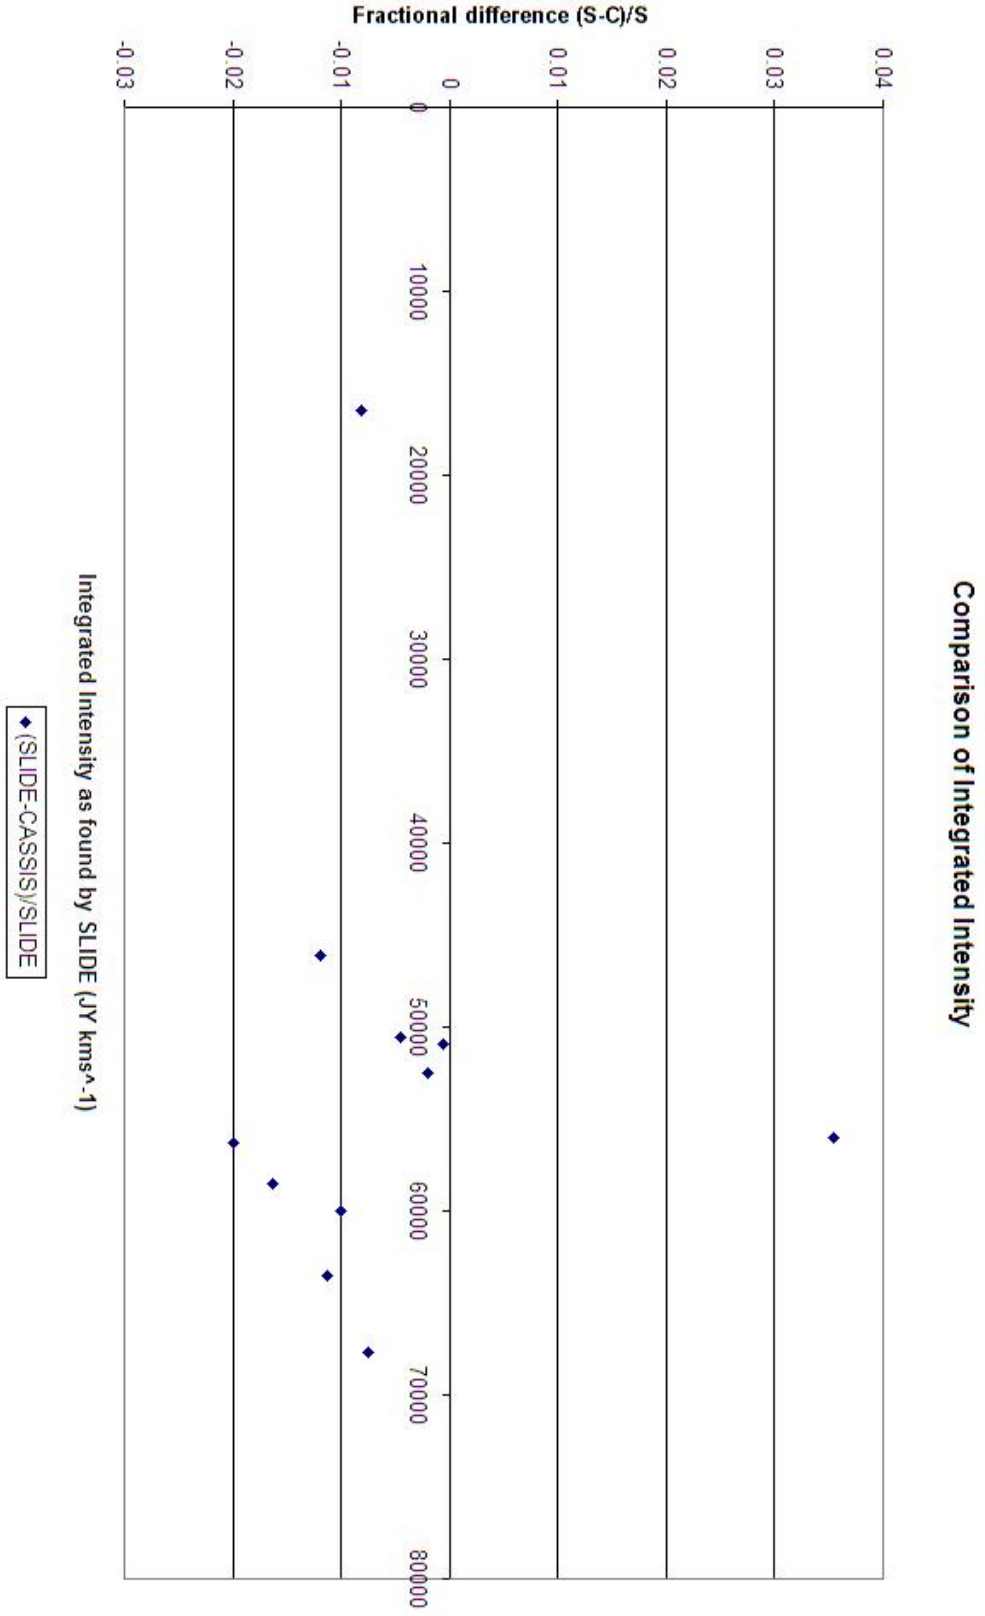

Figure 4.27: Comparison of fractional difference in integrated intensity between SLIDE and CASSIS, by integrated intensity.

Figure 4.26 shows the fractional difference in integrated intensity between SLIDE and CASSIS. Again, we note that integrated intensity is calculated for CASSIS from the line width and FWHM and the errors calculated by quadrature sum by hand. SLIDE does this calculation automatically. SLIDE uses the IDL INT_TABULATED function to calculate the integrated intensity by numerical integration. At most, there is a 2% deviation – again, ignoring the problematic $15cm^{-1}$ line. Although there is no obvious hint of a separate relationship for SLW and SSW in integrated intensity or fractional difference, there is a slight trend for the two programs to deviate as wavenumber increases. Figure 4.26 shows that CASSIS is reporting slightly larger integrated intensities as wavenumber increases. This may be due to the way L is handled between the two codes used, or again might be an artefact of imperfect baseline fitting.

Figure 4.27 compares the fractional difference in integrated intensity against integrated intensity. Once again we see that it is difficult to show a relationship between the two. Apart from the problematic line at $15cm^{-1}$, the deviation is small. This again tends to support the suggestion that the deviation between the two results is largely due to imperfect baseline fitting.

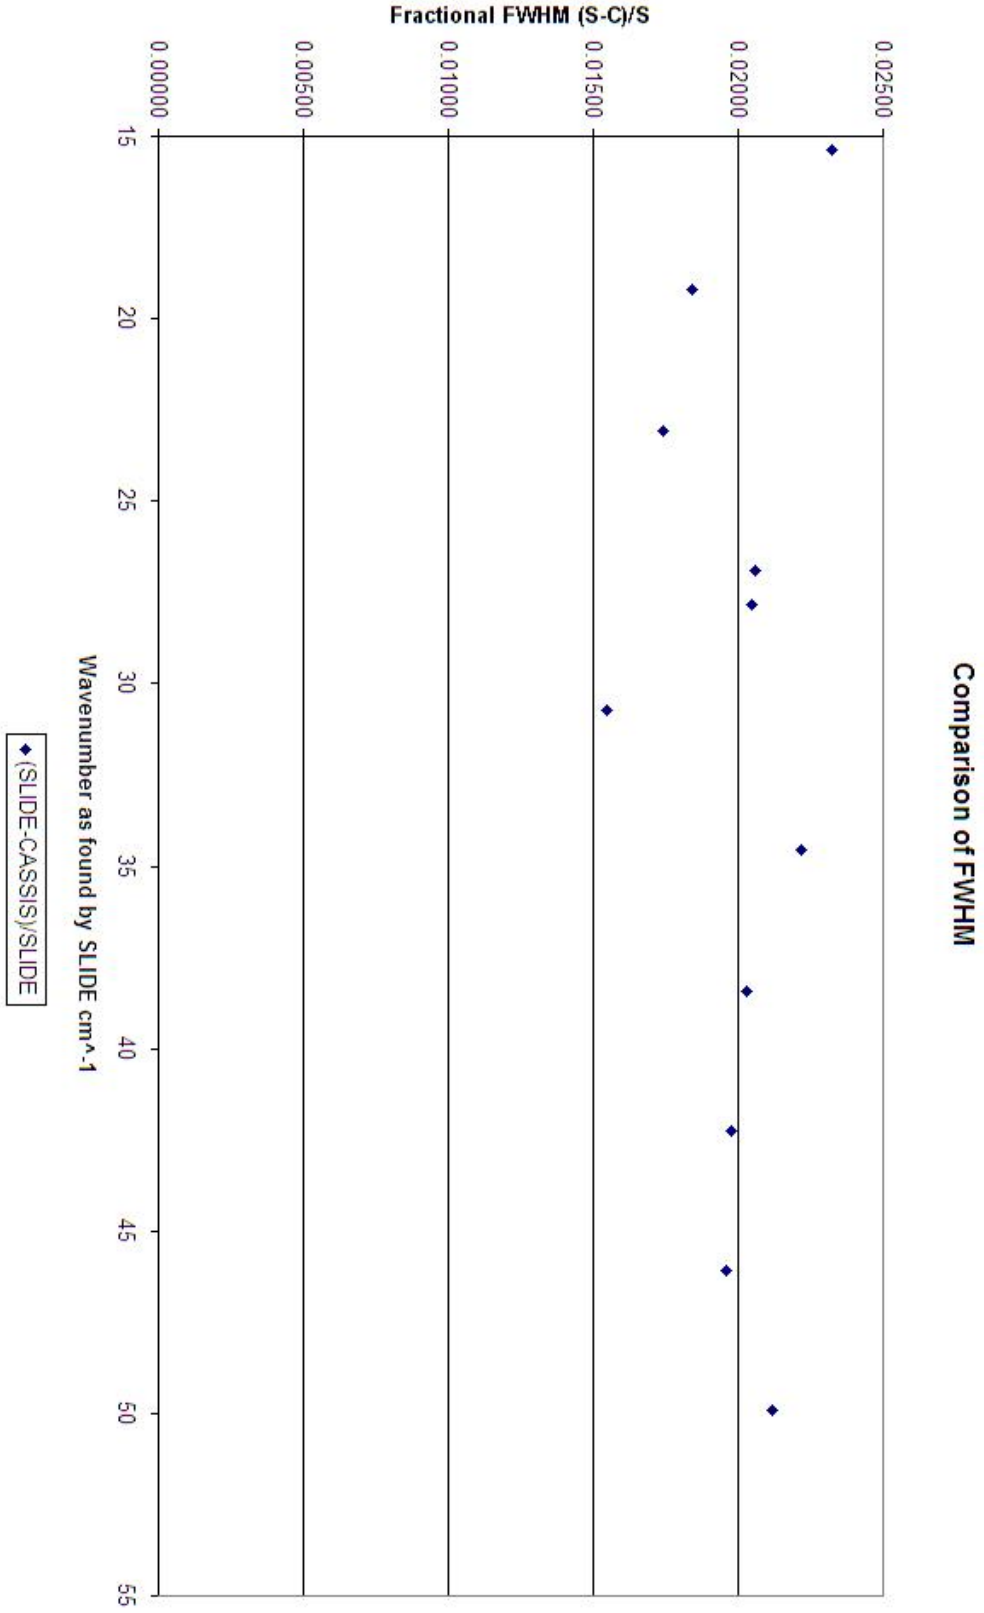

Figure 4.28: Comparison of fractional difference in FWHM between SLIDE and CASSIS, by wavenumber.

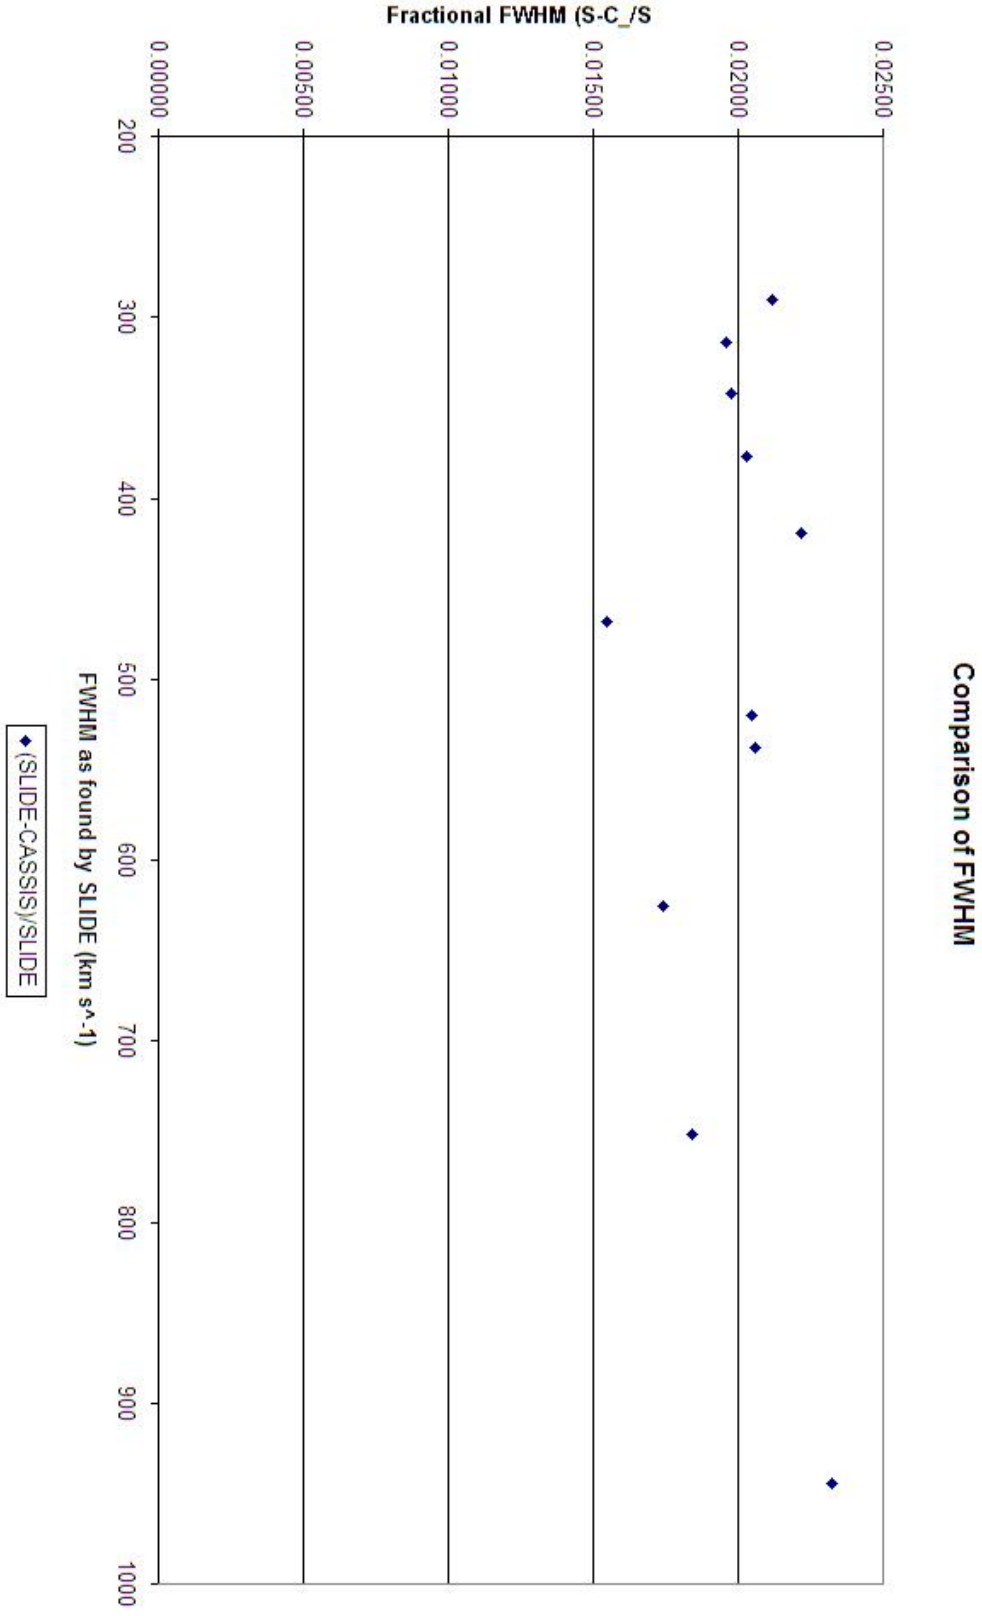

Figure 4.29: Comparison of fractional difference in FWHM between SLIDE and CASSIS, by FWHM.

In this comparison CASSIS used a fixed L and SLIDE used a floating L to generate these fits, and so any comparison between the two routines must take this fact into account. As shown by Figure 4.28, it seems SLIDE returns slightly larger FWHM than CASSIS, of the order of $1-2\%$. There is a hint of a slight difference between the SLW and SSW bands, especially if one ignores the two lines at $27cm^{-1}$ where a CI line lies very close to a ^{12}CO line. In that case, the relationship becomes clear. However, yet again one must stress the dangers of applying this relationship to so few lines. Figure 4.29 compares FWHM to the fractional difference in FWHM, grouped by line FWHM. We see no evidence of a trend in either Figure 4.28 or 4.29 for the relative differences in returned FWHM for wavenumber or FHWM.

4.4.2 Summary of SLIDE and CASSIS comparison

The results of CASSIS and SLIDE show good agreement with each other. Differences of $1-2\%$ can almost certainly be accounted for by differences in the continuum fit. The only area where there is a real difference is in the centroid position error, where CASSIS seems to show a power-law relationship between centroid position and its associated error. It cannot be stressed enough that the number of lines is too few to ascribe a physical interpretation to this.

In this comparison SLIDE used a floating line width, and CASSIS a fixed width. However, for this source this does not pose too much of an issue and the agreement between codes for the source is not compromised by this difference in approach. Of particular interest is the obvious disparity in results obtained across SLW and SSW, where the line fitting codes are revealing some differences between the two spectral bands. SLW is prone to fringing and structure below $\approx 20cm^{-1}$.

So we are possibly seeing how this structure is affecting the line fitting routines of both codes in different ways.

Although SLIDE is less mature then CASSIS, line recovery parameters between the two codes show a remarkable agreement. Even with such a small scale comparison SLIDE has performed well. Testing against other packages has enabled us to refine SLIDE and narrow significantly any areas of difficulty in line retreival.

4.5 SLIDE Output

SLIDE takes as its input a SPIRE-reduced SED, and outputs a variety of different files labelled with an identifying suffix. Each suffix also has an identifier to label each file as being from an SLW or SSW band SED, as well as a line type identifier for SCG, Gaussian, and Sinc line profile fitting. There is also a CMP line type identifier for the Composite fit function fitting, where both Sinc and SCG are fitted simultaneously.

The file types output by SLIDE and their contents are:

INFO contain a list of retrieved line parameters (with errors) and a parameterisation of the initial continuum fit.

SEDD contain a text file SED, with columns for wavenumber, flux and error (if available).

CSBS are the continuum subtracted spectra used for line fitting.

RESD are a three column text file of similar format to the _SEDD_ file but contain the the result of subtracting the fit from the continuum-subtracted

spectrum.

**SFIT** are the generated spectra from the recovered line parameters, or synthetic line spectra without noise.

**BASE** are the baselines subtracted from the spectra to form the CSBS minus any extra polynomial subtractions.

**OVER** are graphical overplots of the fitted lines onto the continuum-subtracted baselines for visual analysis.

**XCSV** are a Comma-Seperated-Value form of the INFO files suitable for import into a spreadsheet or other computer package.

**LINE** are truncated SEDs containing the summed-fit and the continuum-subtracted baseline data for a narrow range around each line centroid position.

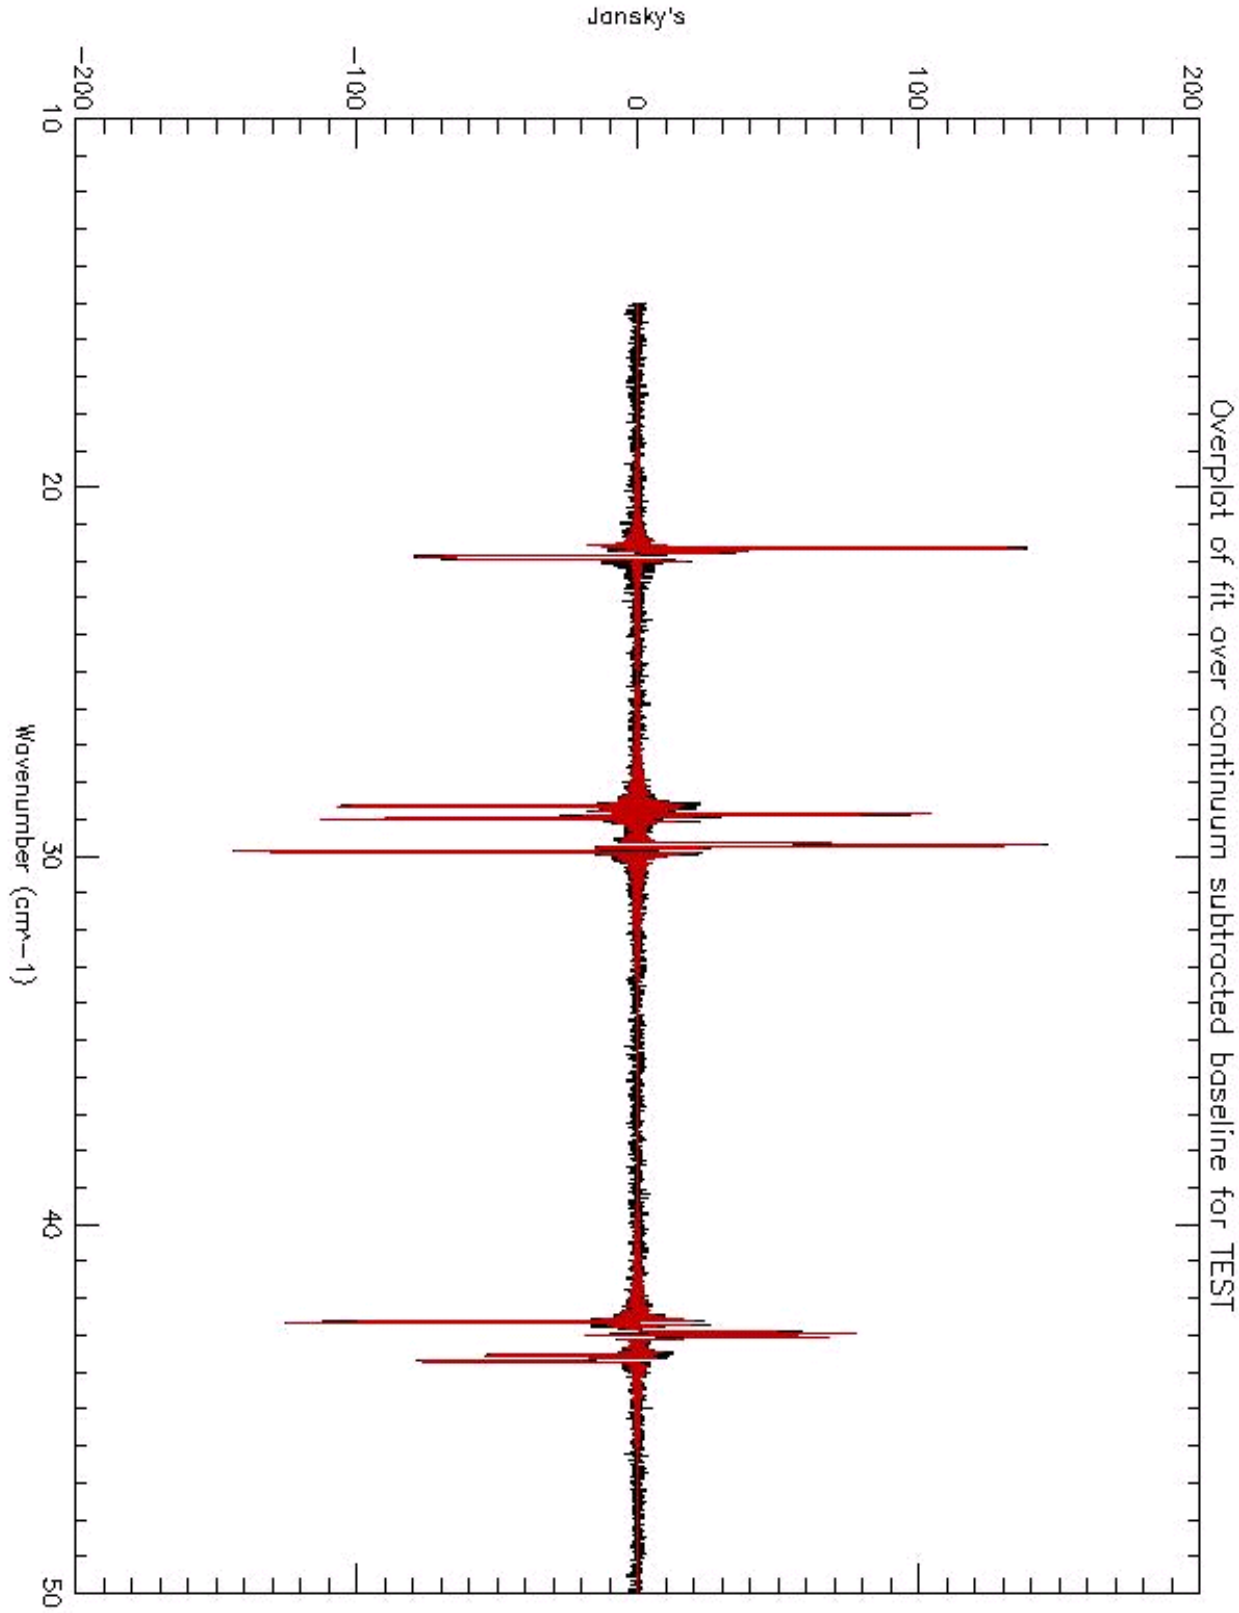

Figure 4.30: An example of an _OVER_ file produced by SLIDE.

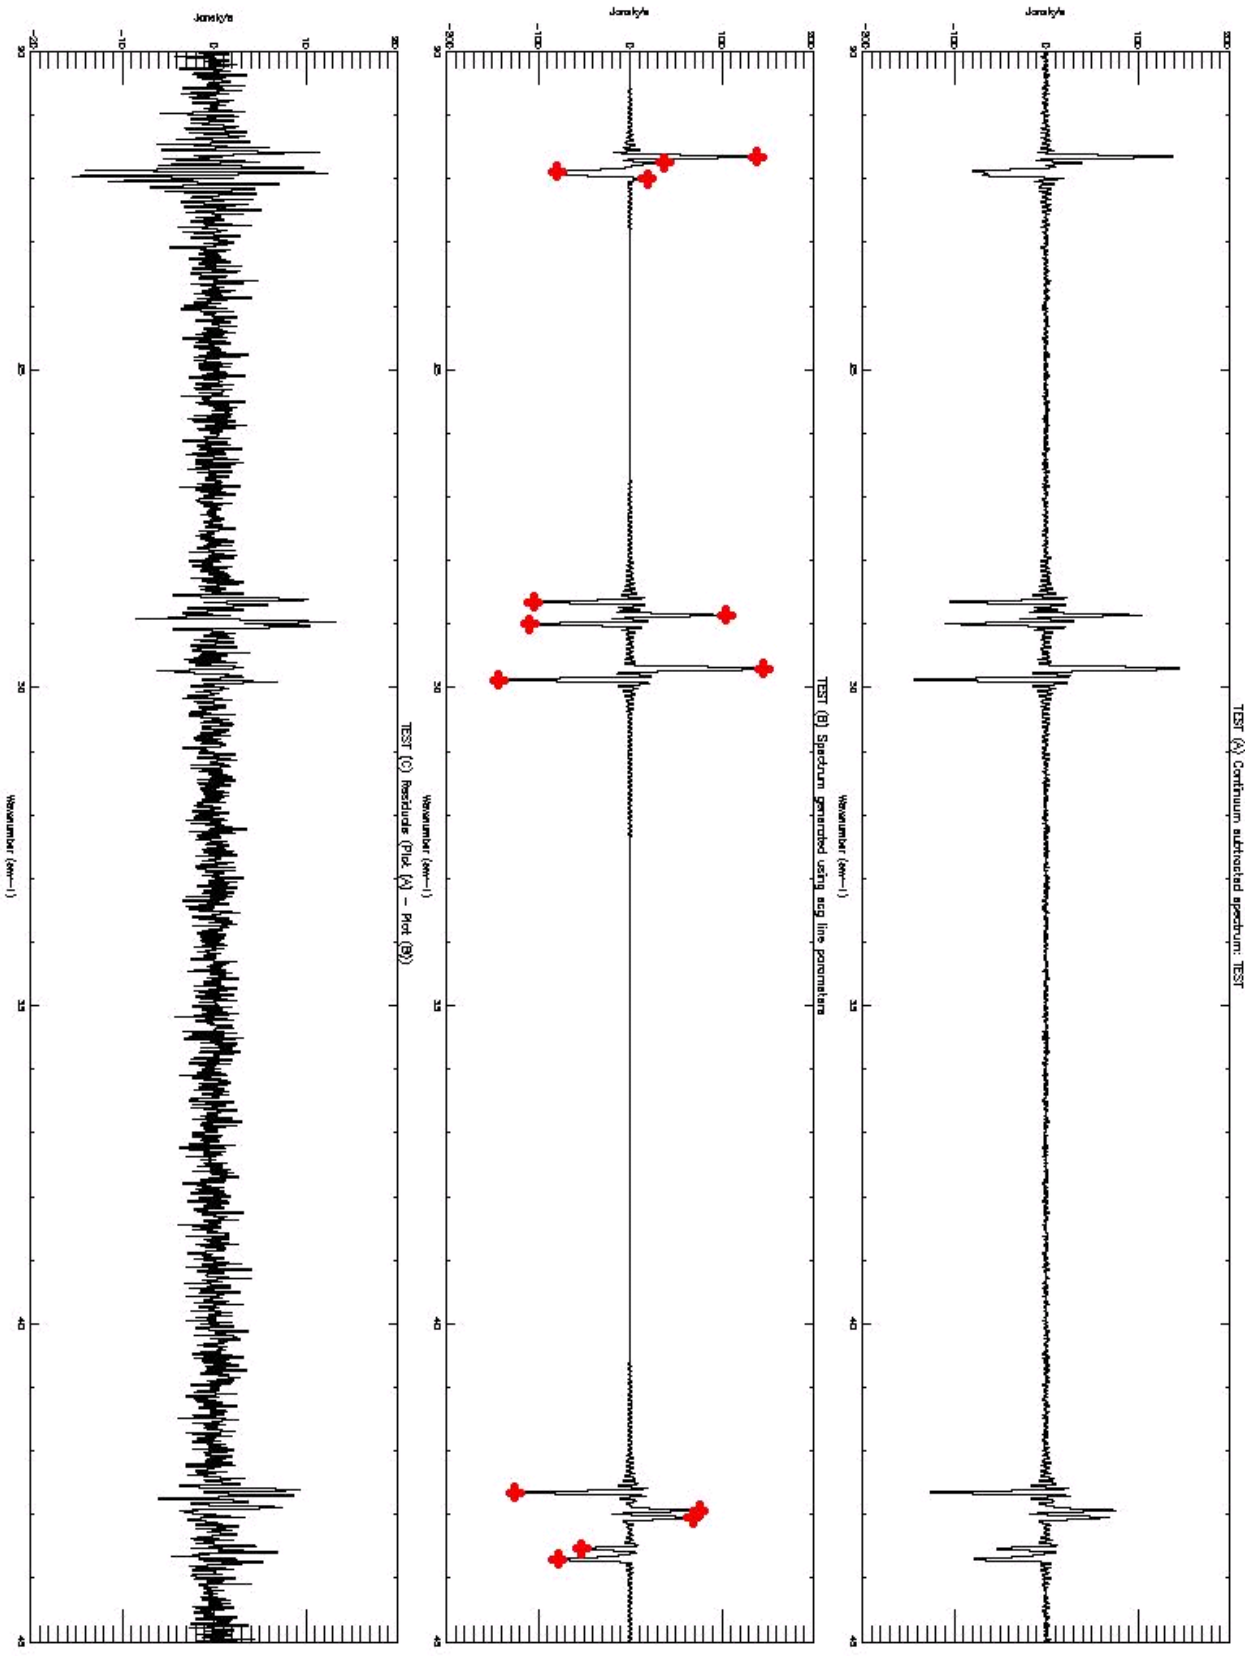

Figure 4.31: A SLIDE 3 way plot.

Figure 4.30 shows a sample *_OVER_* file as produced by SLIDE for a test, artificial SED input file. The black line is the original SED, with its continuum subtracted. Overplotted in red is the SCG fit as performed by SLIDE in this case. Figure 4.31 shows another plot available from SLIDE, which is user selectable during the fitting process. After line fitting, the user can select to output this three-way plot which shows, from top to bottom, the continuum subtracted baseline, the synthetic line SED generated from the fitted line parameters and the residuals from the continuum subtracted baseline after fit. Note the middle plot uses large circles overlaid on the plot to aid visual identification of the line positions found by SLIDE.

4.6 SLIDE: The Future

The future plan for SLIDE is for the code to be merged with a similar tool produced by a team at the University of Lethbridge, Canada – (Naylor, 2010; Jones et al., 2009). This merged code is to be converted into a module for HIPE written in Python. In comparative tests between the two codes, the differences were of a similar magnitude to that between SLIDE and CASSIS. On the same data, centroid positions deviate by around one twentieth of a wavenumber and fluxes by around $10 - 20\%$. Baseline subtraction differences are the largest source of difference between the results of each tool. As far as the user is concerned, the major difference between SLIDE and the UL tool is that SLIDE is meant to be fully interactive, whereas the Lethbridge code is more batch orientated and is configured by text-based configuration files.

At the present, SLIDE has an offshoot version meant to deal with PACS data.

This works in wavelength rather than wavenumber and does fixed- or floating-width FWHM Gaussian fits. This version is being developed and tested with help from teams analysing PACS data, and comparisons between SLIDE_PACS and ELF/DIPSO on PACS data yield consistent results in line centroids to within a hundreth of a micron and flux consistency to around 10%, although this offshoot is still being tested and worked upon.

4.7 Summary

SLIDE was developed as an IDL tool to quickly and accurately identify line parameters in SPIRE SED's, although it has progressed far beyond the original line identification tool. As such, it is used in conjunction with the processing pipeline as opposed to being any replacement for it.

A number of useful continuum fit functions were identified and tested, although SLIDE primarily uses continuum-fitting as a means to improve line-fitting rather than to perform detailed continuum analysis. In examination of SED's produced by SPIRE and through continuum fitting, fringing proved to be problematic. The accurate removal of fringing to leave a smooth baseline is important to accurately identify line parameters. Issues with fringing of the SED as a result of artefacts of processing within the pipeline were identified and analysed. Tests were performed to identify the effect such fringing would have on line identification. As such 'zones of avoidance' were found where the identification of weak lines would be increasingly sensitive to continuum subtraction.

SLIDE was tested using noise-free and noise-added input SED's with various line profiles, and through this testing the continuum fit and line profile fit routines

were developed to be as robust as possible. Issues with SCG line fitting below a specific Gaussian width component were identified, and a new 'Composite' fit routine was developed to surmount the problem. SLIDE was originally developed to perform simplistic line identification on text-file SED's, but now SLIDE has a variety of continuum- and line-fitting routines to process SPIRE data and can accept text- or FITS-file input. It outputs graphical and textual information of the input SED.

SLIDE's initial design tests were complemented by using the University of Lethbridge's line-fitting tool for comparison,(Jones et al., 2009; Naylor, 2010). Over the course of SLIDE's development, sharing of ideas and approaches with Lethbridge were achieved and both tools share many ideas. As a result, these are the two tools that are to be merged into the HIPE tool. All that remains of the major differences between the two tools is the approach. SLIDE is designed to be interactive, whereas the UL tool is configured by text files and can be batch run. There is also a slight difference in approach to line subtraction between the two tools. Currently users are free to use both, compare results and choose the relevant tool as needed.

Although the intention is for SLIDE to be incorporated into a HIPE tool as part of the SPIRE pipeline processing, it is foreseen that SLIDE will still be used and developed as a tool for data processing. A fork of SLIDE called SLIDE_PACS is available for use with PACS data. SLIDE has been released for use for anyone involved in SPIRE data processing, and has been developed to be as user-friendly as possible. Over a number of years, SLIDE has become feature-rich for line identification purposes. It has been used on a number of papers to generate useful line parameters, see Section A.10. As the code has been released, it is possible for users to customise SLIDE for particular applications and to develop further.

Chapter 5

Line Contamination effects in SPIRE spectroscopy

5.1 Introduction

The SPIRE instrument on Herschel consists of a multi-pixel FTS and a three-band photometer. A review of the instrument by Griffin et al. (2006) notes that the SPIRE instruments have a field of view for the photometer of 4 x 8 arc minutes, the largest possible given the layout of the instrument. Its observational bands are centred on 250, 360 and 520 μm for PSW, PMW and PLW respectively. The FTS has a range of $\approx 200 - 650\mu m$ which is split into two bands - SLW and SSW. SSW covers $\approx 200 - 320\mu m$ and SLW covers $\approx 300 - 650\mu m$ simultaneously. The field of view is $2.6'$ for SLW and $2'$ for SSW (unvignetted), and the central pixel for the photometer and spectrometer is co-aligned, see Figure 5.1 .

Photometry observations with SPIRE yield three datapoints, and we can use

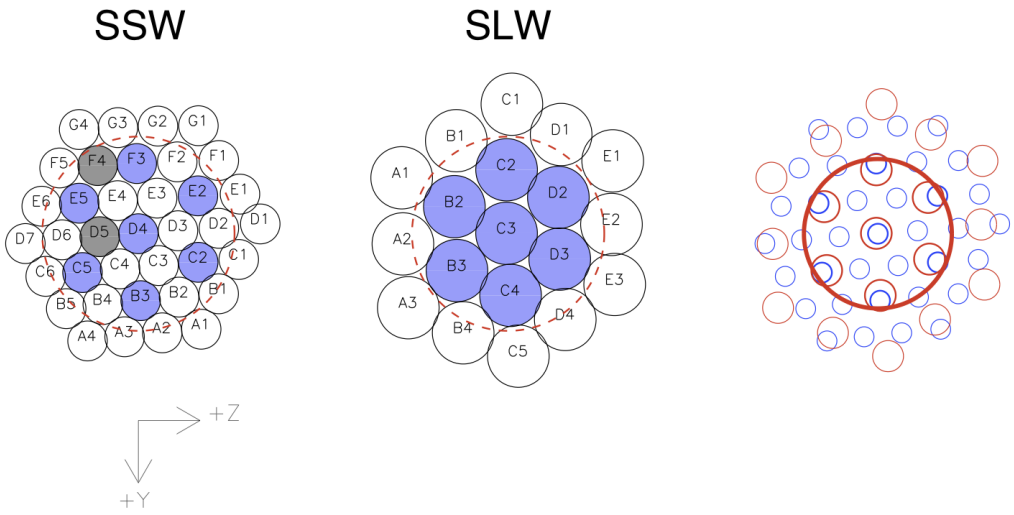

Figure 5.1: SLW,SSW Detector arrays from the SPIRE Observers Manual, European Space Agency (2010).

SCUBA (at $\approx$ 850μm) and IRAM (at $\approx$ 1200μm) frequencies to expand this slightly at longer wavelengths. We can also use PACS observations, at 70μm and 160μm, for expanding to shorter wavelengths. These bands were chosen as they are available for a wide range of sources. Using these five datapoints, we can perform greybody fitting and obtain values for the dust emissivity β and the greybody temperature T for simulated SPIRE observations.

In this region of the infrared spectrum (where the peak of emission corresponds to around 15K or less), there are strong line features from N+ and CO. Photometry observations will sum up the contributions from the continuum and the line features. This presents a problem in using these observations to produce accurate grey body models of the continuum. Line emission will present contamination to our results, and it is therefore essential to be able to assess the scale by which line contamination will affect our fitting. This is important for modelling core masses for pre-stellar cores, as for a spherical, isothermal dust source at a distance d we note that, as per Section A.1

$$M(r < R) = \frac{(S_{850\mu m}(\theta)d^2)}{\kappa_{850}B_{850}(T)}, \tag{5.1}$$

and this relates the total mass $M(r < R)$ contained within a radius R at distance d to the flux density at $850\mu m$ $S_{850\mu m}(\theta)$ integrated over a projected circle of angular radius $\theta = R/d$. We note κ_{850} is the dust opacity per unit mass column density at $850\mu m$, and $B_{850}(T)$ is the blackbody formula at the same wavelength for the dust temperature T. Therefore, any modelling to obtain an accurate measure of the mass will be affected by line contamination. To begin, we can use existing SPIRE FTS observations and use these to obtain the approximate scale of line contamination we expect to see in photometry observations.

5.2 Line Contamination effects

To examine the scale of line contamination, we studied a variety of galactic and extragalactic sources. These are named in Table 5.1. SPIRE observations of these sources were used to examine the scale of line contamination we would expect to see in typical astronomical sources. These are sources observed during the Science Demonstration Phase of the Herschel Mission and each of the sources is part of a Guaranteed- or Open-Time Key Programme.

These comprise a variety of type of sources, both galactic and extragalactic. They represent wide and narrow line sources, as well as sources with a large number of lines. For extended and quasi-extended sources we used SPIRE data calibrated for extended sources, and this introduces a discontinuity between SLW and SSW fluxes due to differing beam sizes.

Figure 5.2 and Figure 5.3 show the SPIRE observations for IC342 and IRC10216 used. We show these two spectra as an example as they show the extremes of line contamination we see in this group of spectra. Figure 5.2 shows IC342, it is an extended-source calibrated observation and shows the discontinuity between SLW and SSW bands due to the different beam sizes. It also has the lowest degree of line contamination seen in this group, at under 5%. Figure 5.3 shows IRC+10216 which has the largest amount of line contamination in this sample at over 25%. It is a point-source calibrated observation and does not show this discontinuity. The discontinuity is important, as the discrepancy in beam size means we can only consider the PLW photometer band in the point-source calibrated models. This is to ensure that we do not introduce errors in the calculation for line contamination from attempting to correct for the difference in beam sizes. This means that for extended-source calibration observations, we consider SLW and SSW separately to study PSW and PLW. In the case of point-source calibrated observations, we combine them into a single observation using the SSW data for the overlap and study all three photometer bands with them.

Point Source	Extended
IC342	Arp 220
DR21	CRL 618
G29.96	IRC 10216
G32.80	Mrk 231
M82	NGC 7027
VY CMa	

Table 5.1: Sources used to study Line Contamination.

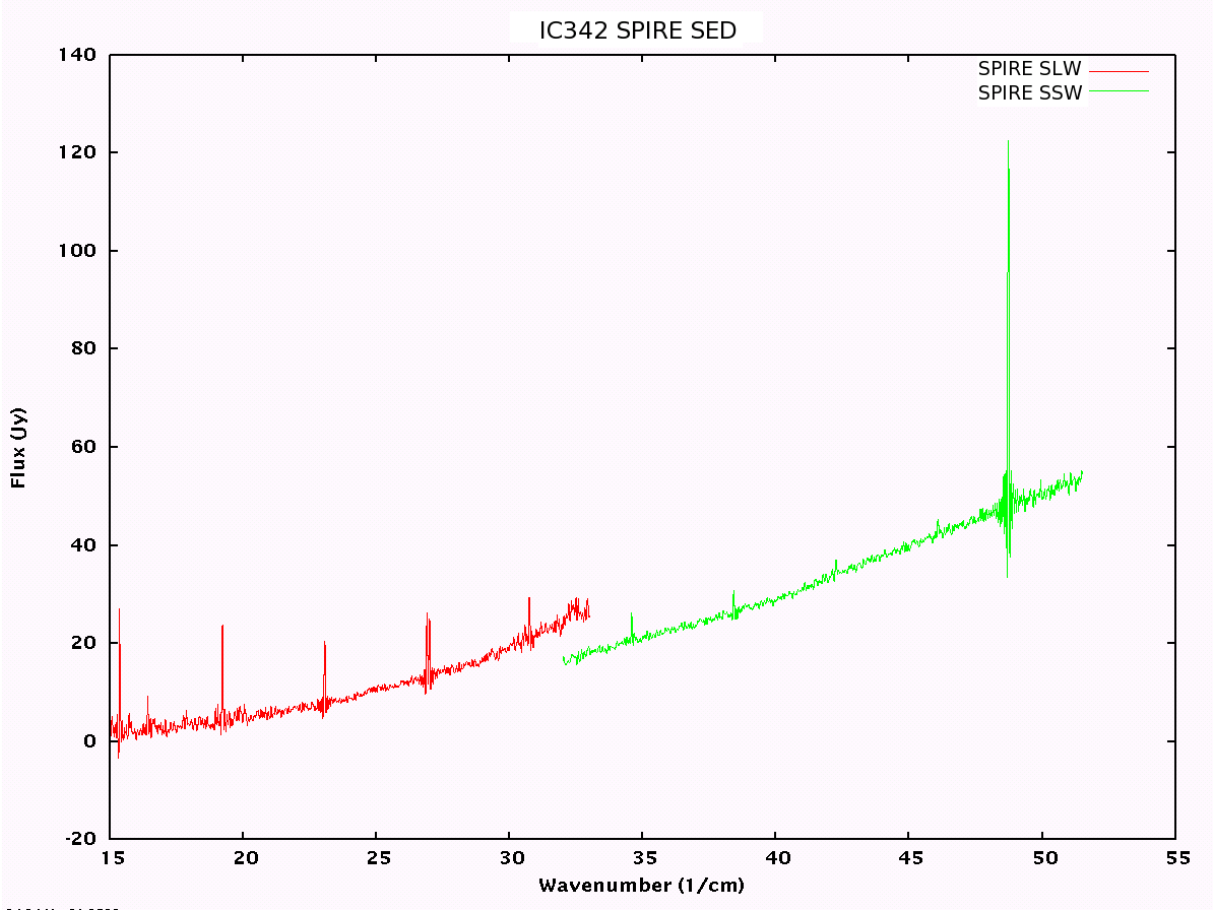

Figure 5.2: IC342 SPIRE observation used for line contamination study.

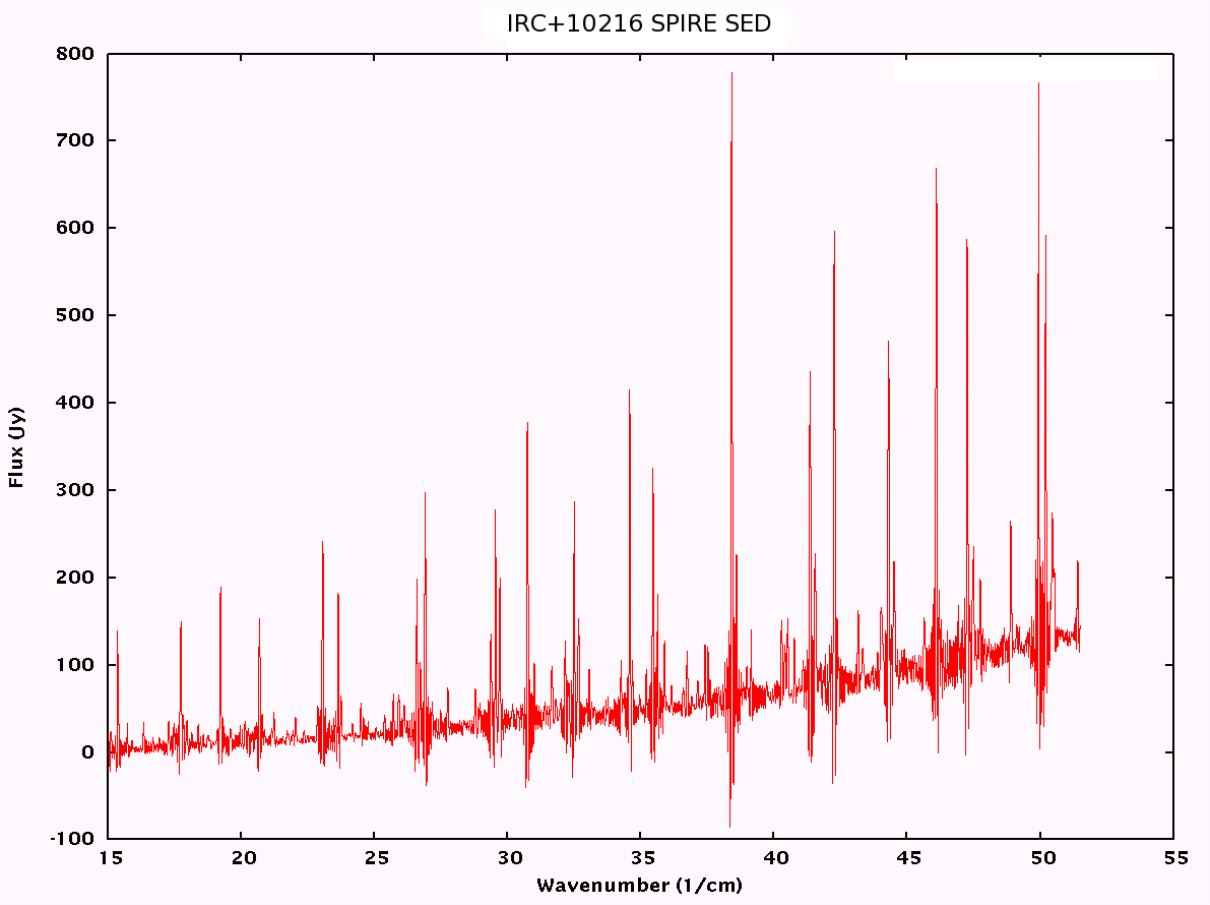

Figure 5.3: IRC+10216 SPIRE observation used for line contamination study.

5.3 Photometer Relative Spectral Response

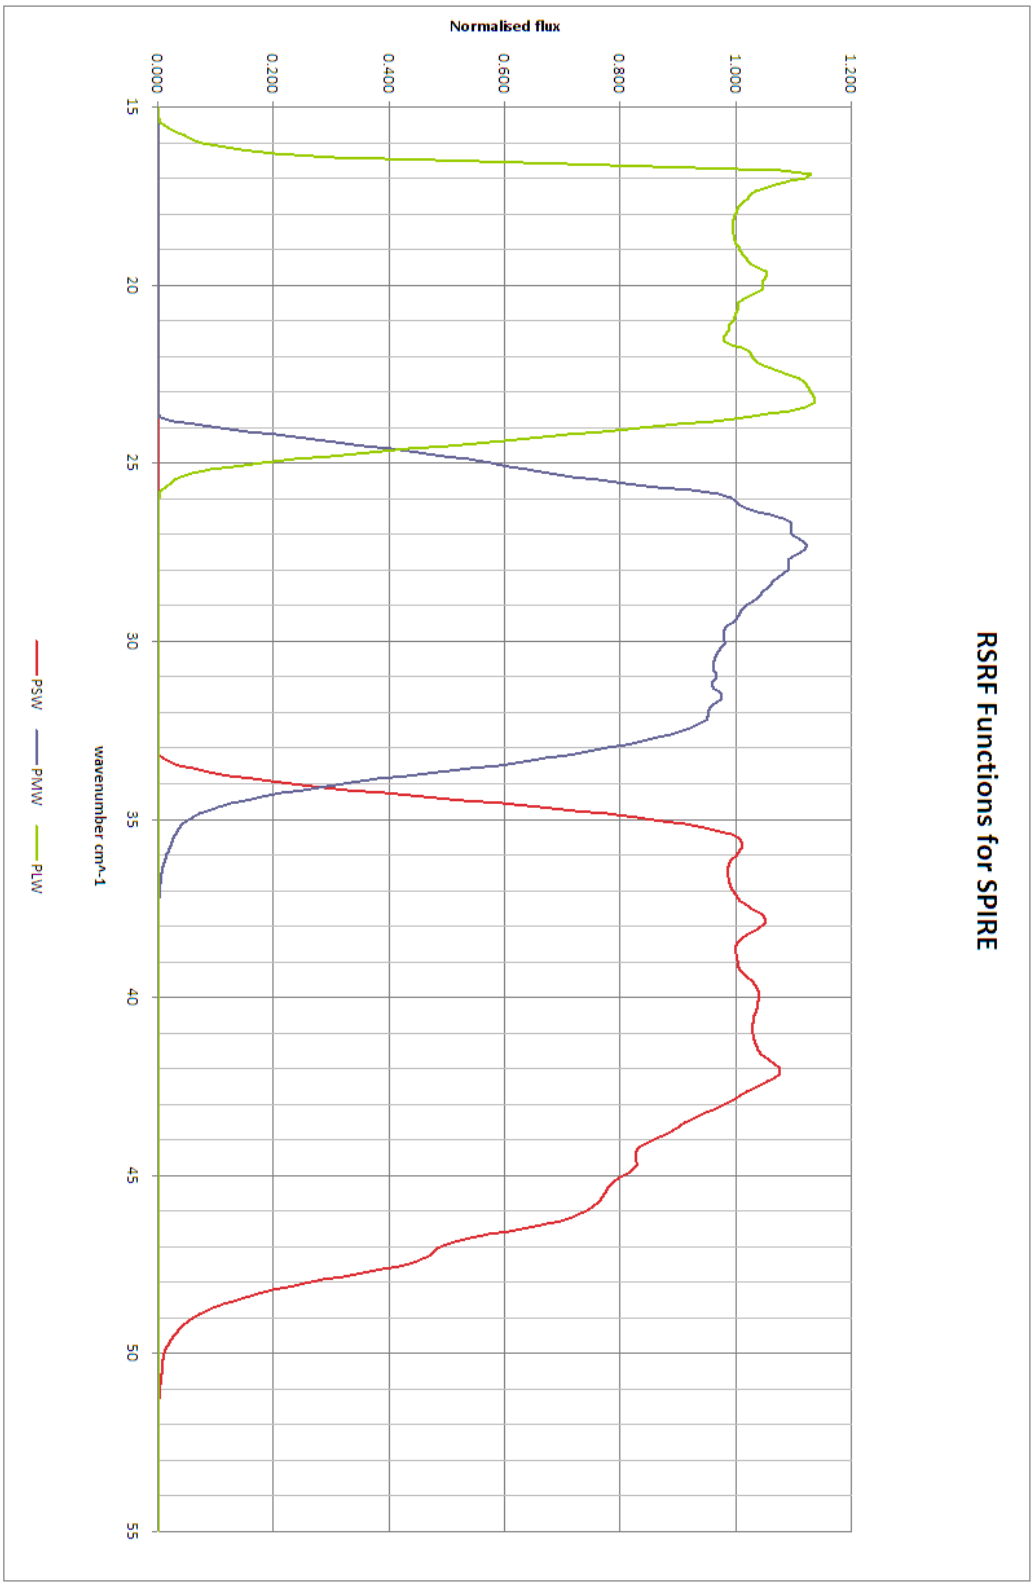

Figure 5.4: SPIRE RSRF.

In simulating photometry observations from spectrometer data, one needs to know the relative spectral response of the filter bands for the photometer. This function is known as the RSRF, the *Relative Spectral Response Function* of the instrument. For the SPIRE photometer, this data has been provided as a dataset of transmission across wavenumber and is shown in Figure 5.4 as normalised-area plots.

We began the investigation of the RSRF by modelling them as a set of polynomial functions of sixth-degree. The RSRF datasets are interpolated to a fine, regularly spaced wavenumber grid across the whole SPIRE band and a polynomial fit applied to the interpolated functions. Figures 5.5 to 5.7 show the RSRF functions with their corresponding polynomial fits overlaid. PLW and PMW were modelled by three polynomial functions, where we fit the rise, plateau and decay of the RSRF seperately. For PSW, we found that the plateau was best fitted by two seperate polynomials.

The choice was made to split each RSRF in order to easily model the functions as polynomials of order six or less. Comparison of the area under the polynomial curves and the original interpolated functions verified that the difference is marginal – less than a few percent. The coefficients for the polynomials used to fit the RSRF are given in Table A.1 We then selected five sources to test how well the polynomial forms of the RSRF compared to the interpolated, original forms of the RSRF using SLIDE, our custom line fitting and continuum recovery tool (covered in detail in Chapter 4). For these sources, we used point-source calibrated spectra and manually joined the two bands together. In the overlap region, we decided to use SSW data. These results are quoted in Table 5.2. Here we compare the continuum subtracted baseline (marked here as 'Baseline'), the continuum (fitted by SLIDE) and the original spectrum.

We noted there are negative values for the continuum subtracted baseline values for several sources. This is caused by the summing process magnifying any errors in the continuum fitting process. We corrected for this by creating artificial spectra, based on SLIDE output of the summed sinc fit (as a synthetic continuum subtracted baseline) and add this to SLIDE's fitted continuum. We shall explore this method in more detail later, but only note here it creates a synthetic, noise free approximation to the SED. We also note that the errors between the two methods are within a few percent, an acceptable limit and hence gives validation of the polynomial form.

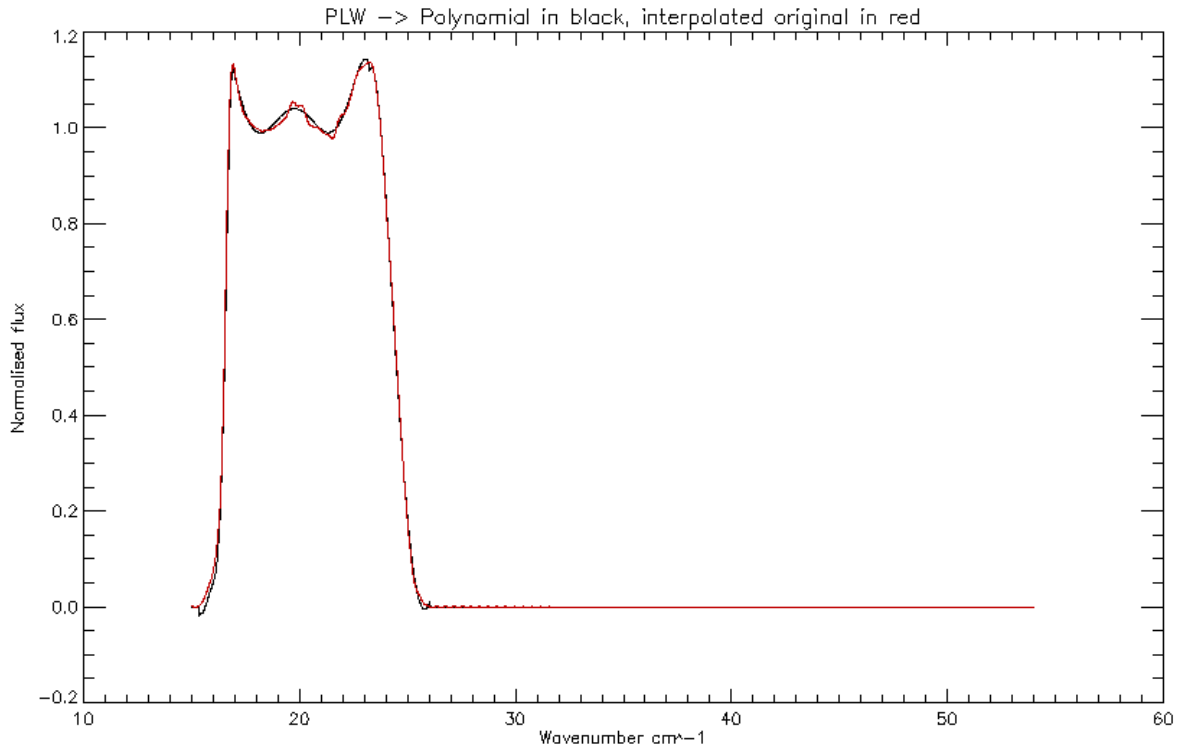

Figure 5.5: PLW comparison of polynomial fit to original.

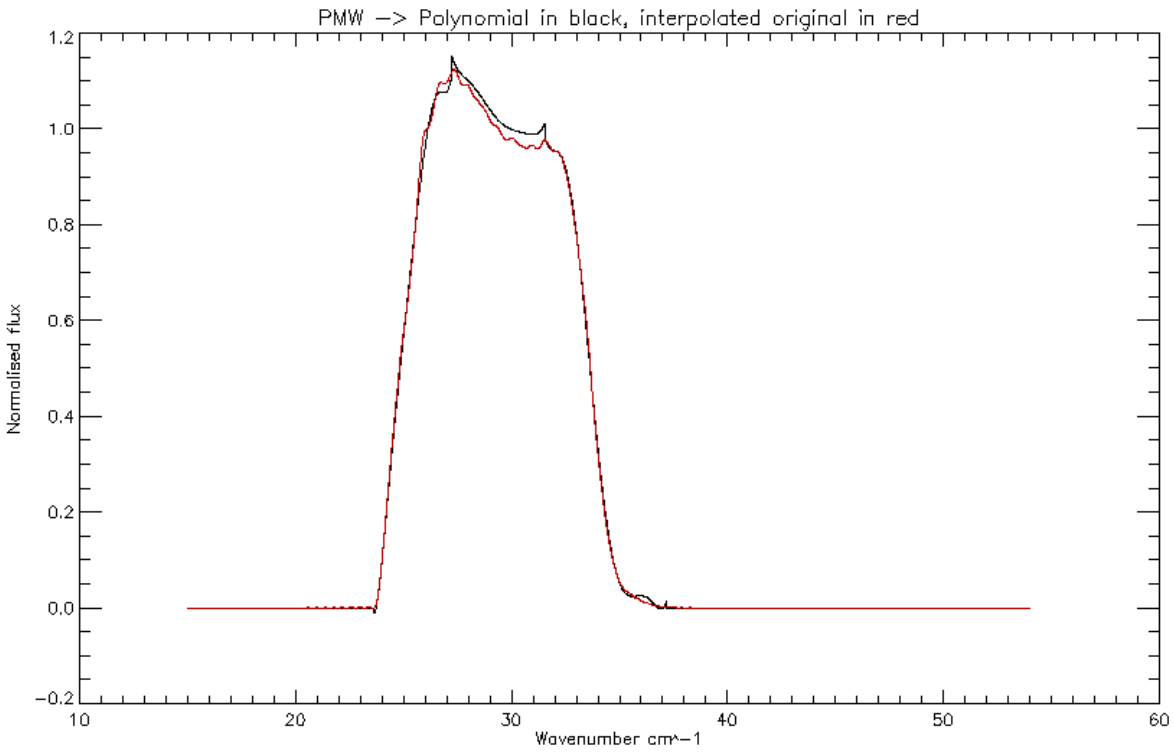

Figure 5.6: PMW comparison of polynomial fit to original.

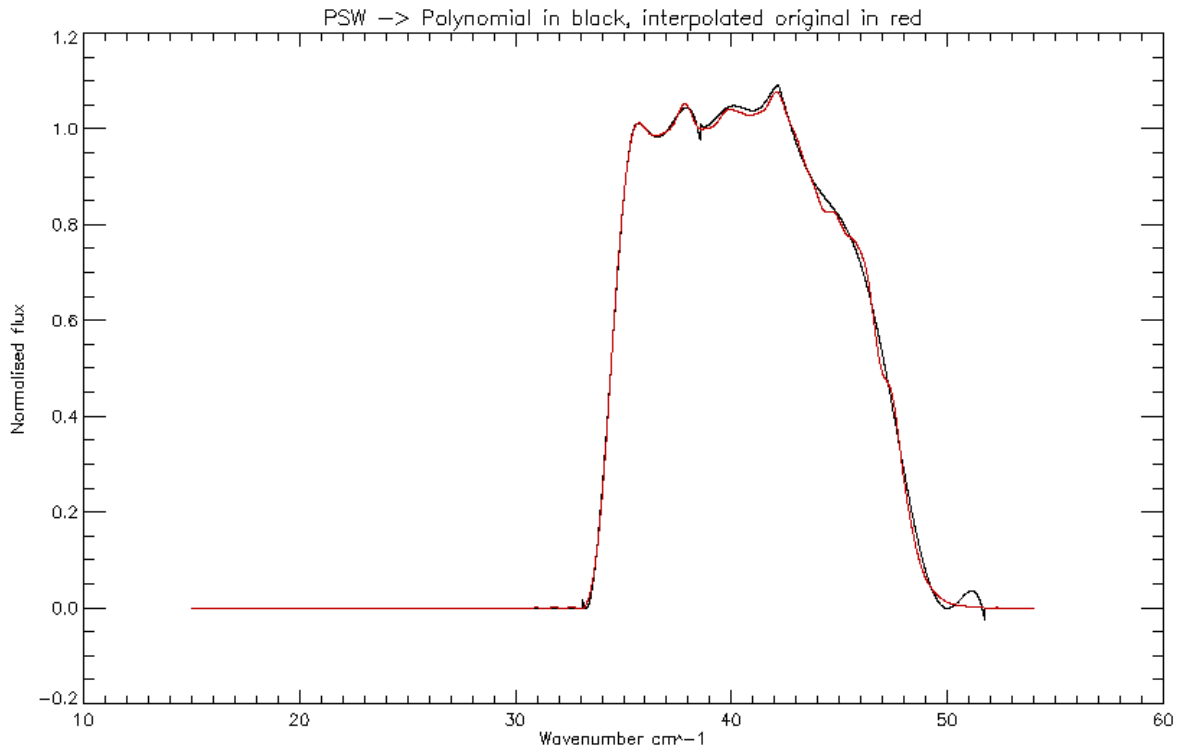

Figure 5.7: PSW comparison of polynomial fit to original.

Relative Error between Interpolated and Polynomial Form									
	Baseline $Jy\ cm^{-1}$			**Continuum** $Jy\ cm^{-1}$			**Original** $Jy\ cm^{-1}$		
Source	PLW	PMW	PSW	PLW	PMW	PSW	PLW	PMW	PSW
Arp 220	0.023	0.004	-0.001	0.025	-0.010	-0.010	0.024	-0.010	-0.010
DR21	0.057	-0.021	0.020	0.025	-0.011	-0.011	0.026	-0.010	-0.010
IC342	0.028	-0.009	0.014	0.025	-0.010	-0.010	0.025	-0.010	-0.010
M82	0.027	-0.009	-0.058	0.025	-0.010	-0.010	0.025	-0.010	-0.010
VY CMa	0.023	-0.013	-0.006	0.025	-0.010	-0.010	0.024	-0.010	-0.010

Table 5.2: Comparison of polynomial and interpolated RSRF.

5.3.1 Simulating observations

A custom tool in IDL was then created, through which we can process SPIRE spectra. For each spectrum, we multiply the flux by our analytical function for each filter band and then sum the result. This gives us a simulated photometer observation for each band. To tabulate the amount of line contamination we would see in our simulated observations of the SPIRE sources, we use the custom line fitting tool SLIDE (SPIRE Line IDentification and Extraction). We discussed SLIDE in more detail in Chapter 4. SLIDE provides not only line parameters, but also continuum parameters. From SLIDE, we use the continuum subtracted baseline, as a line information only spectrum, and the original spectrum, which contains the line and continuum information.

Using our IDL custom filter tool, we then create simulated photometer observations for the continuum subtracted baseline spectrum (or CSBS for short) and the original spectrum. The CSBS provides us with the line contribution across the three bands, whereas the original spectrum contains the line and continuum contribution. We can then examine the ratio between the CSBS simulated observation (P_L) and the simulated observation of the original spectrum (P_O). P_O/P_L gives us a measure of the fractional contribution of the lines to the observation.

Tables 5.3 and 5.4 give line contamination percentages for point-source and extended-source calibrated SPIRE spectra. For point source spectra, PMW data for the midband was obtained by joining the two spectra together and using SSW data for the overlap. With extended sources, the discontinuity meant that only PLW and PSW was considered. The beam size is a factor of four times bigger for the SLW band than the SSW, making an accurate simulation of the PWM band difficult without extensive knowledge of the emission distribution in the extended

sources. PLW and PSW are unaffected as they lie completely within SLW and SSW respectively.

Source		P_0/P_L PLW	P_0/P_L PMW	P_0/P_L PSW
CRL 618	Polynomial	0.014	0.029	0.049
	Interpolate	0.013	0.029	0.050
IRC 10216	Polynomial	0.302	0.217	0.154
	Interpolate	0.301	0.217	0.156
Mrk 231	Polynomial	0.077	0.008	0.004
	Interpolate	0.076	0.009	0.004
NGC7027	Polynomial	0.014	0.026	0.066
	Interpolate	0.014	0.026	0.068
VY CMa	Interpolate	0.161	0.134	0.115
	Polynomial	0.159	0.134	0.116
Arp220	Polynomial	-0.001	-0.003	0.000
	Interpolate	-0.001	-0.004	0.000

Table 5.3: Point Source calibrated Line Contamination data.

Source		P_0/P_L PLW	P_0/P_L PSW
IC342	Polynomial	-0.010	0.000
	Interpolate	-0.012	0.000
DR21	Polynomial	-0.003	0.000
	Interpolate	-0.004	0.003
G29.96	Polynomial	-0.002	0.001
	Interpolate	-0.002	0.001
G32.80	Polynomial	-0.003	0.001
	Interpolate	-0.003	0.001
M82	Polynomial	0.032	0.001
	Interpolate	0.032	0.001

Table 5.4: Extended Source calibrated Line Contamination data.

This procedure was then repeated using the baseline output file and the summed sinc fit files produced by SLIDE. These are synthetic spectra where the baseline output file is the fitted baseline SLIDE subtracts from the original spectrum to produce the CSBS. This becomes the synthetic, noiseless continuum. The summed sinc fit file is the total sum of all the sinc fits used by SLIDE to recover the files. This is a synthetic, noiseless CSBS amd we term this P_{SYN-L}. Summing the synthetic continuum to the synthetic CSBS, we produce a synthetic spectrum P_{SYN-O}. We can again take the ratio of P_{SYN-O}/P_{SYN-L} to obtain a measure of the fractional line contribution for our synthetic observations. This is useful in cases where weak lines, combined with a slightly offset continuum fit, yield negative values for P_O/P_L for a photometer band. It is also useful to help remove any artefacts caused by ripples or noise in the FTS spectra, or features which are instrumental effects.

The results are included in Tables 5.5 and 5.6. For the point source cali-

brated observations, we combined the SLW and SSW datasets into a combined dataset using the SSW data to cover the overlap. For extended source calibrated datasets, the discontinuity due to beam sizes means we consider only PLW and PSW. To accurately model PMW in this case would require a detailed analysis of the differential beam sizes between the SLW and SSW data.

Figures 5.8 and 5.9 show the synthetic observation, which can be visually compared to Figures 5.2 and 5.3. We see there is far less noise and other processing artefacts, as these SED have been artifically generated using the greybody fit of the original SED as a synthetic continuum. We then summed up the sinc fits of each line in the SED to generate a synthetic set of representative lines. We then generated the synthetic SED as a simple summation of all sinc profile lines and the greybody fit.

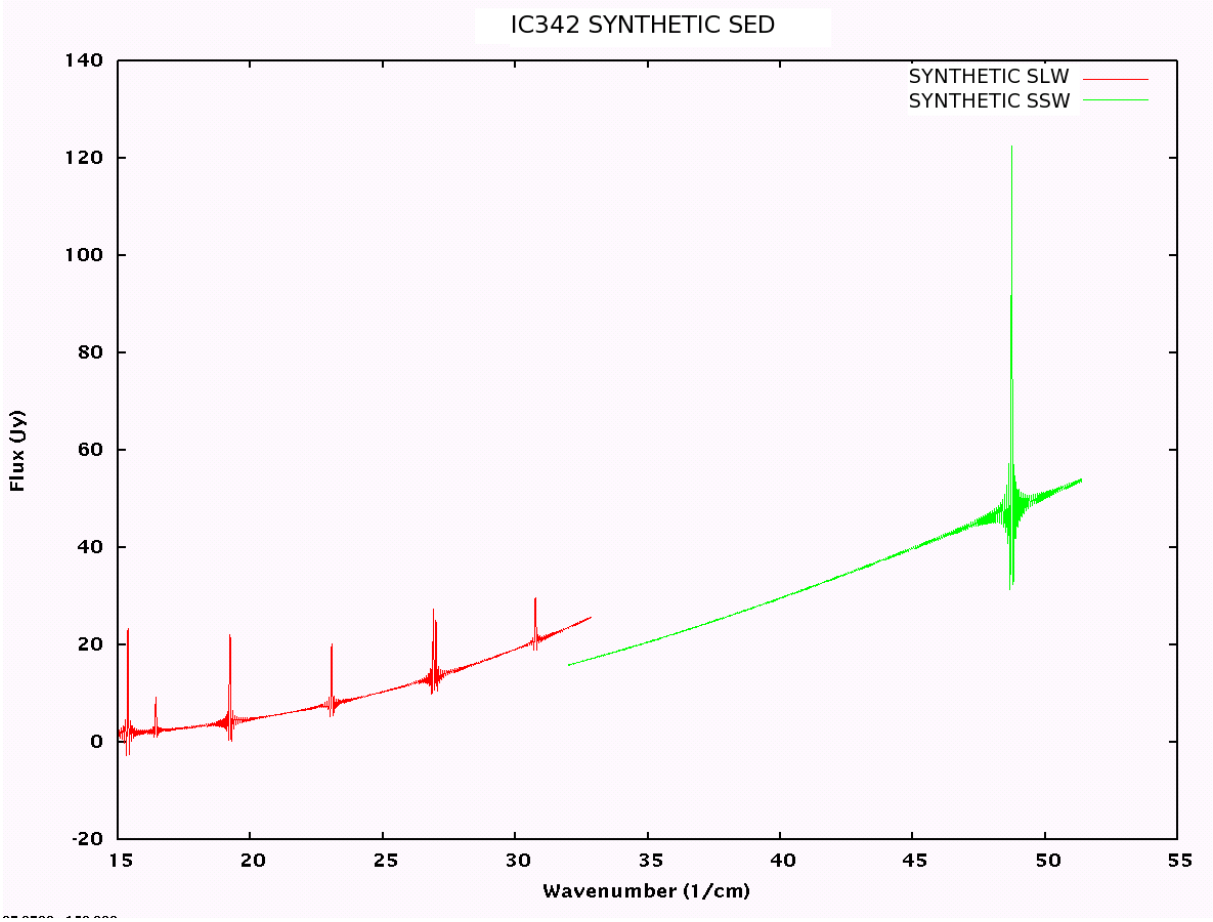

Figure 5.8: IC342 synthetic observation used for line contamination study.

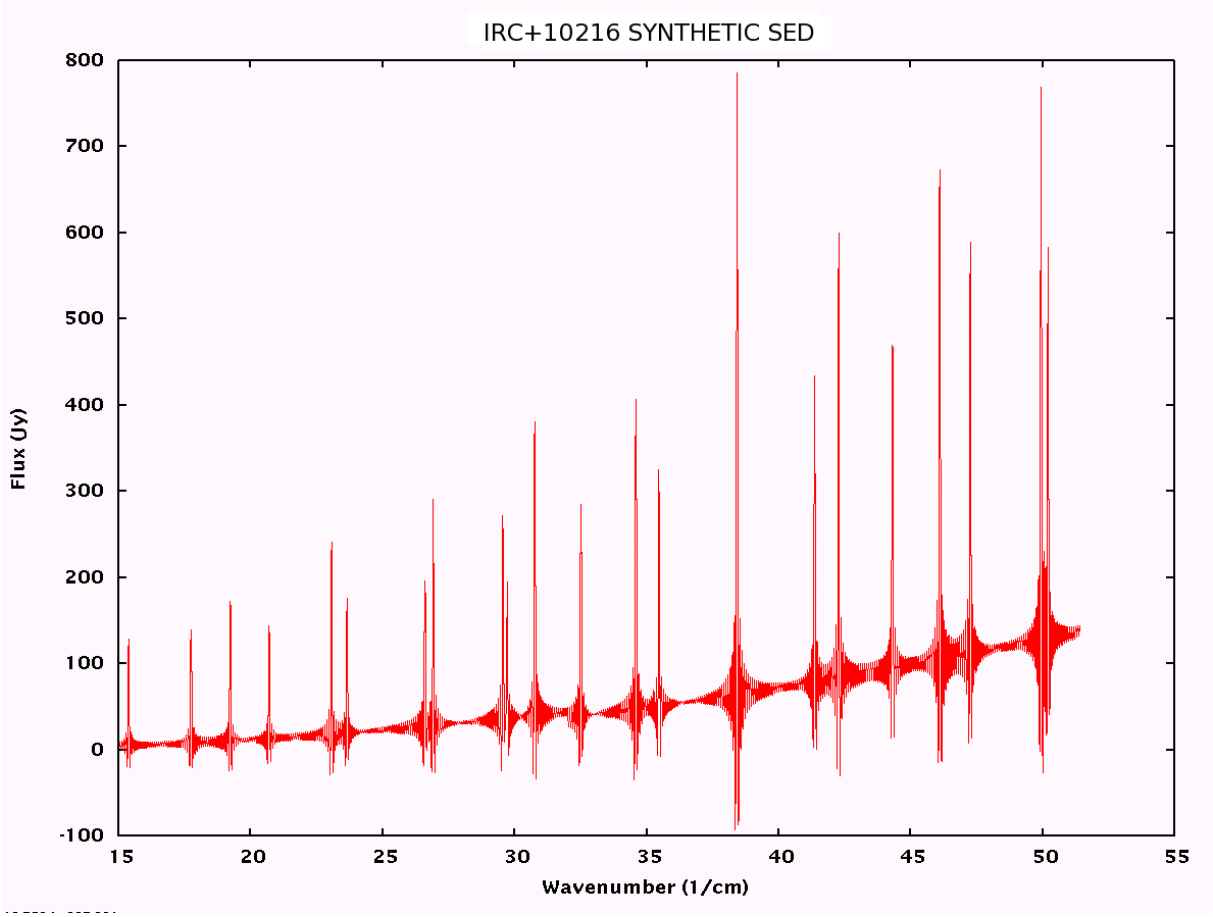

Figure 5.9: IRC+10216 synthetic observation used for line contamination study.

Source		P_0/P_L PLW	P_0/P_L PMW	P_0/P_L PSW
CRL 618	Polynomial	0.046	0.042	0.036
	Interpolate	0.045	0.042	0.037
IRC 10216	Polynomial	0.234	0.162	0.114
	Interpolate	0.233	0.162	0.116
Mrk231	Polynomial	0.024	0.015	0.004
	Interpolate	0.024	0.015	0.004
NGC 7027	Polynomial	0.070	0.094	0.072
	Interpolate	0.069	0.093	0.073
VY CMa	Polynomial	0.140	0.117	0.102
	Interpolate	0.139	0.117	0.102
Arp220	Polynomial	0.018	0.010	0.003
	Interpolate	0.018	0.010	0.003

Table 5.5: Point Source calibrated Line Contamination data for synthetic spectra.

Source		P_0/P_L PLW	P_0/P_L PSW
IC342	Polynomial	0.029	0.001
	Interpolate	0.029	0.001
DR21	Polynomial	0.020	0.010
	Interpolate	0.020	0.010
G29.96	Polynomial	0.008	0.003
	Interpolate	0.008	0.003
G32.80	Polynomial	0.010	0.003
	Interpolate	0.010	0.003
M82	Polynomial	0.060	0.005
	Interpolate	0.060	0.005

Table 5.6: Extended Source calibrated Line Contamination data for synthetic spectra.

5.3.2 Results from simulations

From the original spectra we saw some negative line contamination percentages. This was a result of incorrect baseline fitting on sources with weak lines. Using the synthetic spectra, we note that line contamination in the SPIRE bands can be up to around 25% of the flux in that band. Tables 5.3 to 5.6 show that the largest contribution is from PLW. This is to be expected as the continuum emission is proportional to $\nu^{\beta+2}$ and the continuum emission will be weakest in the PLW band. PLW is also the widest band of the three. The line contamination percentage is generally low $(5-10\%)$, with the line-rich IRC10216 providing the largest percentage contribution (around 25%). Figures 5.2 and 5.3 show these extremes graphically.

5.4 Conclusions from RSRF simulations

We can see that the polynomial forms of the RSRF functions are an excellent fit, and so we can use these polynomial fits to simulate the RSRF for greybody recovery with confidence. We now have an approximate upper limit for the line contamination within the SPIRE bands, and can use this to evaluate the effect on greybody fitting. Given the SPIRE range of $\approx 200-650\mu m$, we can use Wien's Law to obtain temperature parameters for the peak emission within the band. This range equates to blackbody temperatures of between $5-15K$ and so we use a range of greybodies from $5-25K$ to adequately cover this range and to provide simulations where the peak is outside our range of interest.

Simulations where the peak is outside this range will help evaluate how well

fitting behaves when we do not have the greybody peak within our simulated data. We can then compare how well fitting occurs with or without the peak in the band to help constrain the fit. We also evaluate $\beta \approx 1 - 2$ in our simulations, and can examine special cases within our dataset. We can allow for T, β to be free parameters, constrain β to 2 (as per Ward-Thompson et al. (2010)) and allow T to float, or conversely constrain T to either $10K$ or $25K$. $10K$ has a peak within the SPIRE band, and $25K$ has a peak outside it. This feature set comprised the initial simulations, and with the information gleaned from this we expanded the simulations appropriately.

5.5 Line contamination and SED analysis

We now outline the process where we analyse simulated photometry observations with varying degrees of line contamination. We then attempt to fit a grey body curve to these observations and compare the results of the greybody fit to the original defining parameters. By this method, we can accurately assess the degree of modification that line contamination will give.

5.5.1 Greybody SED's and simulations

We define our greybody where F_ν is the monochromatic flux density of the grey body at a frequency ν, $B_\nu(T)$ is the Planck function, ν_c is the frequency at which the optical depth is unity, Ω is the solid angle, f is the filling factor of the source within the observing aperture and β is the dust emissivity index. Following Ward-Thompson et al. (2002) and Kirk et al. (2005, 2007) we take ν_c to be $50\mu m$, and approximate the optical depth $\tau \approx (\nu/\nu_c)^\beta$ to obtain Equation 5.2, where we see

emission reduced by a factor of $e^{-\tau}$ due to optical depth effects. Using the following greybody formula:

$$F_\nu = \Omega f B_\nu(T) \left(1 - \exp\left[-\left(\frac{\nu}{\nu_c}\right)^\beta\right]\right), \tag{5.2}$$

we generate a number of greybody spectra using a range of β from 1 to 2, and T from $5K$ to $25K$. These greybody spectra will then be used to create synthetic photometry observations to examine line contamination effects. For the simulations, we can ignore the filling factor and the solid angle and simplify the above equation to:

$$F_\nu \propto B_\nu(T) \left(1 - \exp\left[-\left(\frac{\nu}{\nu_c}\right)^\beta\right]\right). \tag{5.3}$$

We isolate the values of the grey body at the frequency bands of the photometer. We also add $850\mu m$ SCUBA observations and $1200\mu m$ IRAM observations to give us five simulated observations. To simulate noise, we use 10% of the flux at each of our five wavelength bands and use that as a $1-\sigma$ measure of the random noise. We then add a random $n-\sigma$ level of noise to each band. This was chosen as a realistic estimate of the SPIRE noise level. Our simulation uses a fixed random seed to ensure that the noise is consistent across the simulations. From our SPIRE observations, we use a maximum value of 25% as a good estimate of the maximum degree of line contamination (see above). We then included varying degrees of line contamination between 5% and 25%.

This gives us a matrix of greybody spectra with varying degrees of line contamination. We can then take these five datapoints, and attempt to fit a greybody

curve to them. We can investigate the effect of fixing either T (the temperature of fit) or β (the dust emissivity) on the quality of fit as well as using a metric to investigate the quality of the resultant fit by comparing both fitted SED and the original SED. These simulations have several factors of interest:

- ΔT is the difference between the modelled greybody temperature T_{model} and the recovered, fitted temperature T_{rec}, $\Delta T \equiv T_{model} - T_{rec}$
- $\Delta\beta$ is the difference between the modelled greybody emissivity β_{model} and the recovered, fitted value β_{rec}, $\Delta\beta \equiv \beta_{model} - \beta_{rec}$
- Line contamination percentage applied, LC
- Whether line contamination is applied to a single band, or a range of bands

When considering line contamination, we consider it added to a single waveband or a group of wavebands simultaneously. As a result, the challenge from the simulations is to present the recovered information and the trends found in a legible manner. We also seek to examine interesting cases and to consider how we consider the effects on line contamination in photometer observations using SPIRE and other instruments.

5.5.2 Greybody fitting

Using Wien's law, for our five bands the peak of the Planck function moves out of our wavebands at approximately 15K. This means that for warmer greybody temperatures, we do not have the peak in our bands. When the peak is within the bands, the fitting routine has an extra 'feature' to lock onto so we can examine what happens when the peak is not available.

To begin, we can examine the effects of line contamination on recovered temperature when we add line contamination to the SPIRE photometer bands only. We can also consider the effects of line contamination on one band at a time. In the initial fitting, we fixed β to be two. Conversely, for line contamination in the other two bands (SCUBA and IRAM) we fixed the temperature and examined the effect of fitting on β. Here, we considered the effect when T is fixed to $10K$ and $25K$ and again we consider line contamination effects on fit recovery. This is useful because β variations will be seen most strongly in the longer wavelength band due to the $1 - \exp(\nu/\nu_c)^\beta$ term.

For fitting, we generate a matrix of spectra (from T_{model} = 5K - 25K, β_{model} = 1,1.5,2). A custom written IDL tool generates one thousand spectra for each combination of line contamination percentage, line contamination band, greybody temperature and β_{model} with each spectrum having different random noise levels. We then attempted to recover the temperature and β for each spectrum where we allowed both temperature and β to be free parameters in fitting, with β having a restricted range of between one and two. We took the mean recovered temperature and β for each combination and used this to analyse the overall effect of line contamination on recovery where both parameters were free. Error bars on the fit parameters were calculated using a quadrature sum of the mean $1-\sigma$ error returned from the thousand runs, and the $1-\sigma$ of the thousand returned parameters as per the following equation

$$e_{total} = \sqrt{MEAN(e_{1...1000})^2 + \sigma_{1...1000}}. \tag{5.4}$$

This procedure was again repeated, fixing T and β to investigate their effect on fitting recovery. For further exploration of the parameter space, the simulations were expanded using a 40K greybody as the base. As the waveband used for the original simulations would not constrain the peak we discard the SCUBA and IRAM simulated observations and replace them with simulated PACS observations at $70\mu m$ and $160\mu m$.

5.6 Results of greybody fitting

5.6.1 Greybody fitting - initial simulations

5.6.1.1 Contamination in SPIRE bands

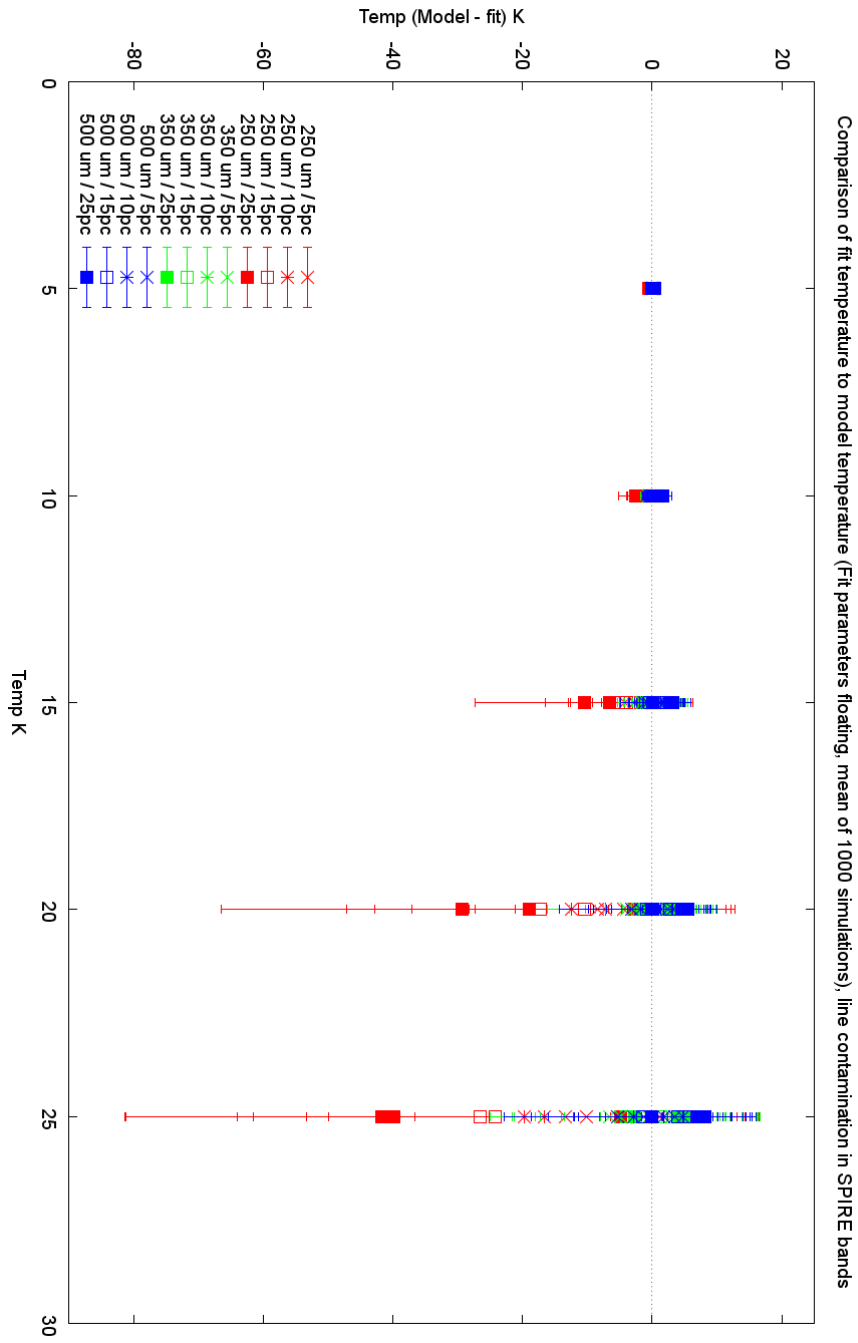

Figure 5.10: Line contamination with free fit parameters across SPIRE bands, overall effect on T_{rec}.

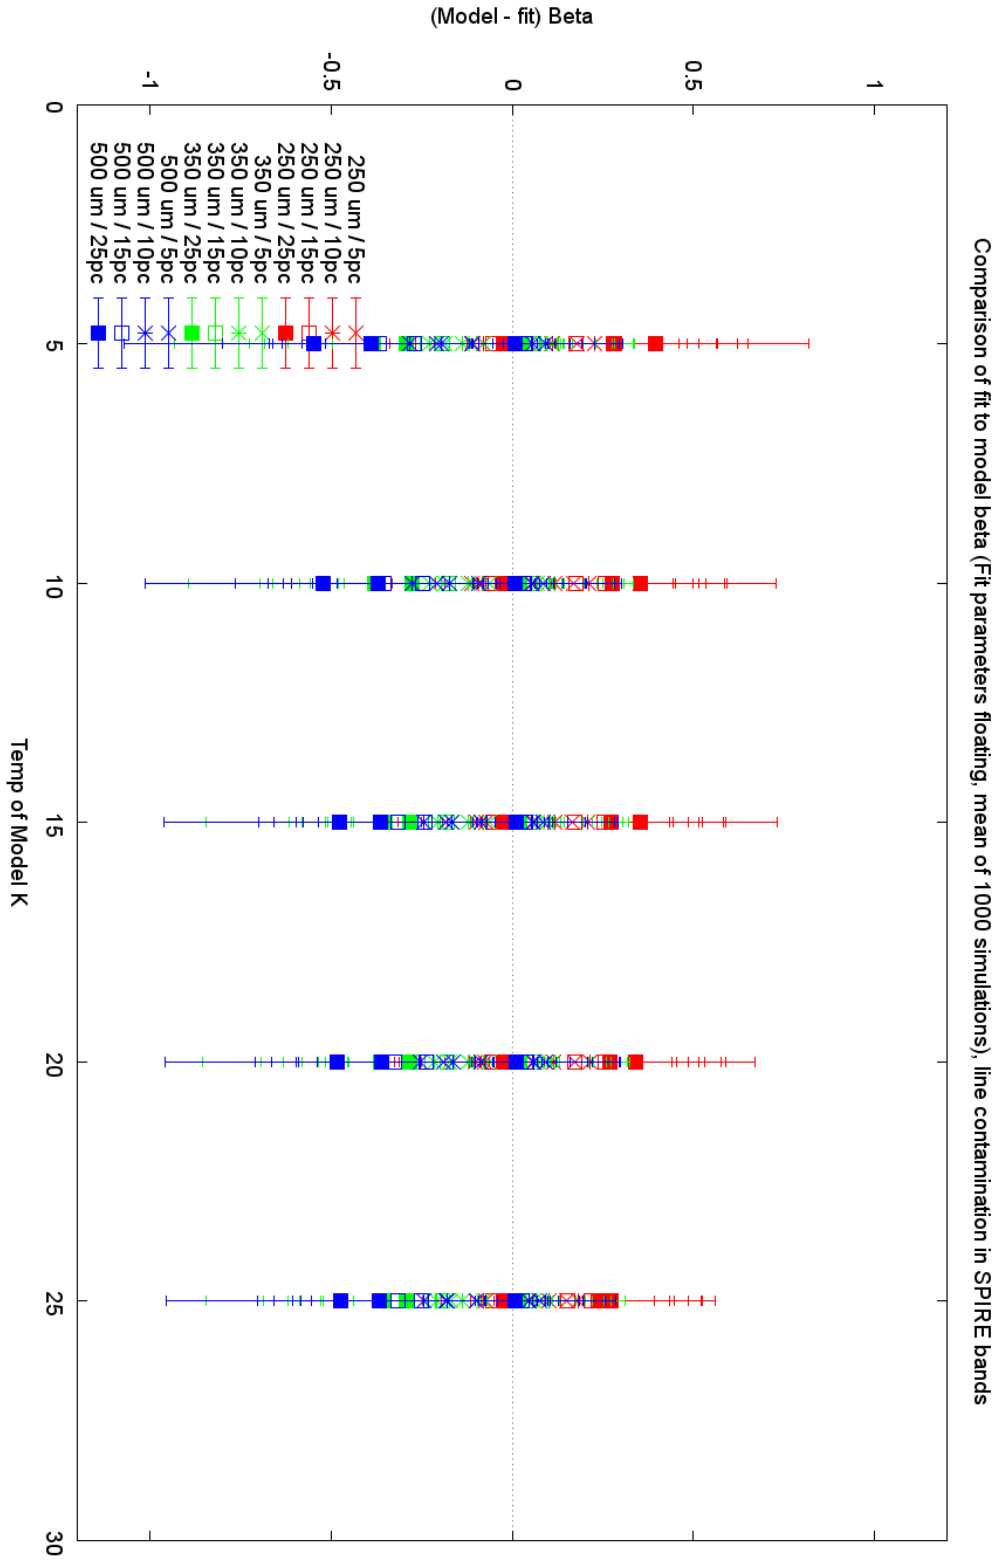

Figure 5.11: Line contamination with free fit parameters across SPIRE bands, overall effect on β_{rec}.

Figure 5.10 examines the effect of line contamination on temperature, examining the overall effect for line contamination in a single wavelength band. The y-axis graphs the difference between the model temperature and the fitted temperature. A positive value for this means that the recovered temperature is lower than the original model temperature. Here the fit parameters of β, T are constrained, with β_{rec} constrained to be between 1 and 2 and T_{rec} having a lower limit of 2K.

From this we note that:

- As the greybody temperature increases, we see an increasing range of recovered temperatures as line contamination is applied
- As the greybody temperature increases, we see that line contamination in the $250\mu m$ band causes the recovered temperature to be driven higher than the model
- Line contamination in the $350\mu m$ and $500\mu m$ bands drives the fit temperature lower
- As the greybody temperature increases, the fit errors increase and the range of recovered temperatures increases (to around 100K for a 25K model)

Increasing line contamination for peak temperatures between 5 - 10K show little deviation from the original model. Here the peak is within our wavelength range, and overall the deviation is under 5K. Line contamination in the $250\mu m$ band drives the fit temperature up, and is expected as here it would appear to 'push' the peak of the grey body towards the shorter wavelength range. Conversely, line contamination in the longer wavelength bands drives the fit temperature down as the peak is shifted towards longer wavelengths. For higher temperatures, where

the peak of the greybody is outside our range, we see a similar effect. However, the effect is magnified. This shows the difficulty of using a small number of points to fit a greybody in this range. Without a peak to help constrain the fit temperature, we see larger errors in the recovered temperature. The temperature essentially becomes unconstrained.

Figure 5.11 does not reveal a lot of information. The recovered β_{rec} varies across the range of the parameters, and there is only a slight tendency for the range to narrow with increasing temperature. For shorter wavelengths, the fit β_{rec} is lower than the original greybody. For longer wavelengths, this is reversed.

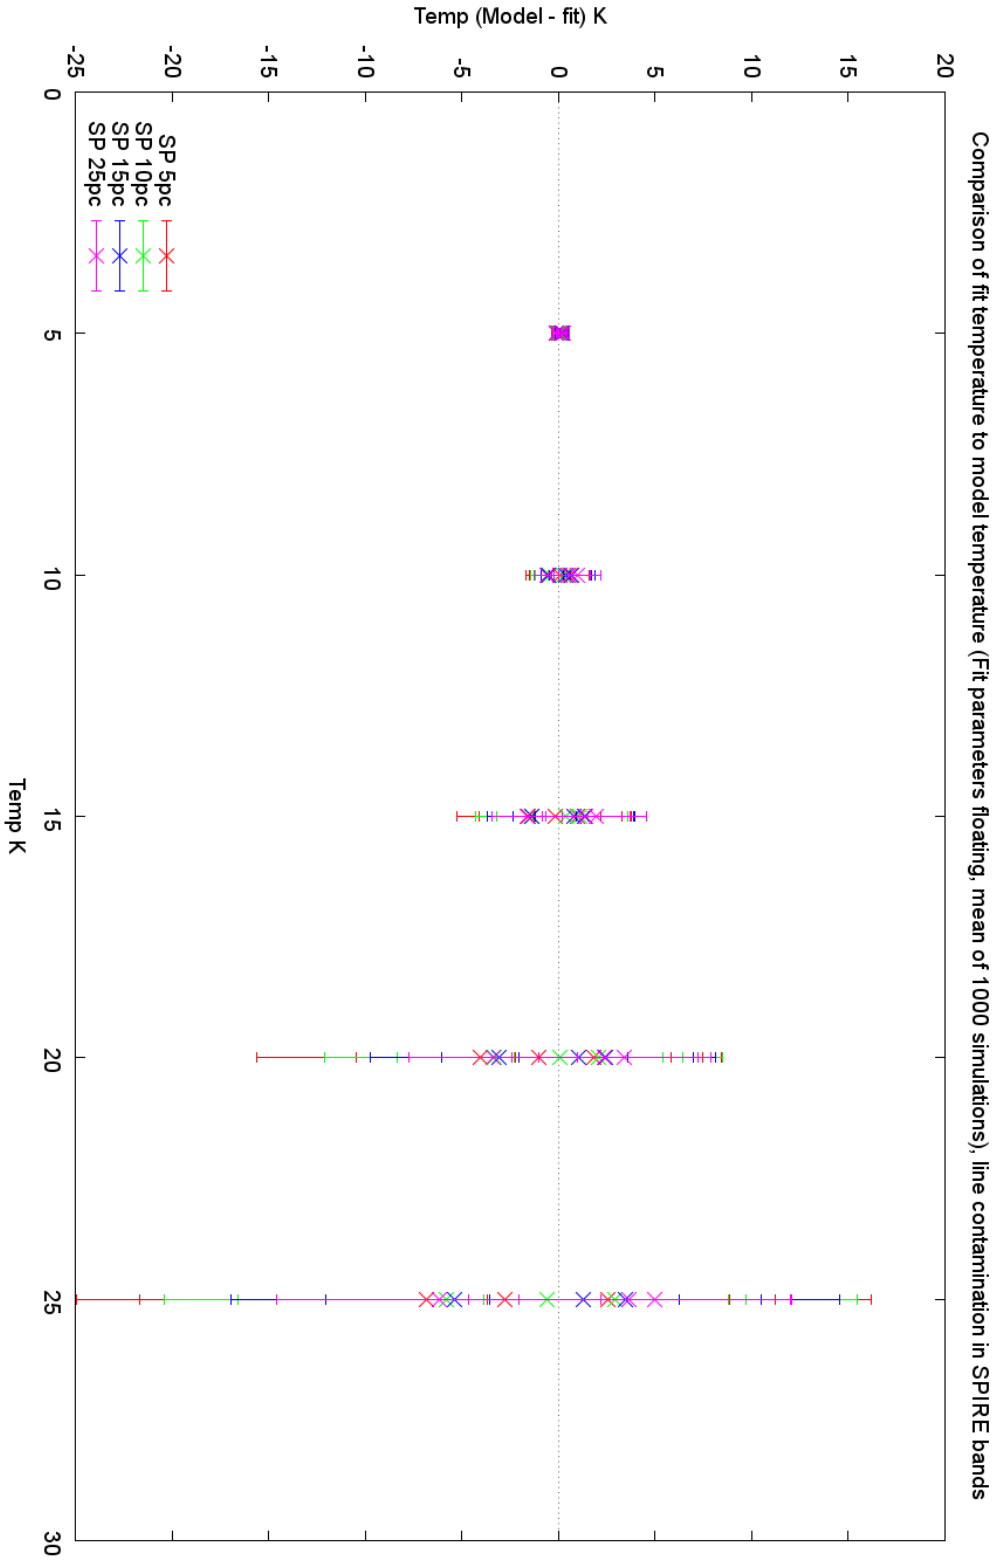

Figure 5.12: Line contamination with free fit parameters across SPIRE bands, effect on T_{rec}.

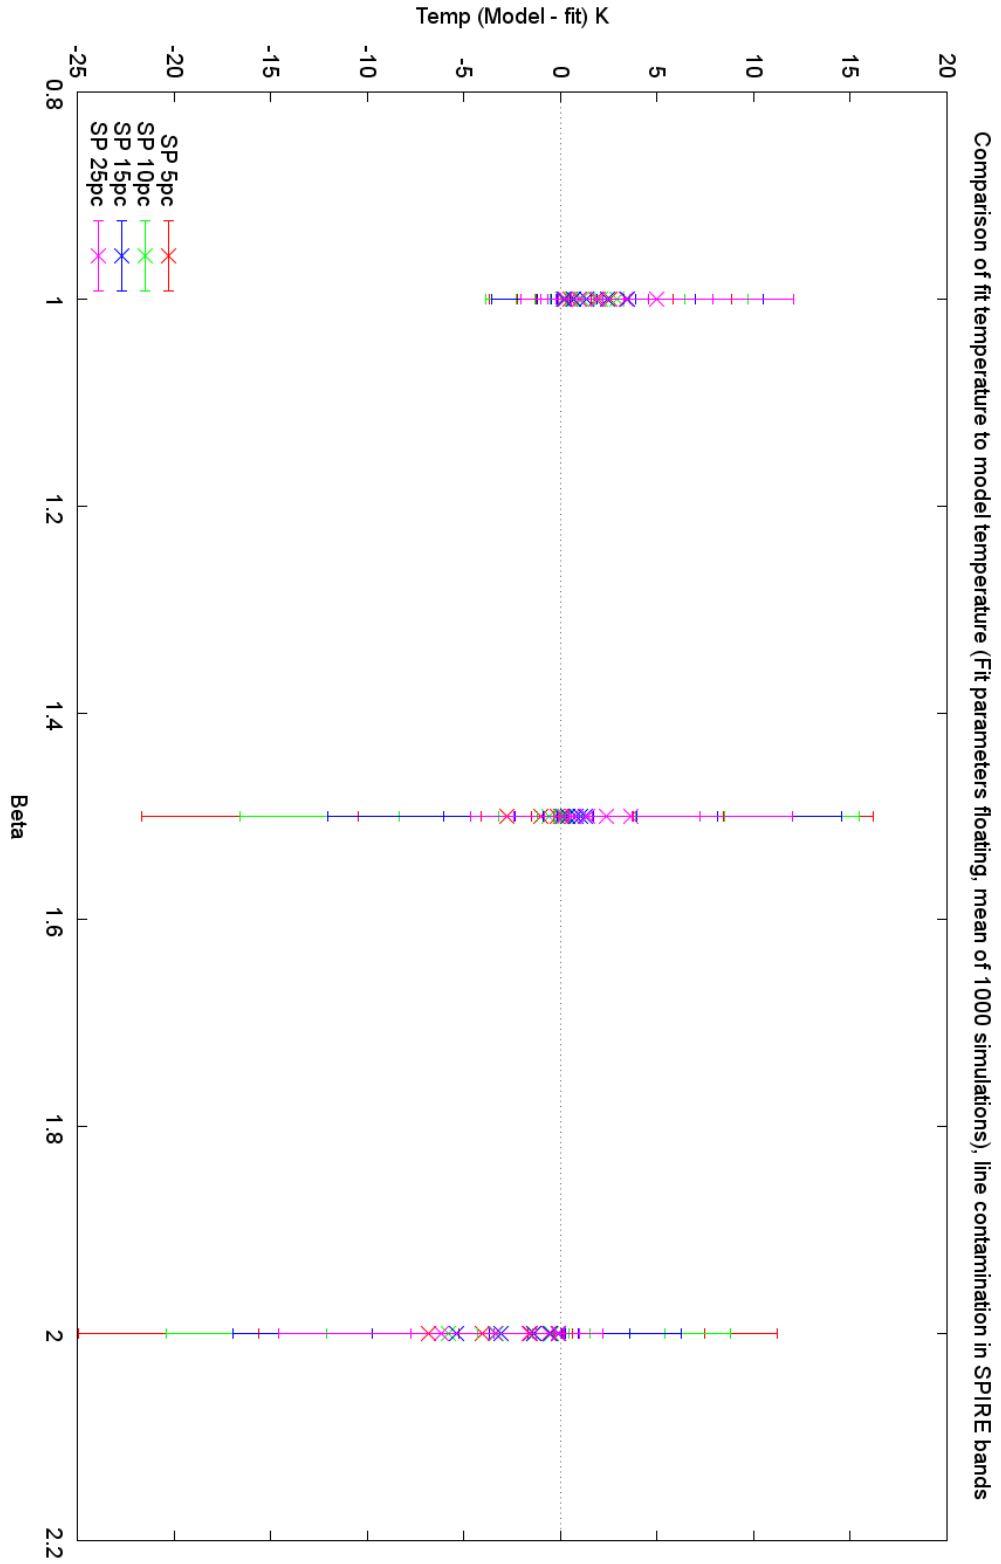

Figure 5.13: Line contamination with free fit parameters across SPIRE bands, effect on T_{rec}, grouped by β_{model}.

Figures 5.12 and 5.13 examine the effect of line contamination when it is applied globally across the SPIRE band. Whereas the previous figures show varying line contamination percentages on a single wavelength, here the same line contamination percentage is applied simultaneously to the three SPIRE wavelength bands. These simulations, where the line contamination is applied to a group of wavelength bands simultaneously are termed GLOBAL. LOCAL simulations are where the line contamination is applied to a single wavelength band, Here, we note that the range of temperature deviation from fit to model is reduced significantly. Figure 5.12 shows that the recovered temperature has a range of 50K for the 25K greybody, compared to a 100K deviation for the individual case.

As the model temperature increases, we see a slight trend towards the fit temperature being higher than the model. Increasing the model temperature also increases the errors on the recovered fit, and this correlates to what we saw in the previous section. Here we also see that increasing the line contamination percentage seems to reduce the fit errors slightly. Examining the effect on β_{rec} we see that Figure 5.13 shows that higher β_{model} values for the original greybody values yield slightly larger temperature fit deviations. However we do see a tendency for higher values of β_{model} in the original greybody models to give higher recovered temperatures once line contamination is applied to the SPIRE bands. When $\beta = 1$, we see a reduced range in the recovered temperature, as well as a tendency for the fitted temperature to be lower than the model. We show the effects on β_{rec} in greater detail in Figures 5.14 and 5.15.

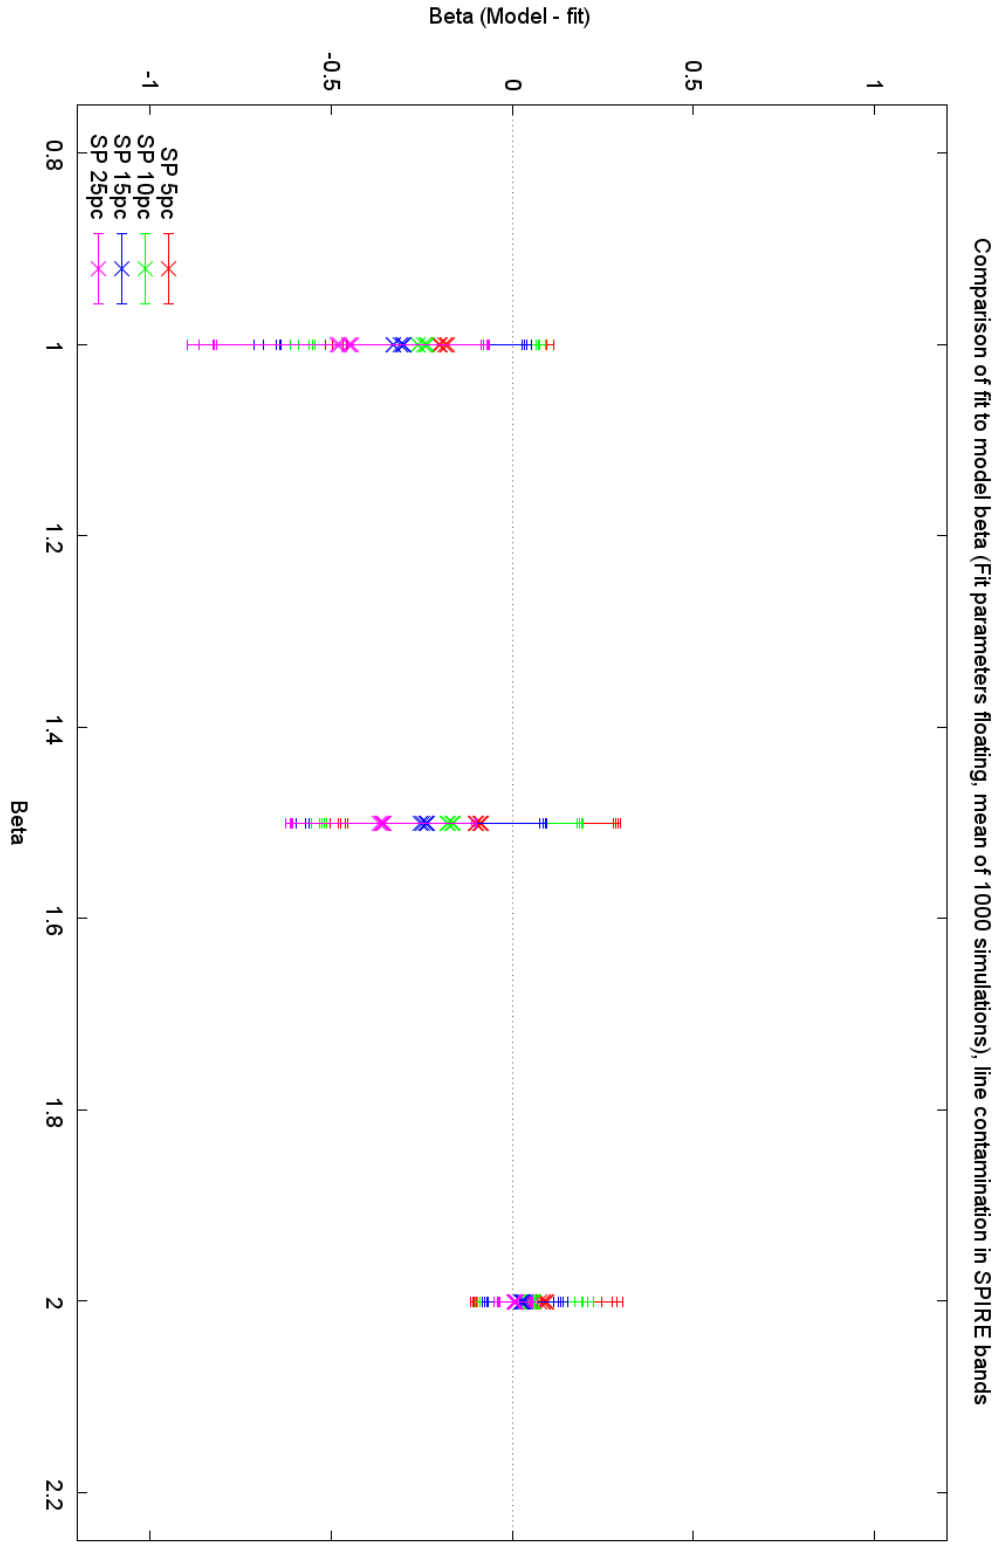

Figure 5.14: Line contamination with free fit parameters across SPIRE bands, effect on β_{rec}.

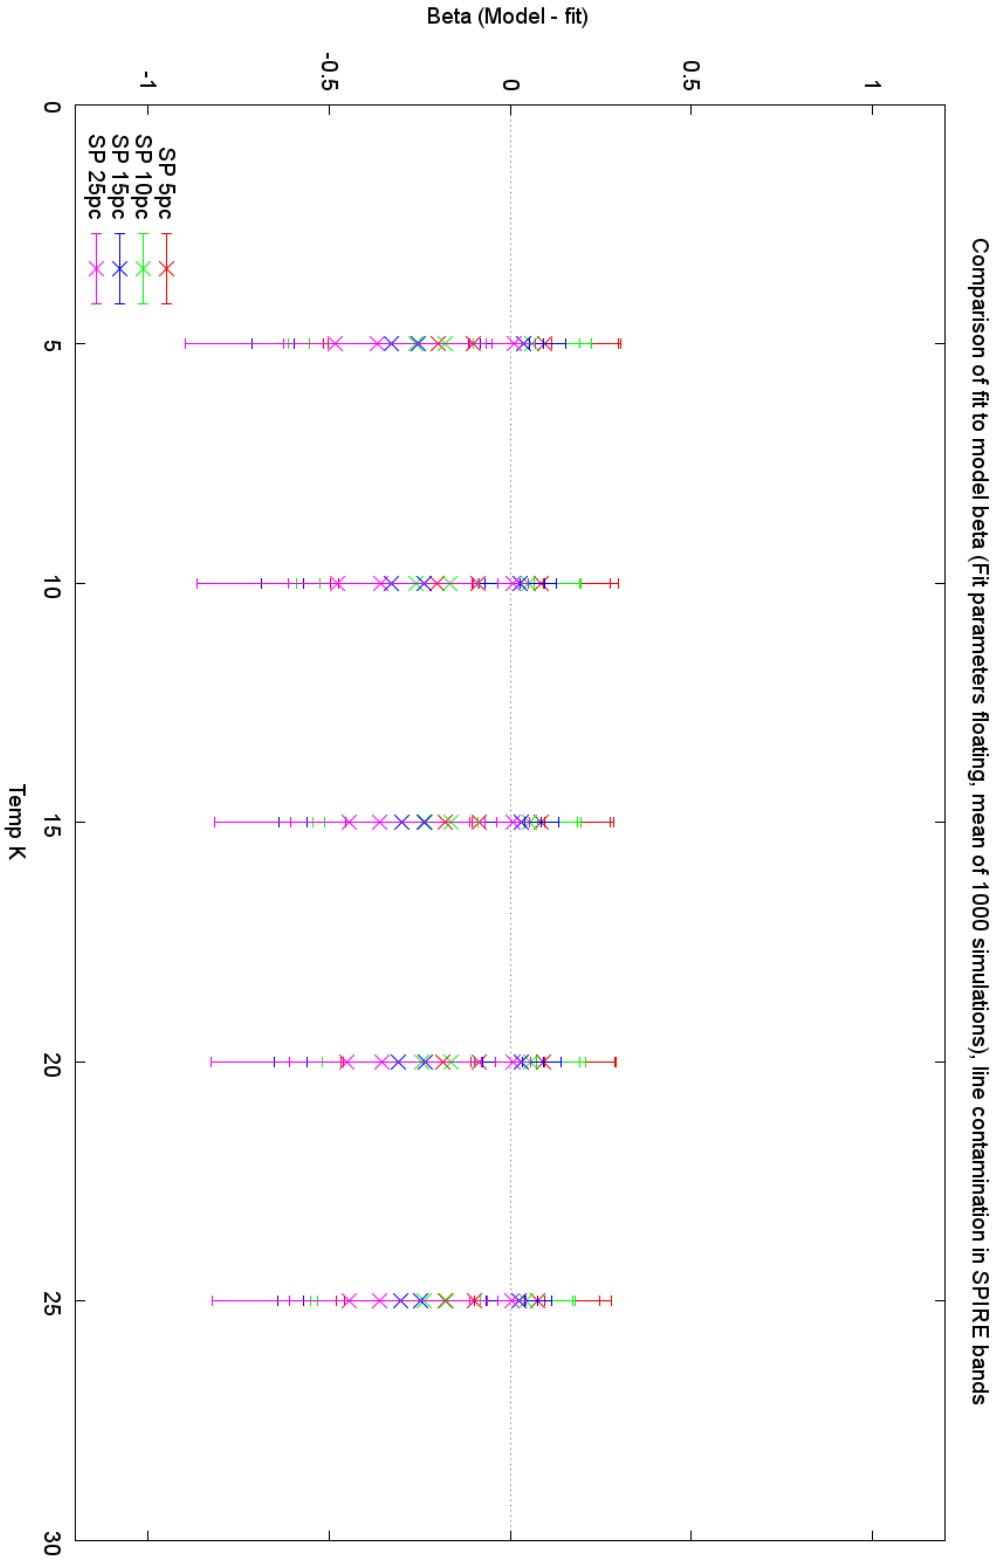

Figure 5.15: Line contamination with free fit parameters across SPIRE bands, effect on β_{rec}, grouped by T_{model}.

We first examine the behaviour of the simulations at the fit limits $\beta_{model} = 1.0$ and $\beta_{model} = 2.0$, where we find β_{rec} is constrained in one direction. For $\beta_{model} = 1.0$ we find that the addition of line contamination in the SPIRE bands leads to the recovered value β_{rec} to be higher than the initial β_{model} used to generate the SED. This is the lower limit of β and so the recovered value can not be less than this. This is accompanied by the recovered temperature T_{rec} being lower than T_{model}. As we increase the line contamination percentage, we see that at the lower β_{model} limit the deviation of T_{rec} from T_{model} increases and T_{rec} is driven increasingly lower. Comparing T_{Model} to T_{rec}, we see it is $5K : 4.87K$ and $25K : 22.41K$ for 5% line contamination and this becomes $5K : 4.74K$ and $25K : 20.0K$ at 25% contamination. In this lower limit of β_{model}, β_{rec} shows only a small range of variation over the temperature parameter space of the simulations for each line contamination percentage. For example, $\beta_{Model} : \beta_{rec}$ is 1.00 : 1.2 for the 5% case and 1.00 : 1.48 for the 25% case as temperature goes from $5K - 25K$. Here we see that line contamination has a weak effect on T_{rec} and a slightly stronger effect on β_{rec}.

Where $\beta_{model} = 2.0$, we recover a higher greybody temperature and a slightly lower β_{rec} value to compensate for the loss of freedom in fitting β_{rec}. Here the behaviour of T_{rec} is largely constant. For example, in the 5% line contamination case we find that $T_{Model} : T_{rec}$ is $5K : 5.13K$ up to $25K : 31.84K$. At the other extreme of 25% line contamination this becomes $5K : 5.14K$ and $25K : 31.18K$ and so we see that line contamination in this case has a weaker effect on recovered temperature than at the lower β_{model} limit. β_{rec} is recovered lower than β_{model} but β_{rec} increases slightly (0.02) as T_{model} increases. Here we find $\beta_{Model} : \beta_{rec}$ is 2.0 : 1.90 at 5K for 5% line contamination and is 2.0 : 1.92 at 25K for the same line contamination.

For the $\beta_{model} = 2.0$ case, we also find that β_{rec} increases slightly for increasing line contamination. Here, we see that $\beta_{Model} : \beta_{rec}$ is 2.0 : 1.99 at 5K and 25% line contamination - compared to 2.0 : 1.90 as shown in the last paragraph for 5% line contamination. We see $\beta_{Model} : \beta_{rec}$ is 2.0 : 1.99 for 25K and 25% line contamination which we can again compare to the last paragraph. We see that at the highest line contamination percentage, the range of increase in β_{rec} is constrained by the fit limit. The increase in β_{rec} in the 25% case from the 5K to the 25K case is only 0.03.

Comparing the behaviour of β_{rec} at the two limits we see that the recovered temperature T_{rec} shows different behaviour dependent more on β_{model} than the line contamination applied. The behaviour of T_{rec} is more strongly dependent on β_{model}. Overall, the behaviour of β_{rec} is largely consistent over the two limits. As expected its behaviour is constrained, but apart from the sign of the deviation the behaviour is comparable.

We see a turnover of behaviour in recovered temperature for $\beta_{model} = 1.5$, where there is greater degree of freedom for fitting β_{rec}. Here we find that for lower line percentages $(5\%, 10\%)$ and for $T_{model} <\approx 15K$ we see $\beta_{rec} \geq \beta_{model}$ and $T_{rec} \leq T_{model}$. As T_{model} increases we see an inversion of this behaviour with T_{rec} becoming $\geq T_{model}$. The changeover temperature for this behaviour is $\approx 15K_{5\%}, 25K_{10\%}$. However β_{rec} stays approximately constant across the temperature range for each line contamination percentage. $\beta_{rec} \approx 1.6$ for 5% contamination and $\beta_{rec} \approx 1.85$ for 25% contamination with a limited range of around 0.02. The behaviour of β_{rec} in the middle range mirrors that of the lower limit. However, as line contamination increases above 10%, we lose this temperature-dependent inversion behaviour and the degree of deviation between model and recovered temperature also increases. The inversion temperature corresponds to roughly where the peak moves out of

the waveband we are considering. For 5% line contamination,$T_{Model} : T_{rec}$ is $5K : 4.96K$ and $25K : 27.75K$ and becomes $5 : 4.83K$ and $25K : 21.33K$ at 25% contamination.

We define a metric χ to judge the quality of fit by comparing the original, noiseless greybody data to the fit to the line-contaminated data. We define this metric as per Equation 5.5 where $S_o(\nu)$ is the original data, $S_{LC}(\nu)$ is the line contaminated data and $\mathcal{E}(\nu)_{rec}$ is the random noise added (treated here as an error in the line contaminated value), yielding

$$\chi = \Sigma_\nu((S_o(\nu) - S_{LC}(\nu))/\mathcal{E}(\nu)_{rec})^2. \tag{5.5}$$

At low line contamination percentages, χ is found to be roughly constant across the range of T_{model} and β_{model}. The fit quality improves slightly for lower β_{model} and degrades for higher β_{model}. This again suggests that the fixed limits of β_{model} causes the fit recovery to degrade at the higher limit. As line contamination increases, the higher β_{model} values show an increasing degradation in the fit quality compared to the lower β_{model} simulations with the same initial T_{model}. When we consider the $1-\sigma$ errors the range of the recovered temperature widens significantly, increasing the deviation from the fitted temperature to the original greybody temperature from $\approx 5K$ to $\approx 25 - 30K$.

In summary, line contamination has a complex effect on the recovery of the SED parameters. By constraining β_{model} to values currently seen in the literature, we note that the behaviour of the recovery changes at the fit limits. At the limits the behaviour of β_{rec} as a fit parameter is constrained, in that β_{rec} has only one direction of freedom and the resultant fit is defined by which limit we are at.

For central values of β_{model} where there is no such strong constraint on the fit parameter, we see behaviour that is dependent on the initial conditions and the degree of line contamination. The slope of ΔT decreases as the line contamination increases.

5.6.1.2 Contamination in non-SPIRE bands

We can now look at the effect of line contamination on the non-SPIRE bands ($850\mu m$ and $1200\mu m$) in our simulations. These are shown in Figures 5.16 to 5.22.

At first glance, these plots would seem to mirror the SPIRE band equivalents. This shows that having so few datapoints to constrain the fit is a problem if we consider line contamination for each wavelength band separately.

Figures 5.16 and 5.17 show the effect on T_{rec} and β_{rec}. Comparing Figures 5.10 to 5.16 shows an overall similarity in results. As T_{model} increases there is a tendency for T_{rec} to be greater than T_{model}. This dependence is not quite so strong for the non-SPIRE case. Figures 5.11 and 5.17 show β_{rec} grouped by T_{model}, and shows there is a slight tendency for β_{rec} to be recovered higher than β_{model} for the SPIRE case, and lower than β_{model} for the non-SPIRE case

Looking at Figures 5.12 and 5.18 we see that although there are similarities, line contamination in the non-SPIRE bands drives the fitted temperature up by a larger margin once errors are considered. The actual range of fitted temperatures is doubled for non-SPIRE line contamination. This has the effect of making Figure 5.18 much more asymmetric along the x-axis than Figure 5.12. A similar comparison of Figures 5.13 and 5.19 shows that for increasing β_{model}, there is a tendency for the temperature to be driven higher for non-SPIRE line contamination.

We see the larger error bars are again showing the difficulty of fitting for so few datapoints. Increasing line contamination in the longer wavelengths will distort the greybody shape, even though the flux for these datapoints is less than for the shorter wavelengths. Figure 5.20 shows a sample greybody, with fit, showing the lower flux for the longer wavelengths.

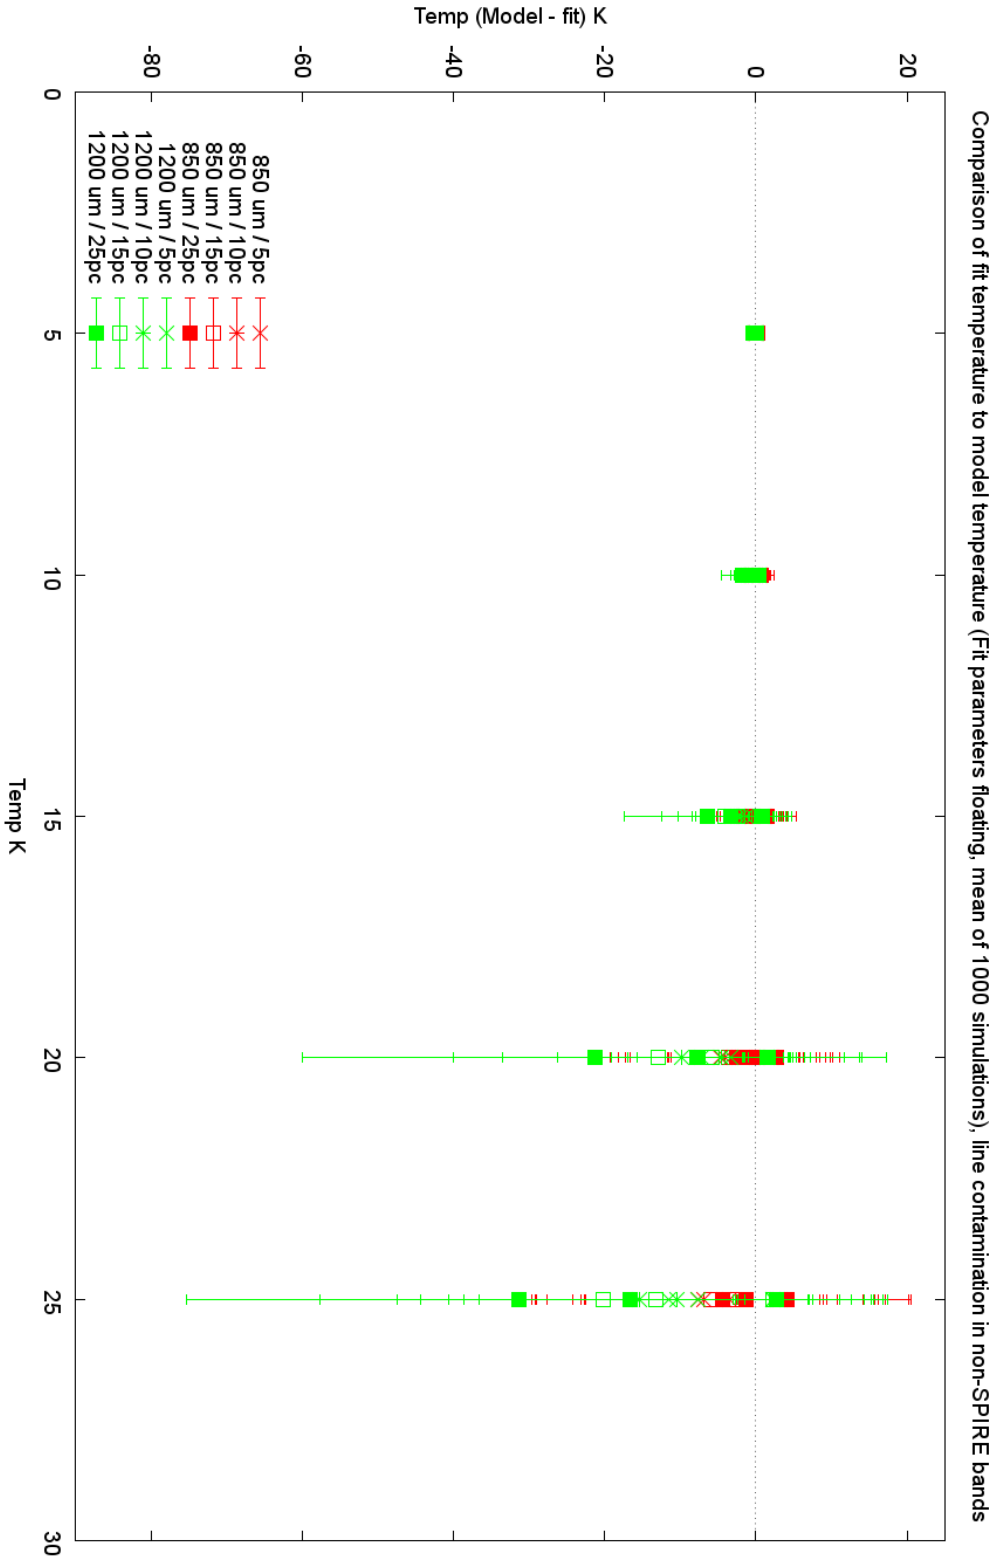

Figure 5.16: Line contamination with free fit parameters across non-SPIRE bands, overall effect on T_{rec}.

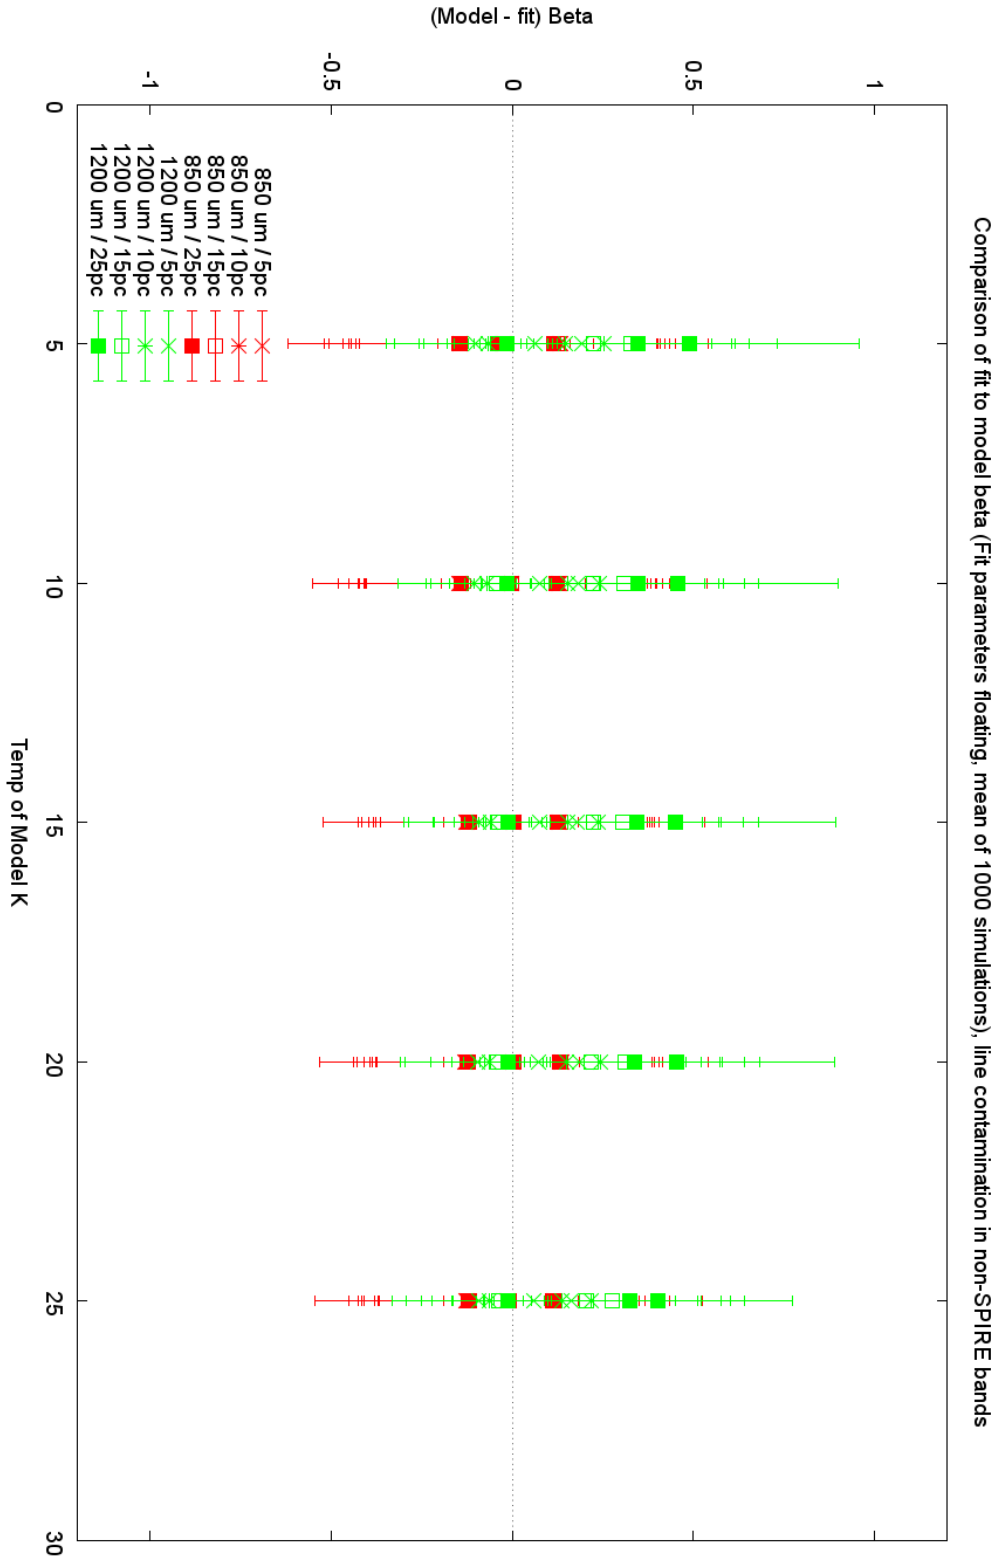

Figure 5.17: Line contamination with free fit parameters across non-SPIRE bands, overall effect on β_{rec}.

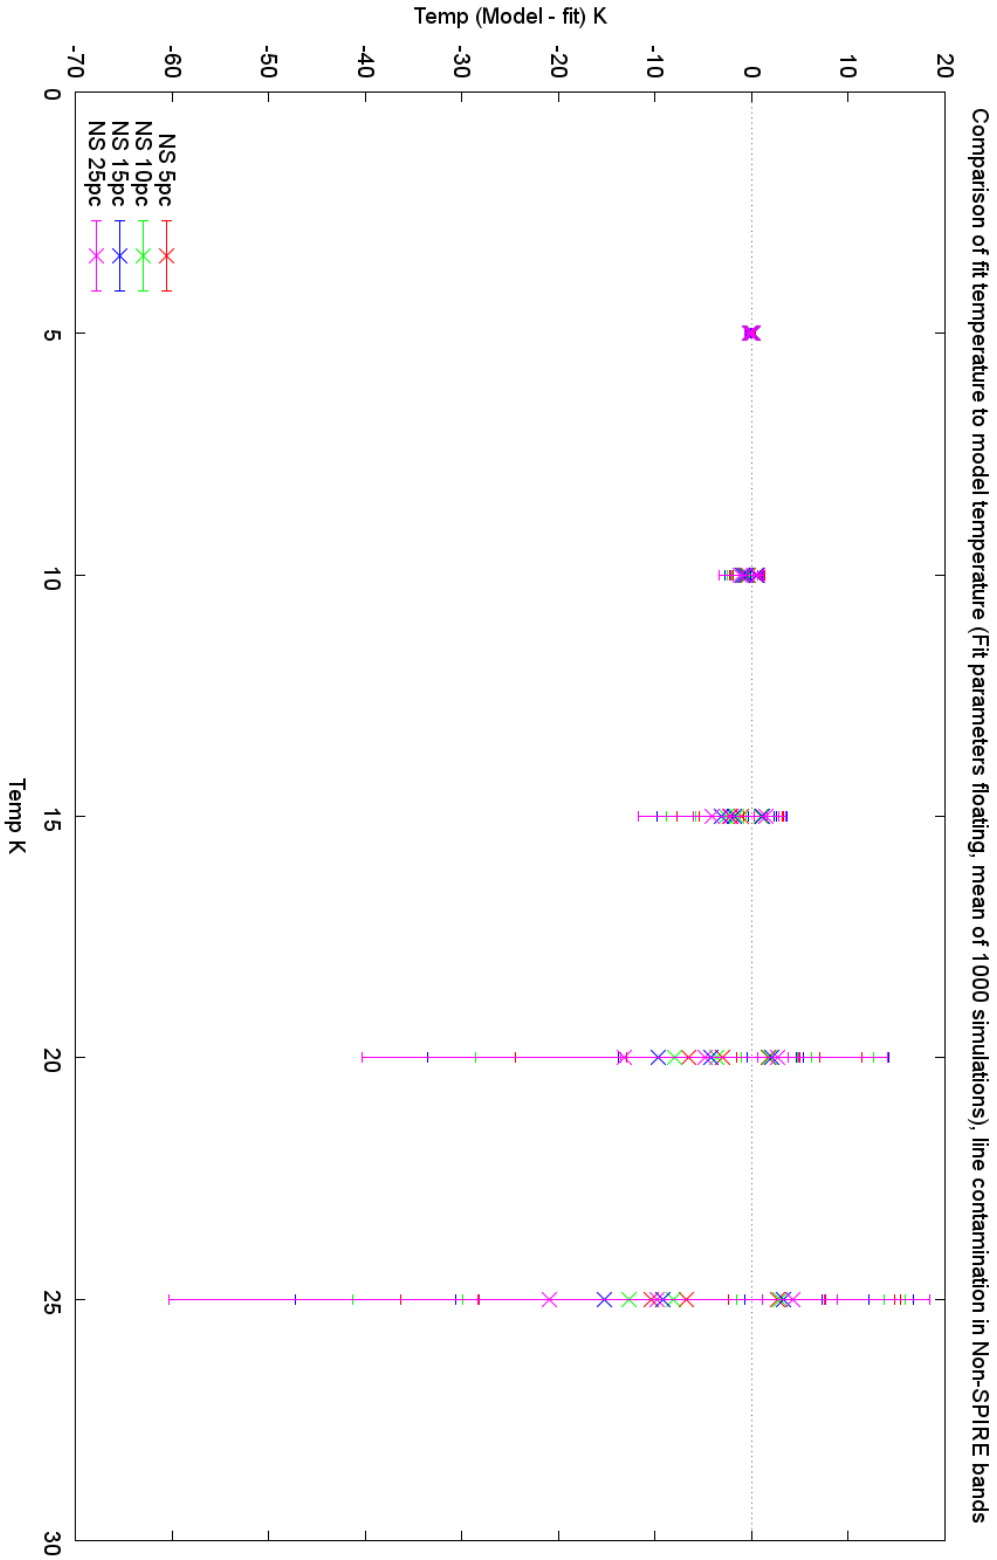

Figure 5.18: Line contamination with free fit parameters across non-SPIRE bands, effect on T_{rec}.

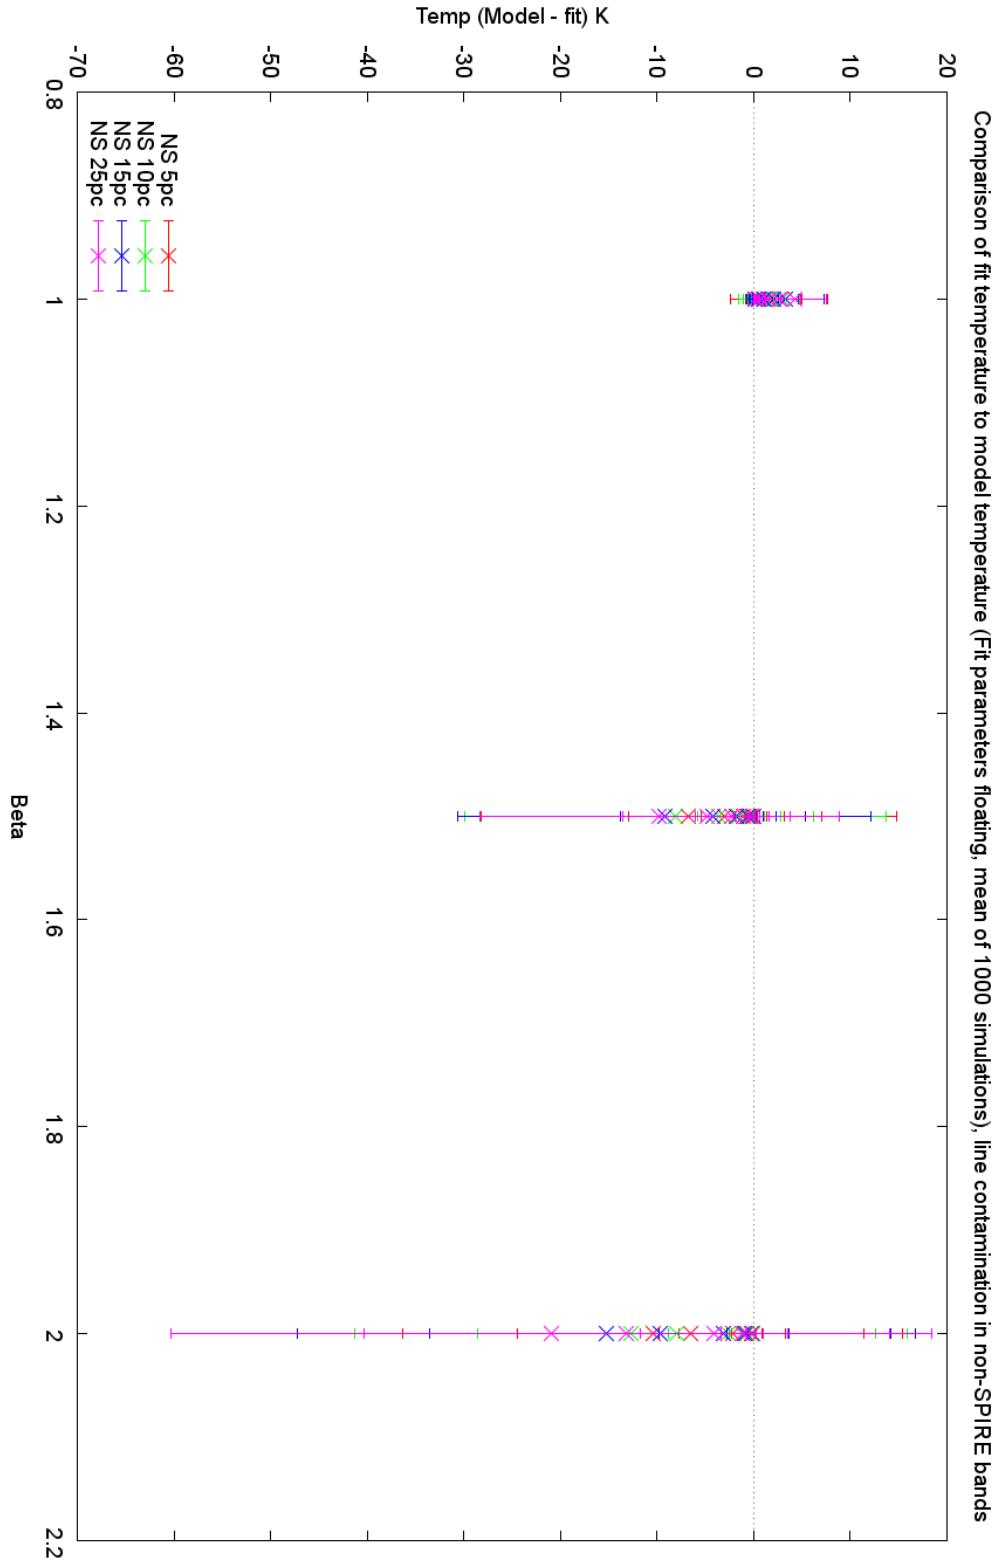

Figure 5.19: Line contamination with free fit parameters across non-SPIRE bands, effect on T_{rec}, grouped by β_{model}.

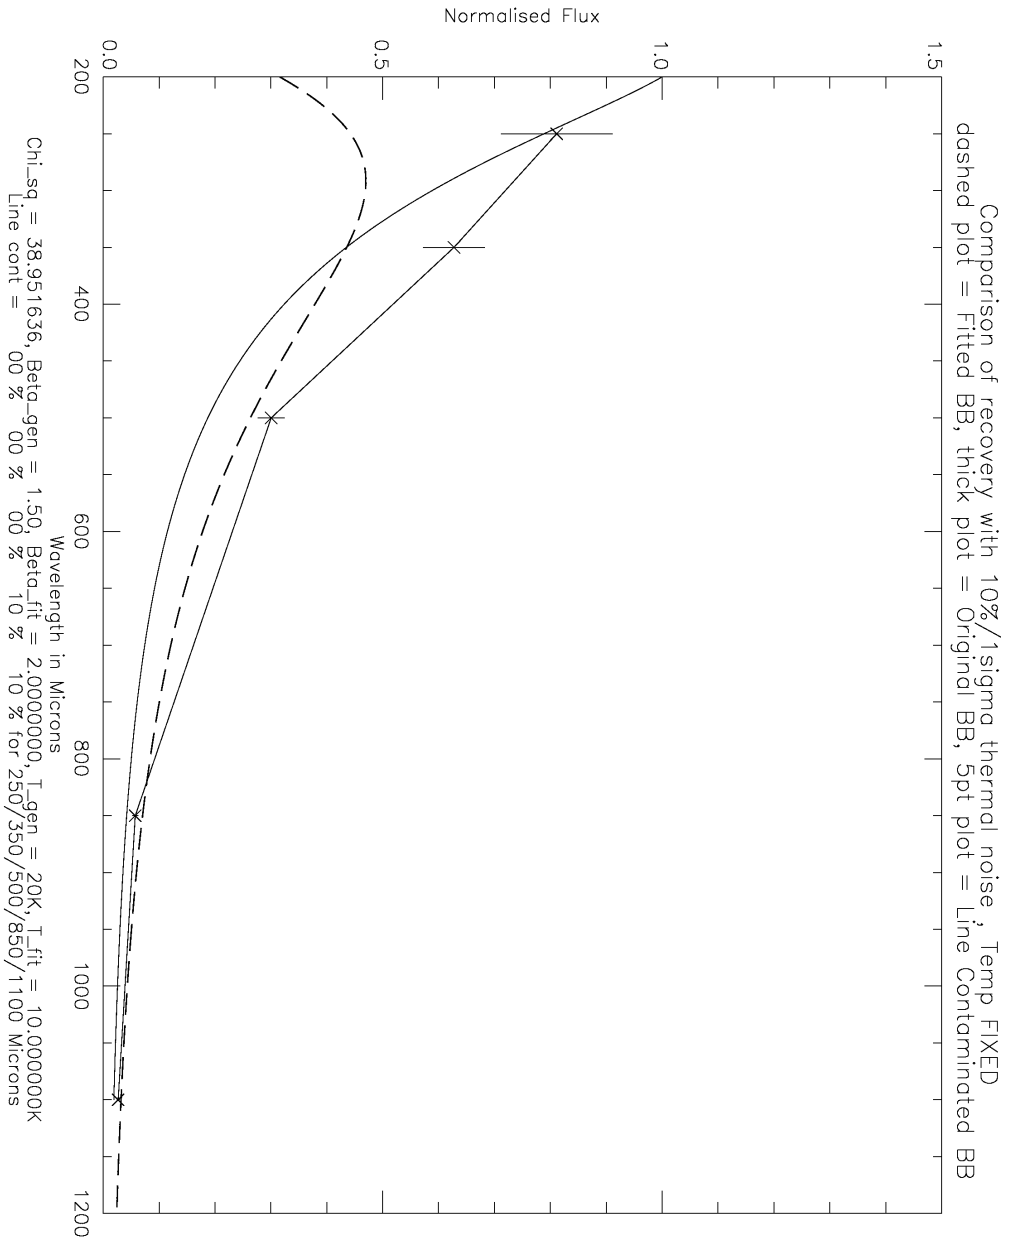

Figure 5.20: Sample from the simulations, showing the original greybody, the greybody with line contamination and the resultant fit.

Again, we can examine the effect on β. Comparing Figures 5.14 and 5.21, we see that there is a reversal. For lower values of β_{model}, the non-SPIRE line-contaminated fits show a smaller range of variation. This is inverted for the higher values of β_{model}. Overall, non-SPIRE line contamination gives a lower value of β_{rec} from the fit. Figures 5.15 and 5.22 show clearly that line contamination in the non-SPIRE bands yield a lower value for β_{rec}.

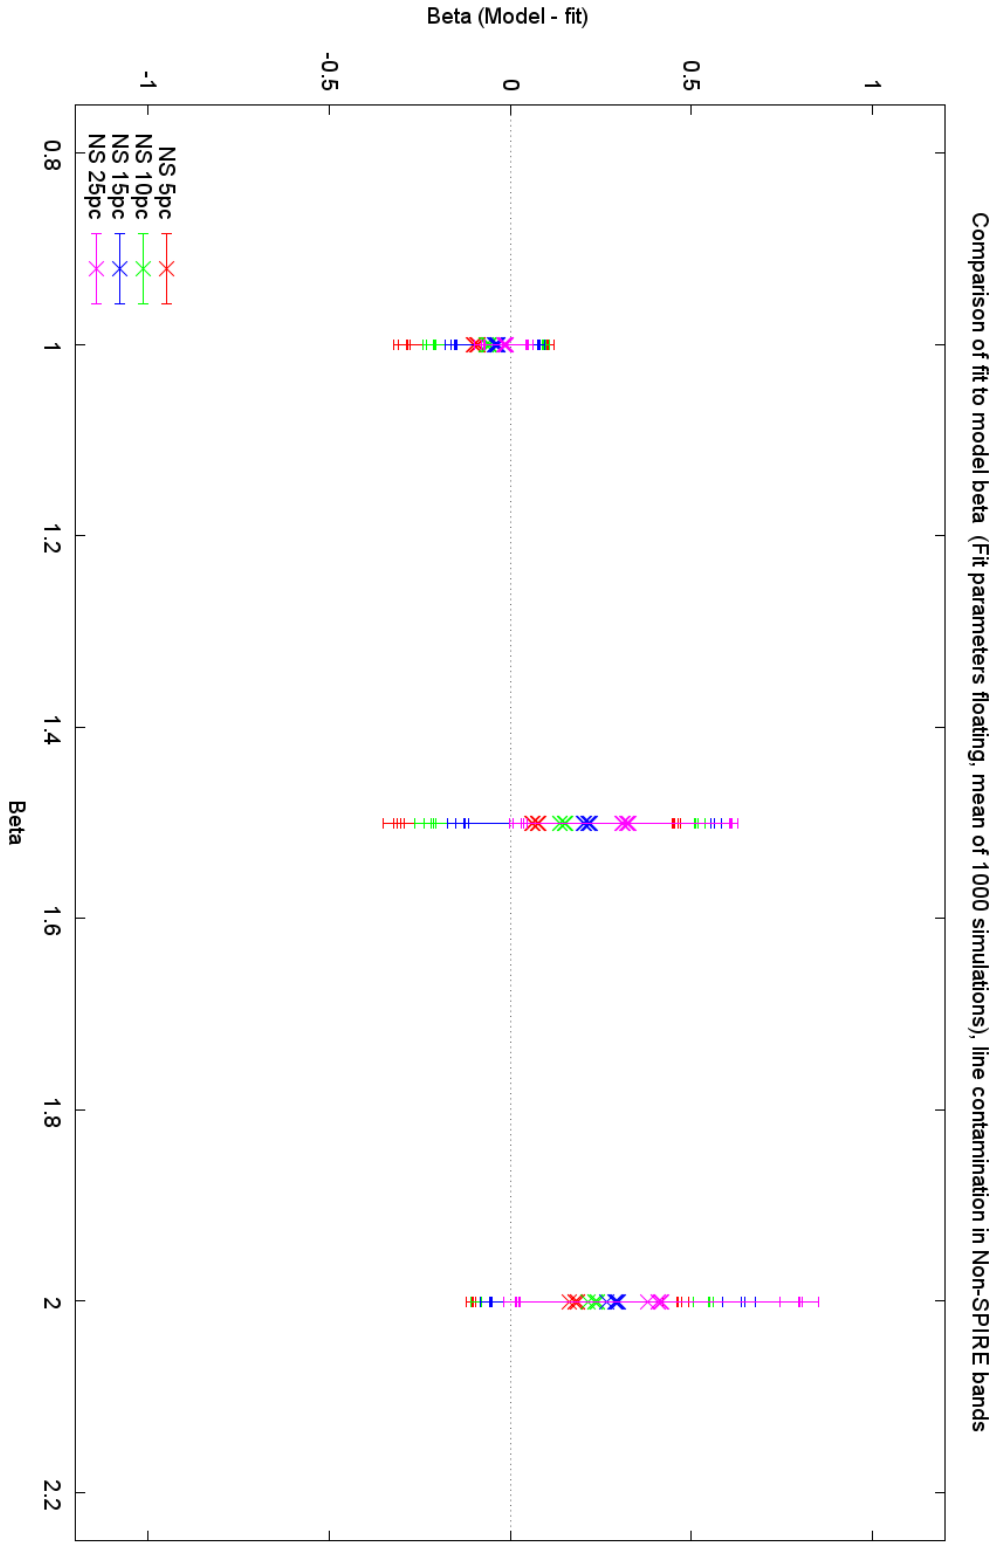

Figure 5.21: Line contamination with free fit parameters across non-SPIRE bands, effect on β_{rec}.

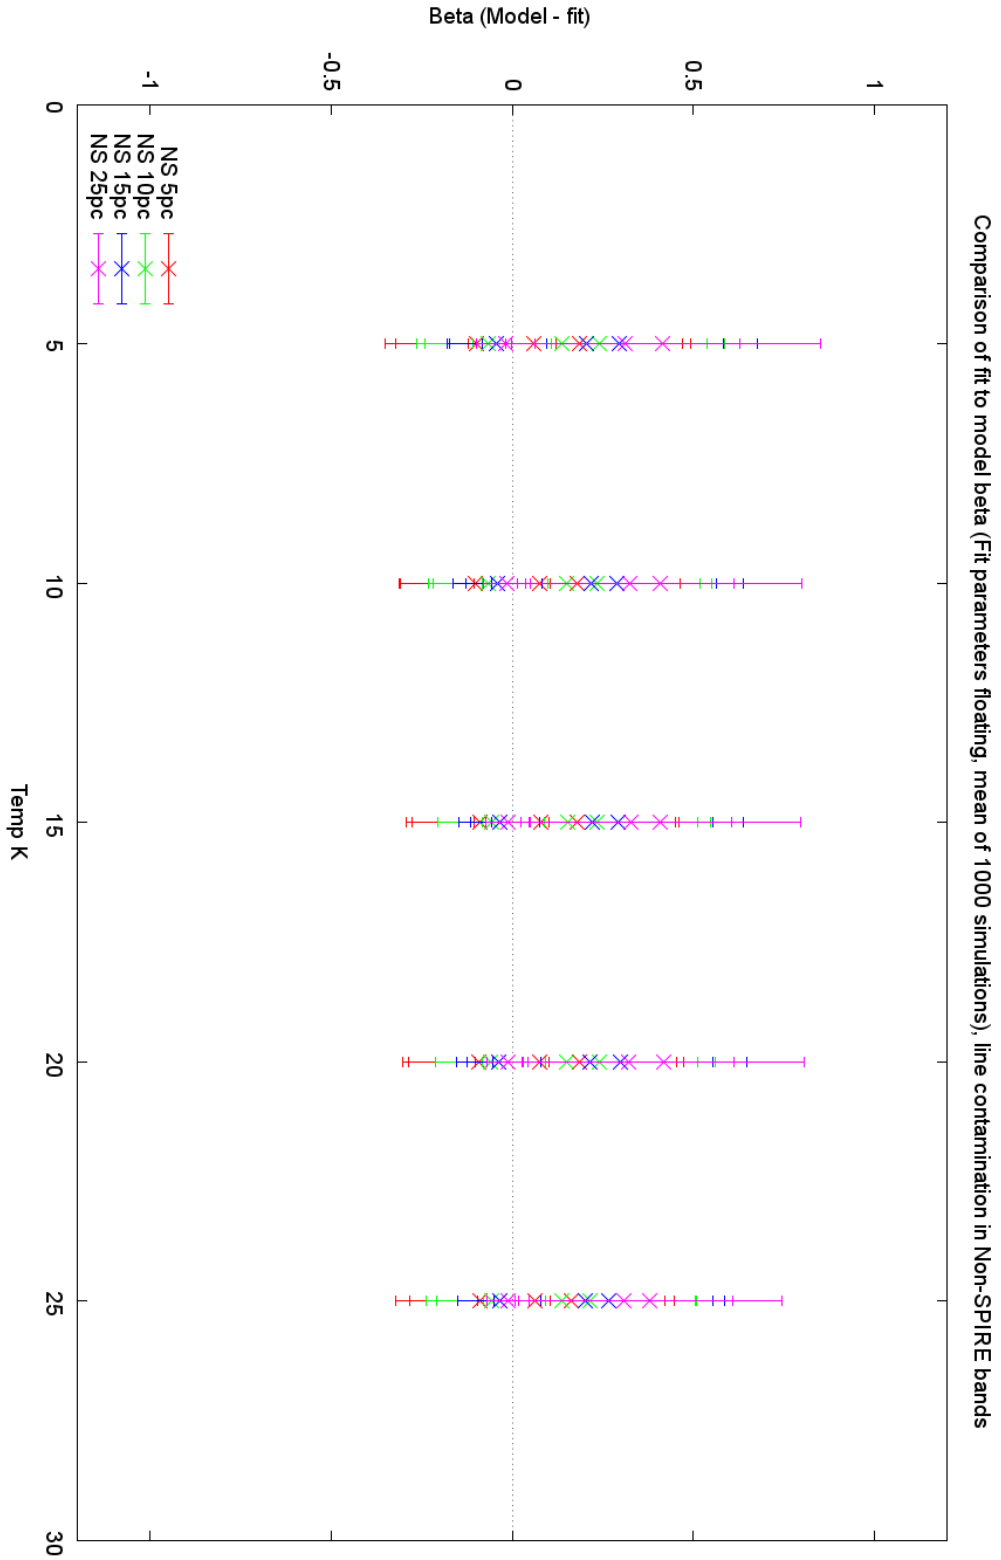

Figure 5.22: Line contamination with free fit parameters across non-SPIRE bands, effect on β_{rec}, grouped by T_{model}.

5.6.1.3 Free fit case - a closer examination

The following contour plots (Figures 5.23 to 5.34) give a clearer view of the effects of global line contamination in the free-fit case. Here we define *global* contamination as applying to all three SPIRE bands (global SPIRE contamination) or both non-SPIRE bands (global non-SPIRE contamination) simultaneously.

For global SPIRE contamination, we examine the effect of increasing line contamination on β_{rec} recovery (Figures 5.23 and 5.24) and temperature recovery (Figures 5.25 and 5.26) in closer detail. We then examine the effect on recovery on non-SPIRE waveband line contamination (Figures 5.27 to 5.30). Figures 5.31 to 5.34 use χ as a metric for quality of recovery, where we use the metric to give a measure of how well the overall fit is. In this case, we calculate χ as the difference of the squares between the line-contaminated greybody and the greybody we derive from the fit coefficients (β_{model} and greybody temperature T_{model}).

Figures 5.23 and 5.24 show how the variation of β_{rec} is affected for increasing line contamination in the SPIRE bands. As the line contamination percentage increases, we see an overall tendency for β_{rec} to increase. We see this as the negative range of the chart increase for increasing percentage. We also see a slope to the graphs in all cases. Although there is a tendency for β_{rec} to increase for higher percentage line contamination, for higher β_{model} greybodies this effect is reduced. So where β_{model} is already 2, we would expect that the fitted β_{rec} cannot increase anymore as this is a fixed limit.

What we see in the data is that for the lower β_{model} values in our simulations β_{rec} is driven high by about $0.3 - 0.4$ and the greybody temperature reduced. For the higher limit of β_{model}, we see β_{rec} driven low, by about 0.1 to 0.01 as the

temperature increases, and the temperature driven high. For 5K greybodies, the temperature increase is about 0.1K and for 25K greybodies it is about 6K. As β_{rec} is constrained by the fit, what this means is that as the fit routine reaches the fixed upper limit the only fully-free parameter left for the fit is temperature.

We see this in Figures 5.25 and 5.26, where the higher temperature greybodies show this change in behaviour more clearly. For greybodies with an initial low temperature, the slope of Figure 5.25 and 5.26 for each value of β_{model} is small. For 5K we see the temperature driven low for low β_{model}, and driven high for high β_{model}, by about 0.1K for 5% line contamination and by about 0.3K for 25% line contamination. Where the initial greybody temperature is at the higher end, we see these results magnified. For 5% line contamination, we see for 25K greybodies a variation of 3K/6K for low/high β_{model}, and for 25% this becomes 5K/6K.

Figures 5.27 to 5.30 show the effect of non-SPIRE band line contamination. Broadly, the effect is similar. For low initial greybody temperatures, we see the temperature driven low for low β_{model}, and driven high for high β_{model} and this effect is again magnified by higher greybody initial temperatures. For the 5K case we see a similar range as for the SPIRE band contamination, but for the 25K case the change is 3K/10K for 5% contamination and 5K/20K for 25%. The behaviour for the central values for β_{model} is not similar to the SPIRE band contamination.

We recall for SPIRE band contamination in the central β_{model} region we find that for lower line percentages $(5\%, 10\%)$ we see that for lower temperatures $\beta_{rec} \geq \beta_{model}$ and $T_{rec} \leq T_{model}$. As T_{model} increases we see an inversion of this behaviour with T_{rec} becoming $\geq T_{model}$. However β_{rec} stays approximately constant across the temperature range. The changeover temperature for this inversion behaviour is $\approx 15K_{5\%LC}, 25K_{10\%LC}$. As line contamination increases we lose this

inversion and the degree of deviation between model and recovered temperature also increases. Above 15% line contamination we see consistent behaviour across the temperature range, in that $\beta_{rec} \geq \beta_{model}$ and $T_{rec} \leq T_{model}$. For the $25K$ simulations the change in temperature ΔT is $+2.75K/5\%$, which becomes $+0.6K/10\%$, $-2.25K/15\%$,$-3.7K/25\%$. Hence for low line contamination percentages, the behaviour of the mid-range β simulations changes as the temperature increases. The inversion temperature is roughly that where the peak moves out of the waveband we are considering. As line contamination increases, the shape of the SED is distorted. These distortions are such that for higher percentages the fit behaviour changes. As the initial greybody temperature increases, we find the fit temperature deviates from the initial temperature by a larger amount. As the peak of the greybody emission moves out of the spectral range, the deviations increase sharply. This is expected as the peak is an extra feature that is useful in constraining our fit.

For non-SPIRE line contamination in the central β values $T_{rec} \geq T_{model}$ and β_{rec} stays approximately constant, although $\beta_{rec} \leq \beta_{model}$ for a given line contamination percentage. As the percentage increases, β_{rec} decreases from ≈ 1.4 at 5% contamination to ≈ 1.2 at 25%. The amount of increase in T_{rec} in the central β_{model} regime increases with both temperature and line contamination percentage - ranging from $0.04K$ at $5K/5\%$ up to $9.75K/25\%$. There is also no turnover, the behaviour is consistent across the range of line contamination percentages for non-SPIRE line contamination. We see that β_{rec} is consistently driven low across the board. For low β_{model}, the effect is to keep $\beta_{rec} \approx 1$, and for higher β_{model} it is driven low by about 0.2 (for 5%) and 0.4 (for 25%). For low β_{model} there is a slight increase as this is a lower limit, so yet again we see a slight β_{rec} increase and T_{rec} decrease, where the fit is again constrained by having only one fully-free

parameter in temperature, and one fit parameter constrained in one direction.

Overall, for SPIRE line contamination we see a tendency for β_{rec} to increase and the fit temperature to decrease for low β_{model}. At the higher β_{model} end this situation is reversed due to the limit. Central β_{model} values show a differing behaviour dependent on line contamination. For non-SPIRE contamination we see that β_{rec} decreases and the fit temperature increases, and here the lower fit limit causes the inversion in behaviour as the constraint, with no central β_{model} behavior deviations. The range of temperature deviation increases as the initial greybody temperature increases, and becomes significantly larger for the higher temperature range of the simulation. This suggests that as the greybody peak moves out of our spectral range, we see increasing issues with the fit.

Figures 5.31 and 5.32 show the quality of fit for SPIRE contamination, using a χ value as a metric by comparing the original spectrum to the resultant fit. For lower line contamination percentages, there is little to distinguish between any of the simulations, although as the percentage increases we see a slight fit degradation for higher β_{model}. This degradation also increases for higher temperatures, but with an unusual feature. There is a slight tendency for the fit quality to improve for lower β_{model} as the temperature increases.

Figures 5.33 to 5.34 show the quality of fit for non-SPIRE contamination. For lower percentages, it is similar to the SPIRE figure. As the percentage increases, we see the inversion we saw earlier. Here, the quality of fit is better for higher β_{model} and as the temperature increases the fit quality is largely stable. Overall, χ shows little variation across the ranges shown.

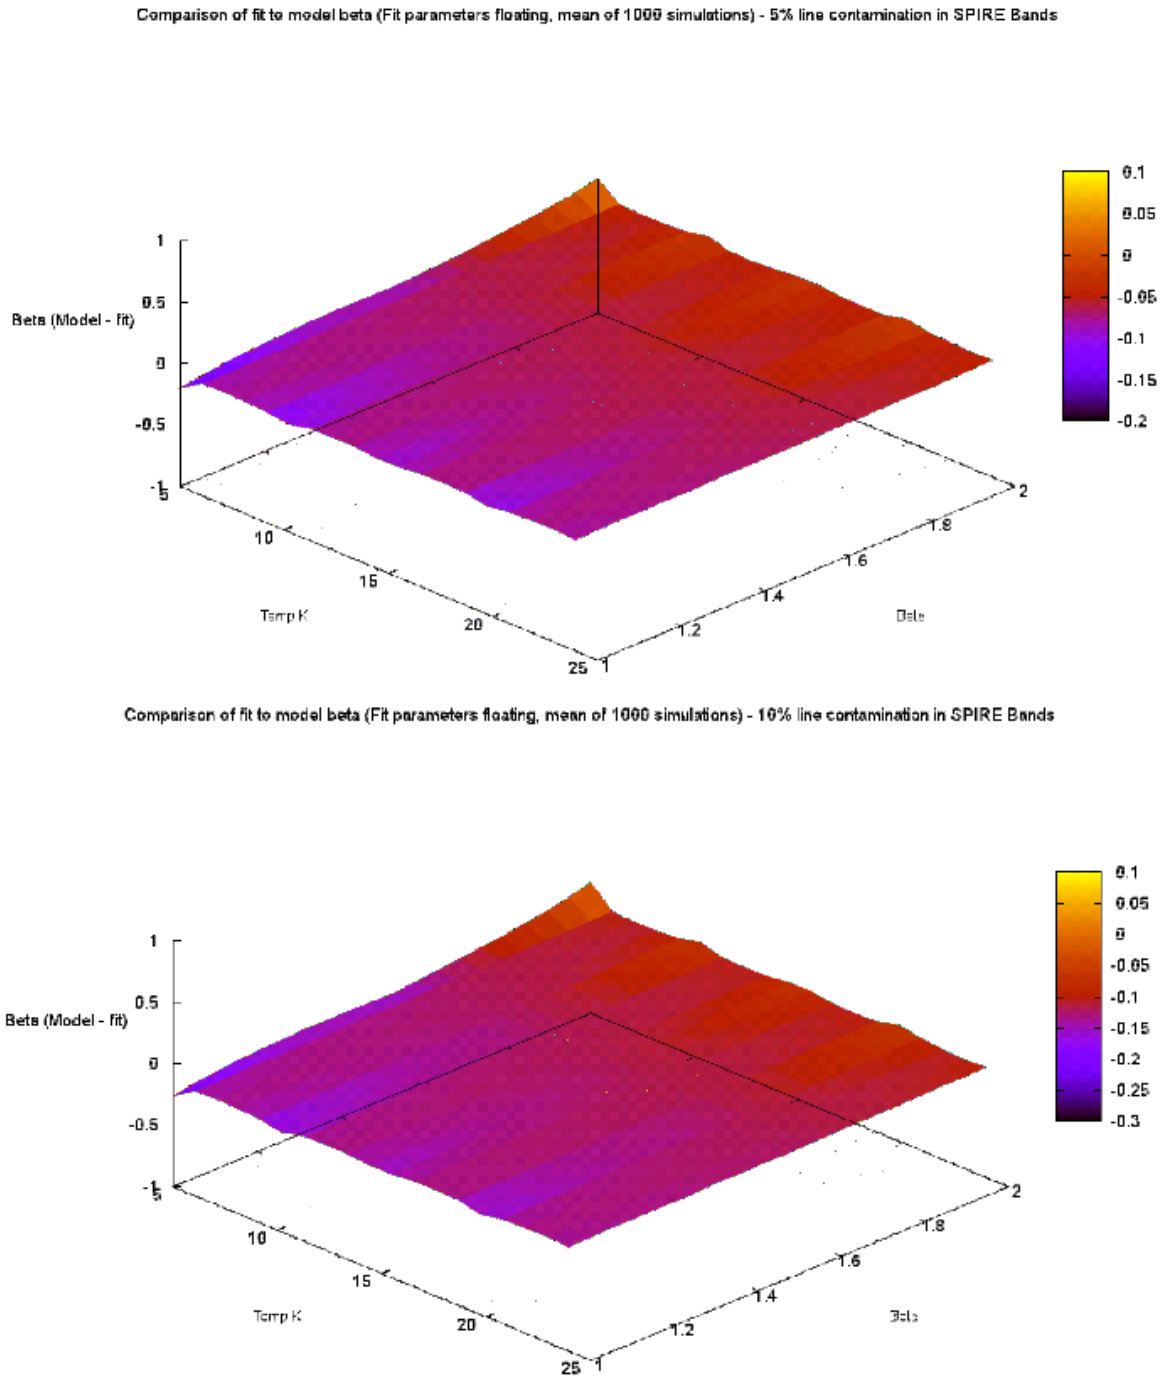

Figure 5.23: Line contamination with free fit parameters across SPIRE bands, effect on β_{rec}, 5% and 10% line contamination percentage.

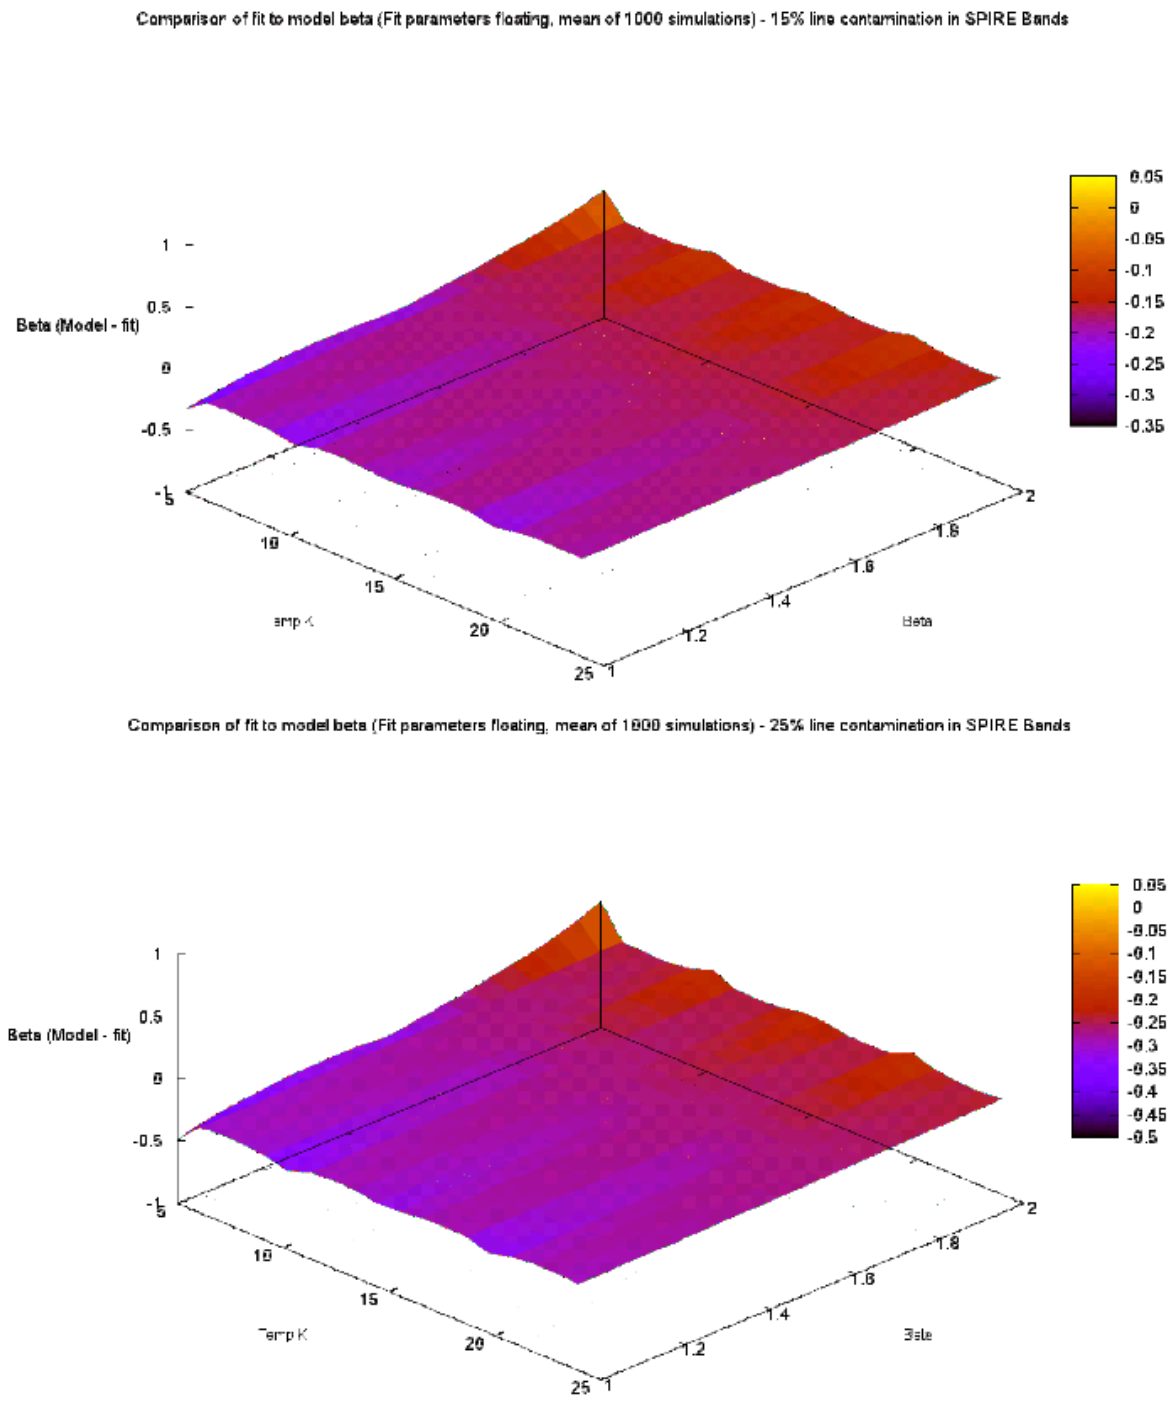

Figure 5.24: Line contamination with free fit parameters across SPIRE bands, effect on β_{rec}, 15% and 25% line contamination percentage.

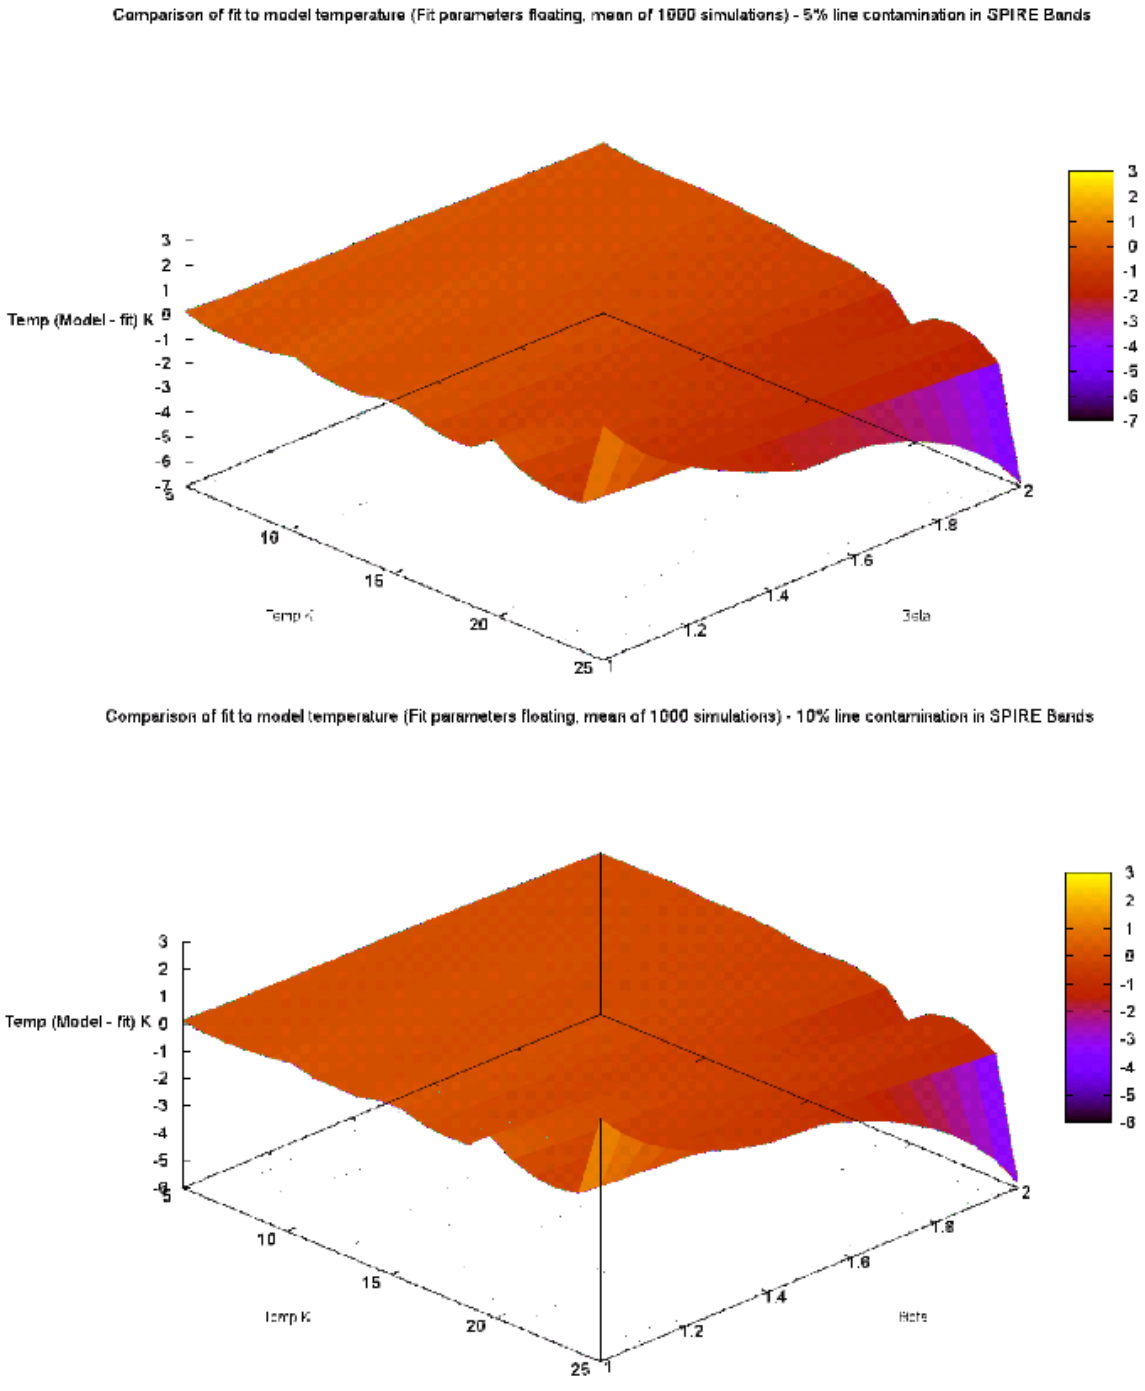

Figure 5.25: Line contamination with free fit parameters across SPIRE bands, effect on T_{rec}, 5% and 10% line contamination percentage.

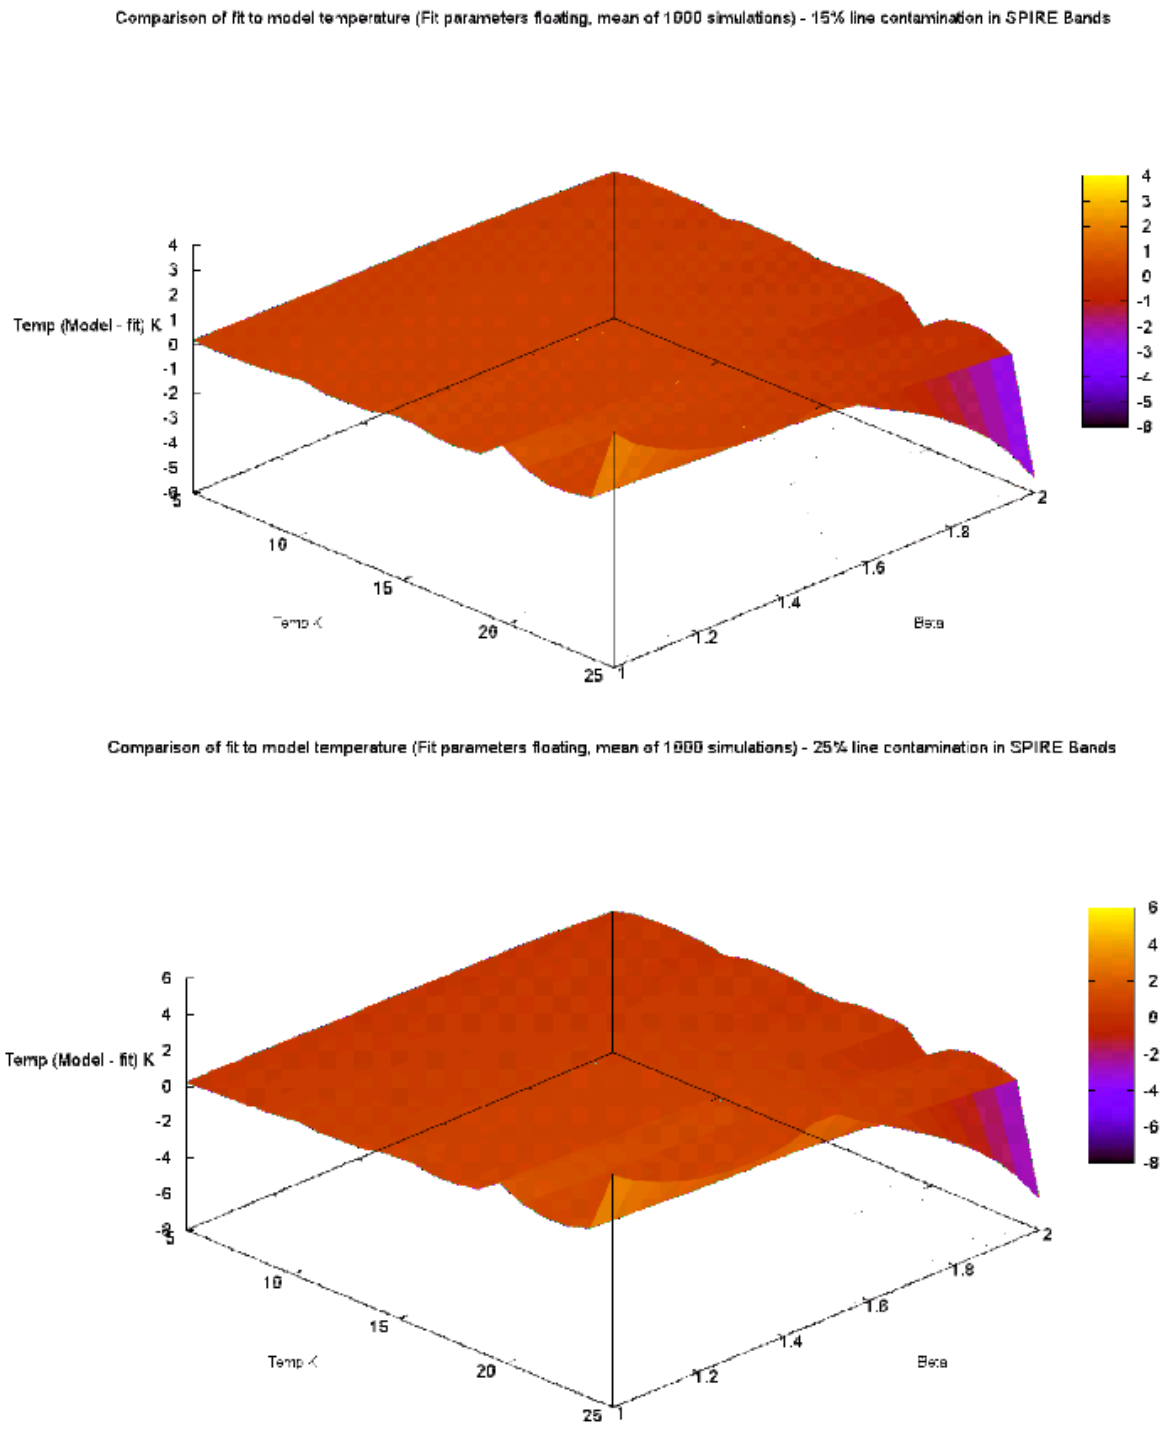

Figure 5.26: Line contamination with free fit parameters across SPIRE bands, effect on T_{rec}, 15% and 25% line contamination percentage.

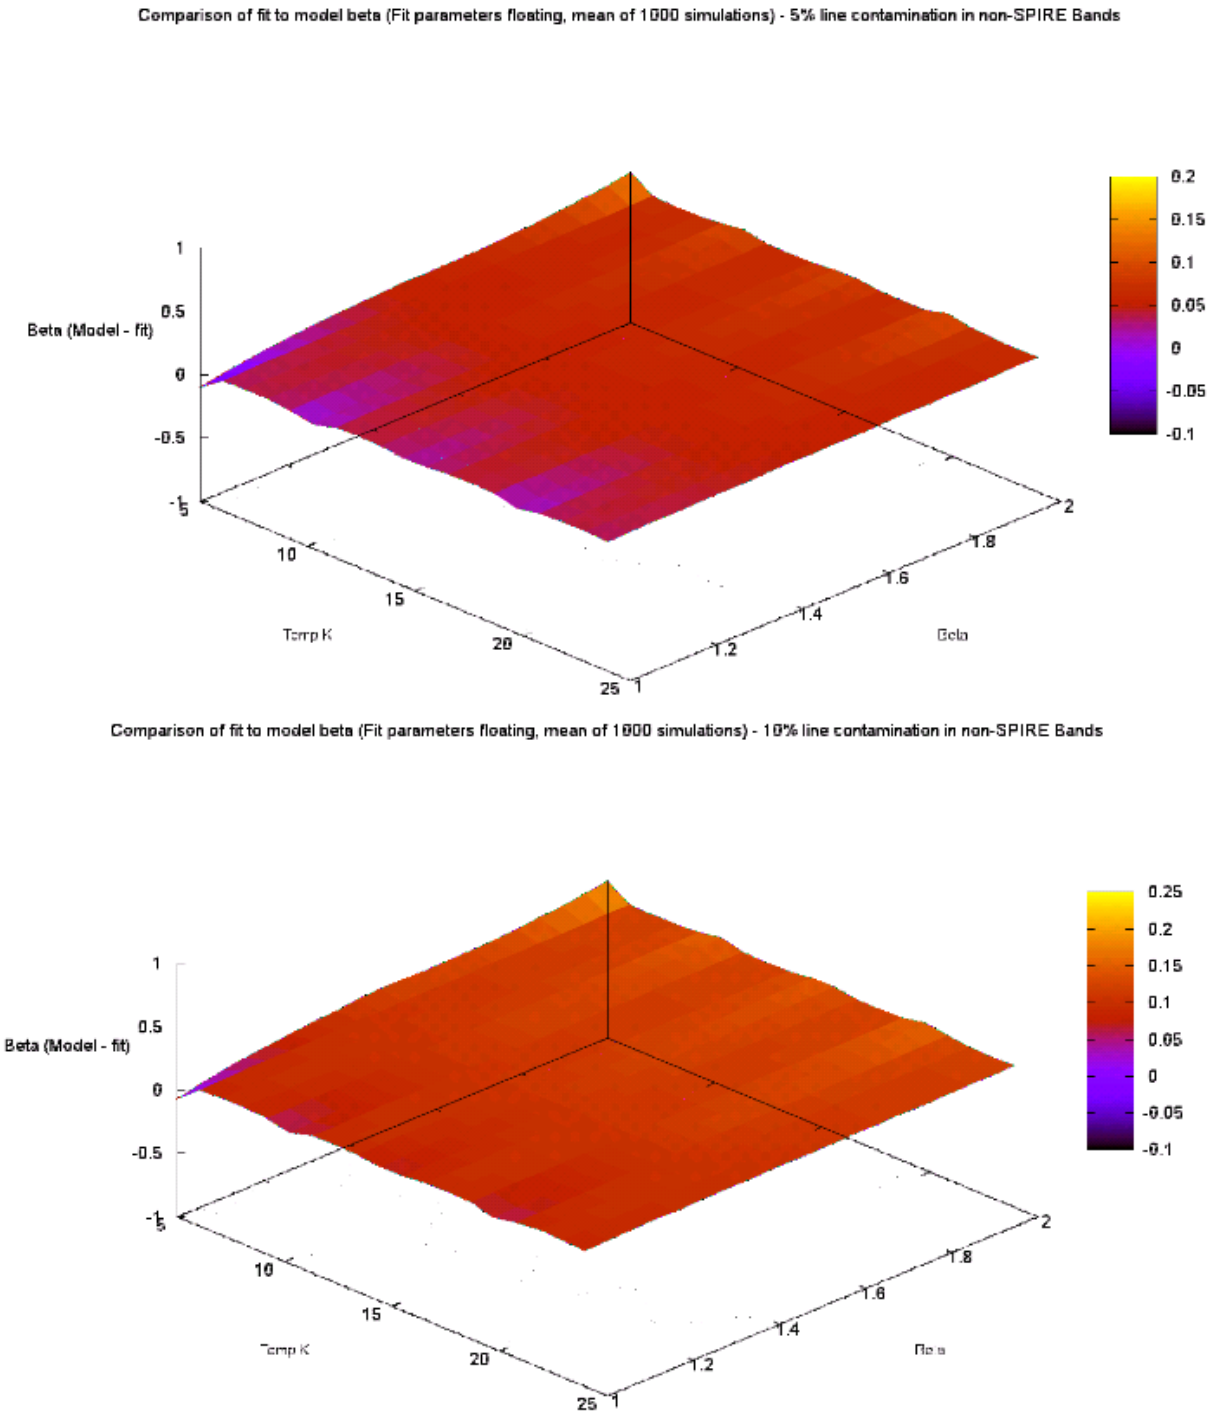

Figure 5.27: Line contamination with free fit parameters across non-SPIRE bands, effect on β_{rec}, 5% and 10% line contamination percentage.

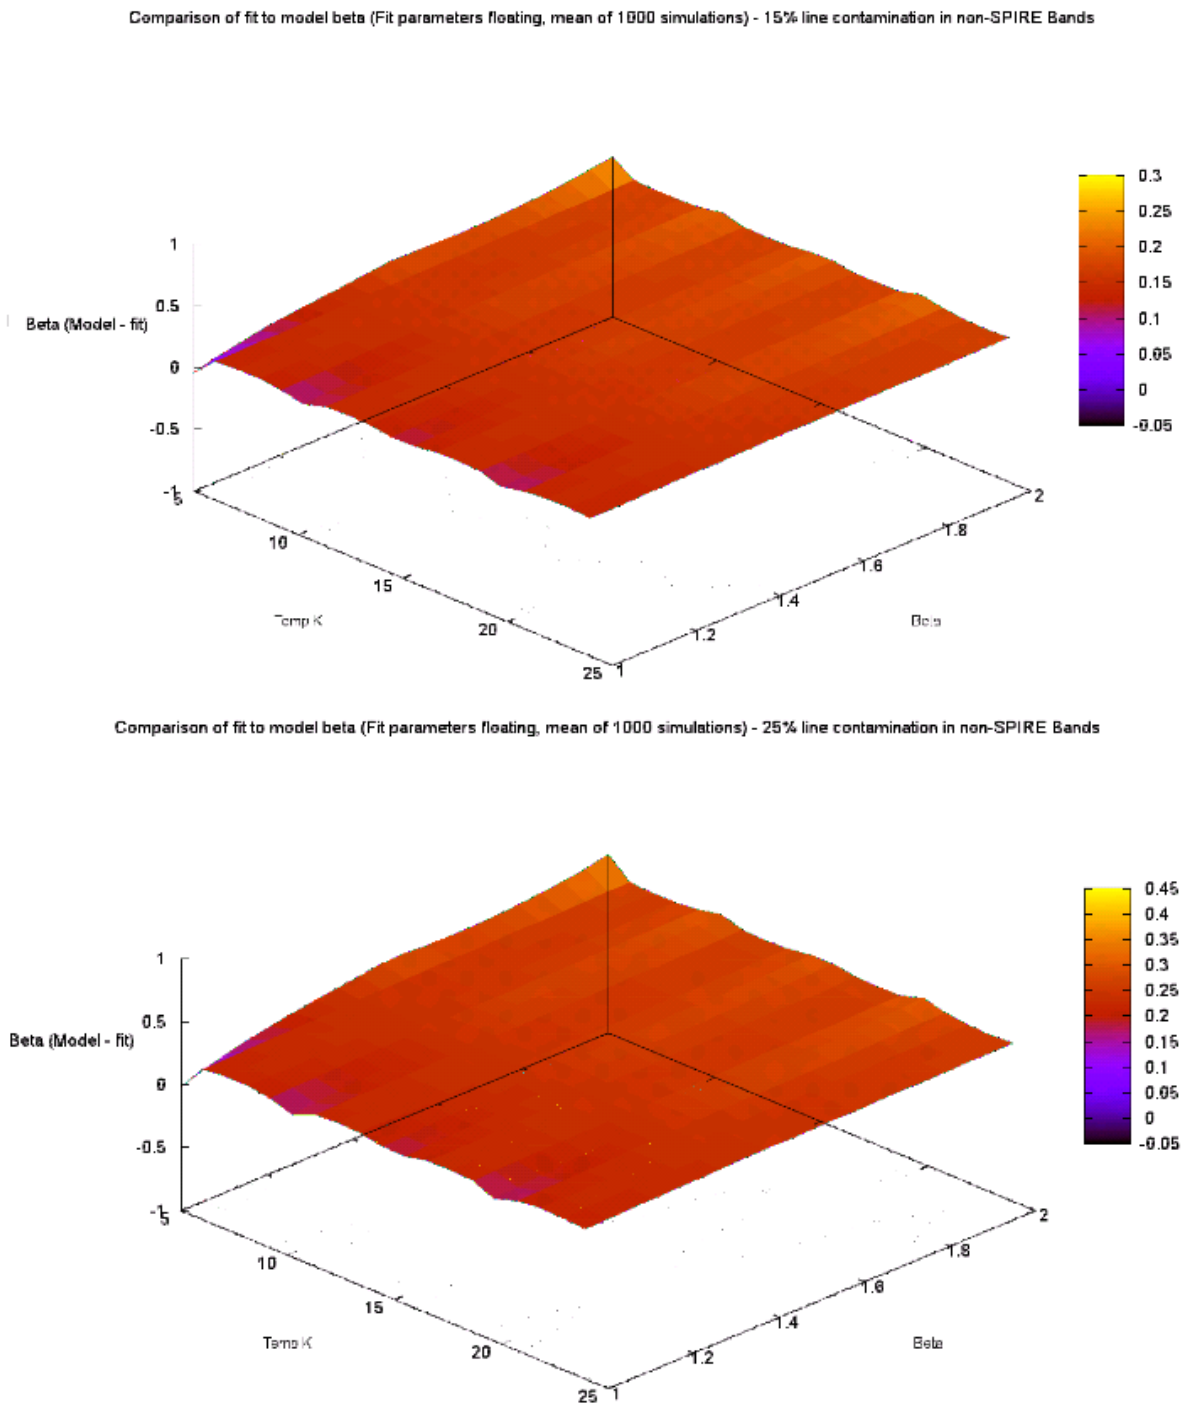

Figure 5.28: Line contamination with free fit parameters across non-SPIRE bands, effect on β_{rec}, 15% and 25% line contamination percentage.

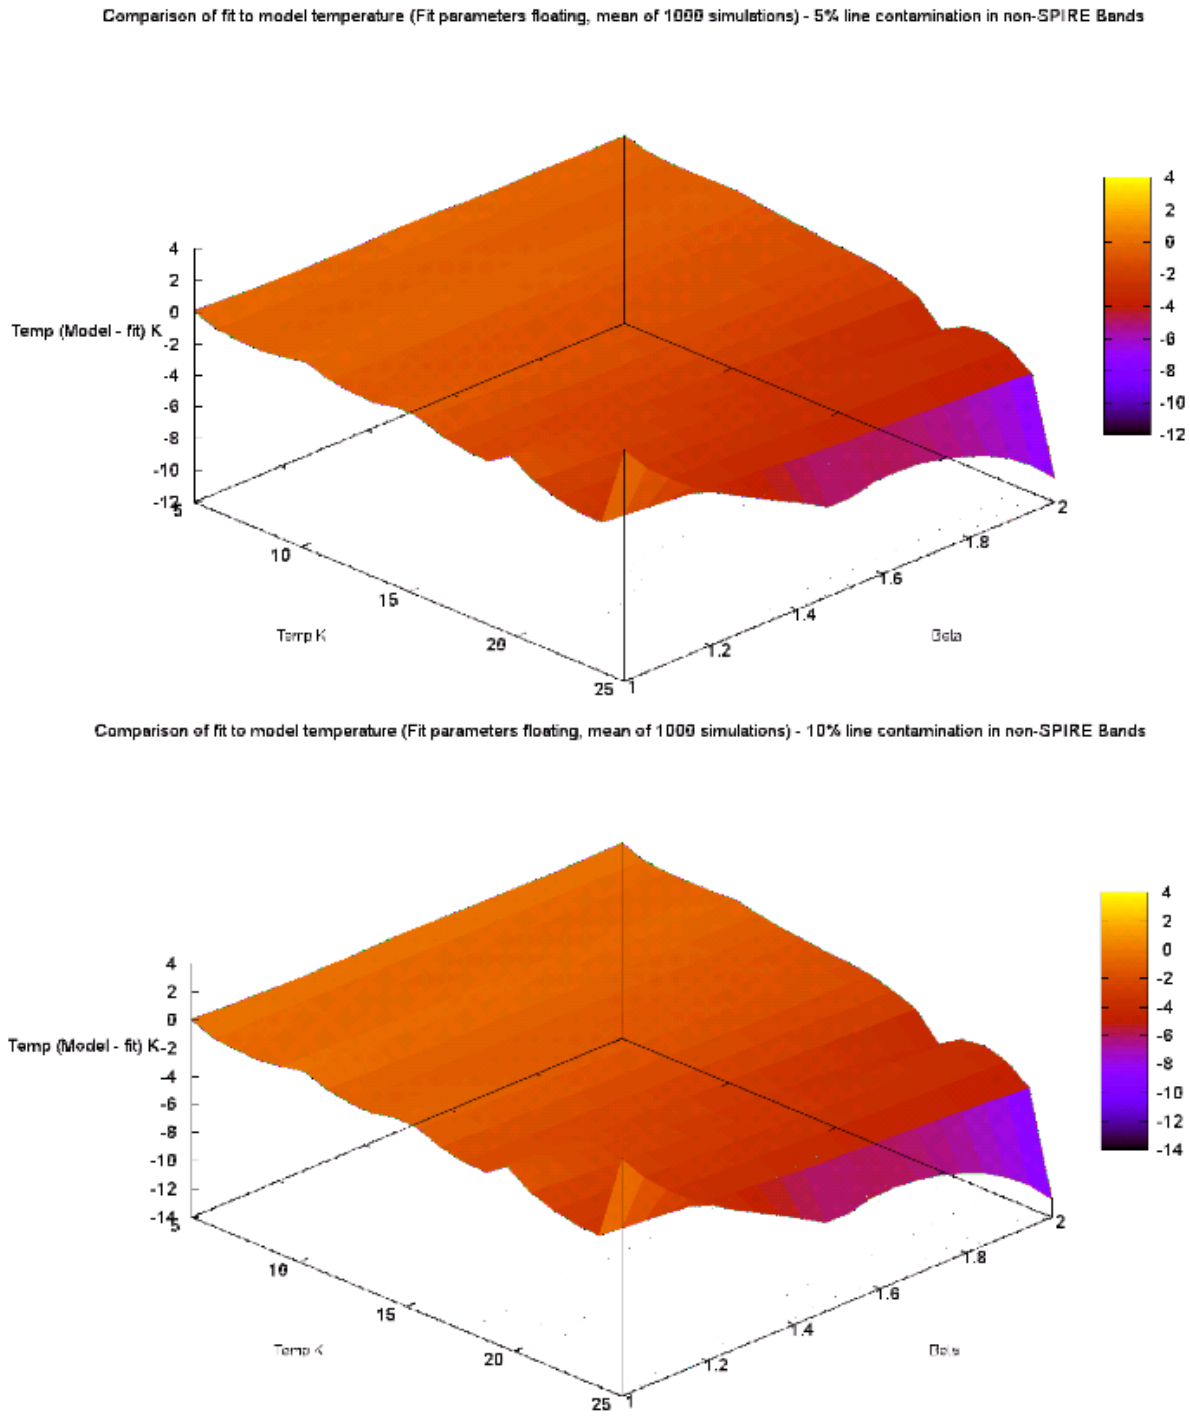

Figure 5.29: Line contamination with free fit parameters across non-SPIRE bands, effect on T_{rec}, 5% and 10% line contamination percentage.

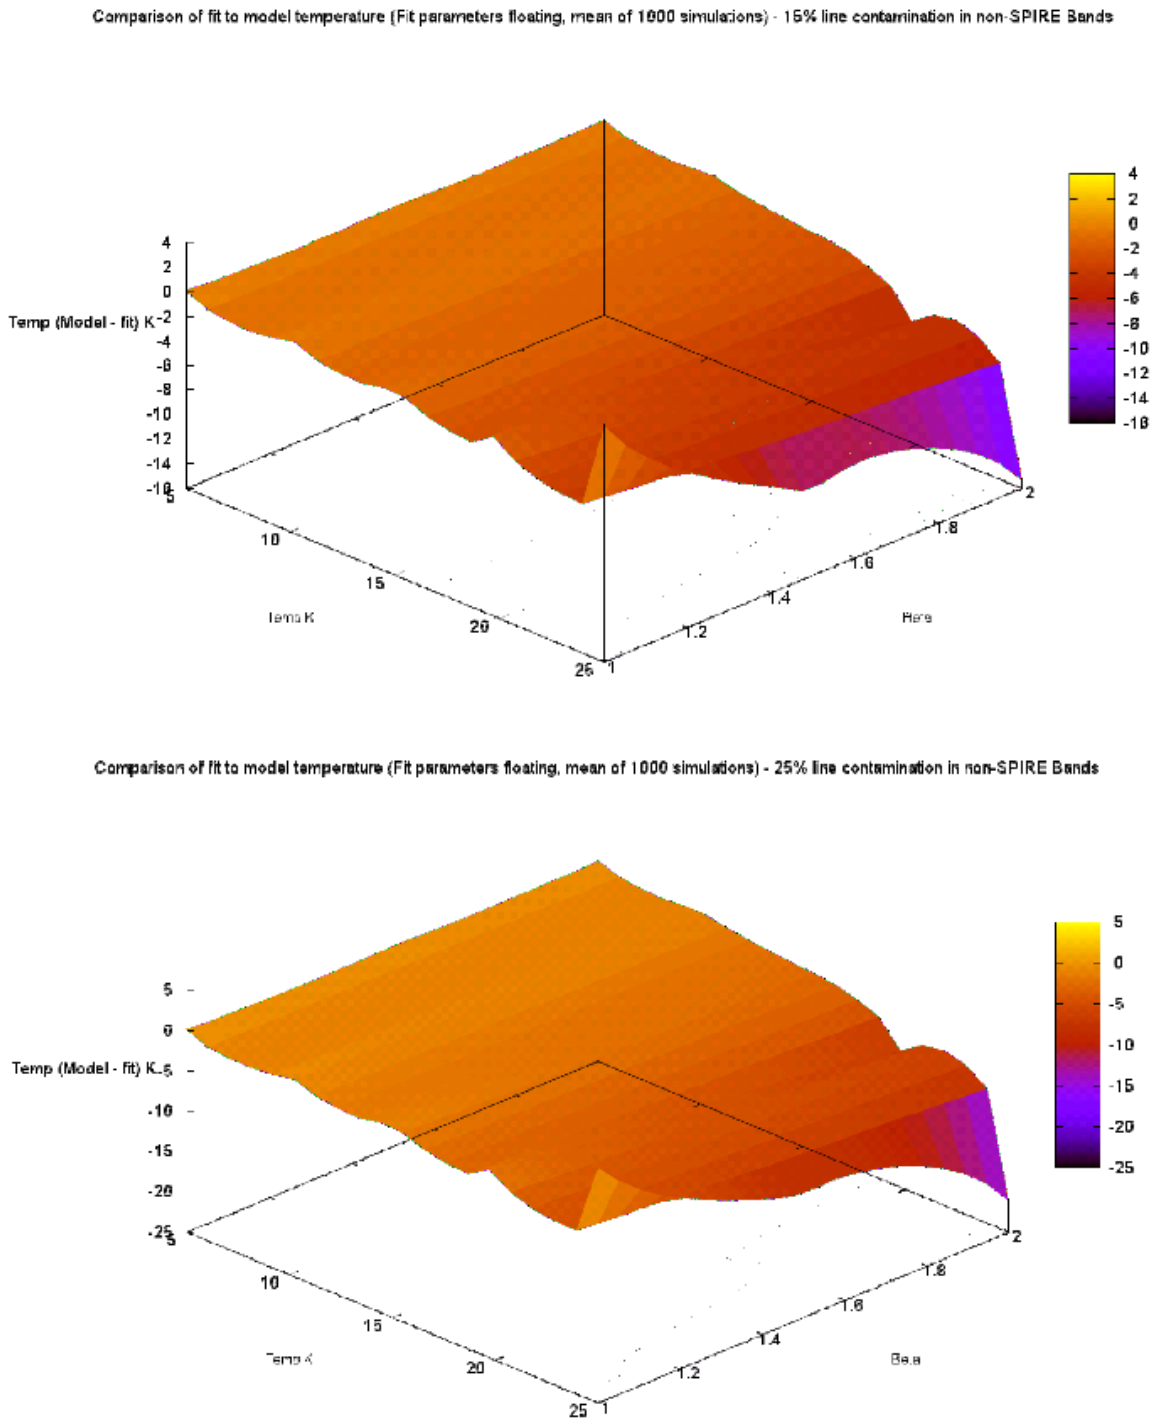

Figure 5.30: Line contamination with free fit parameters across non-SPIRE bands, effect on T_{rec}, 15% amd 25% line contamination percentage.

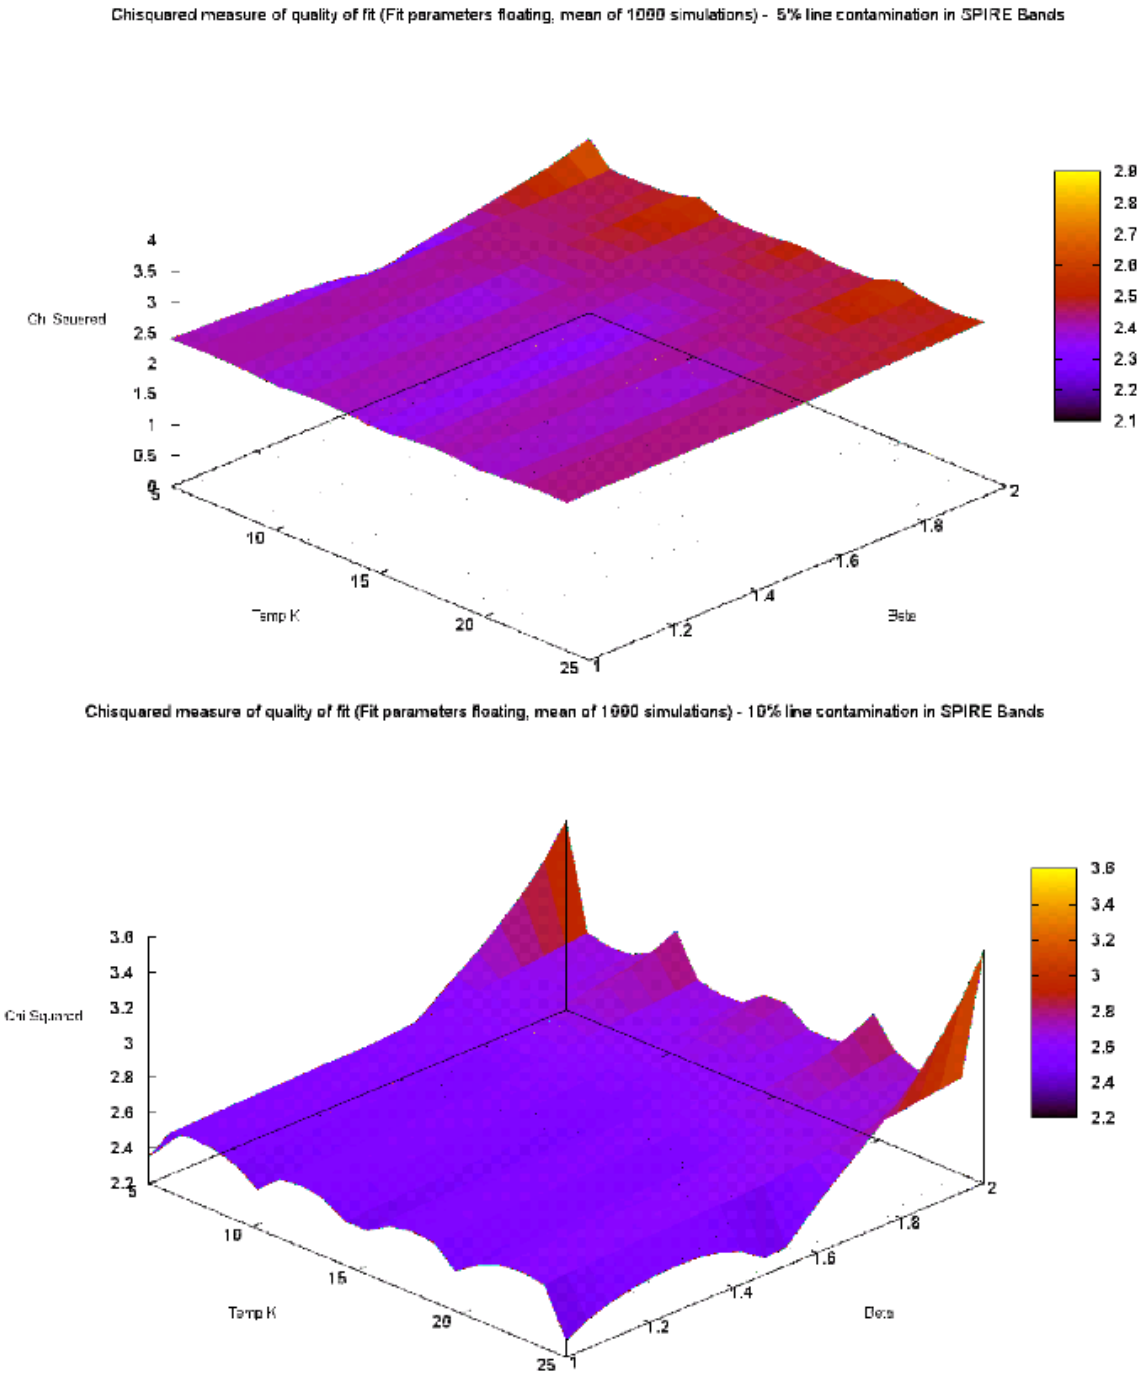

Figure 5.31: Line contamination with free fit parameters across SPIRE bands χ, 5% and 10% line contamination percentage.

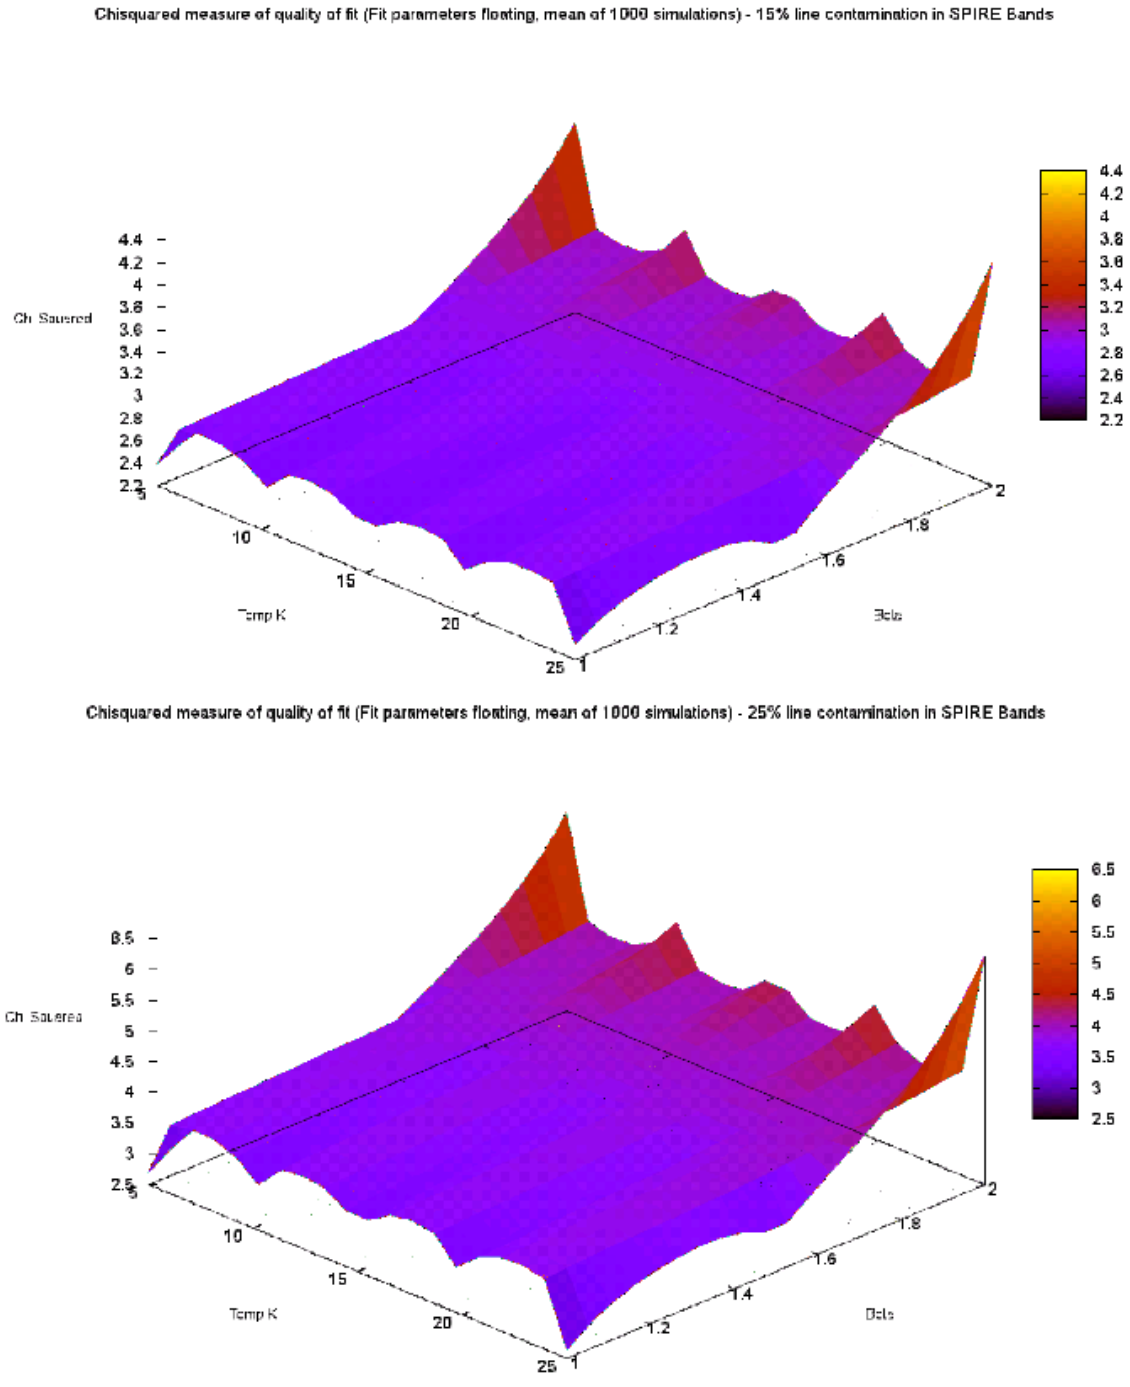

Figure 5.32: Line contamination with free fit parameters across SPIRE bands χ, 15% and 25% line contamination percentage.

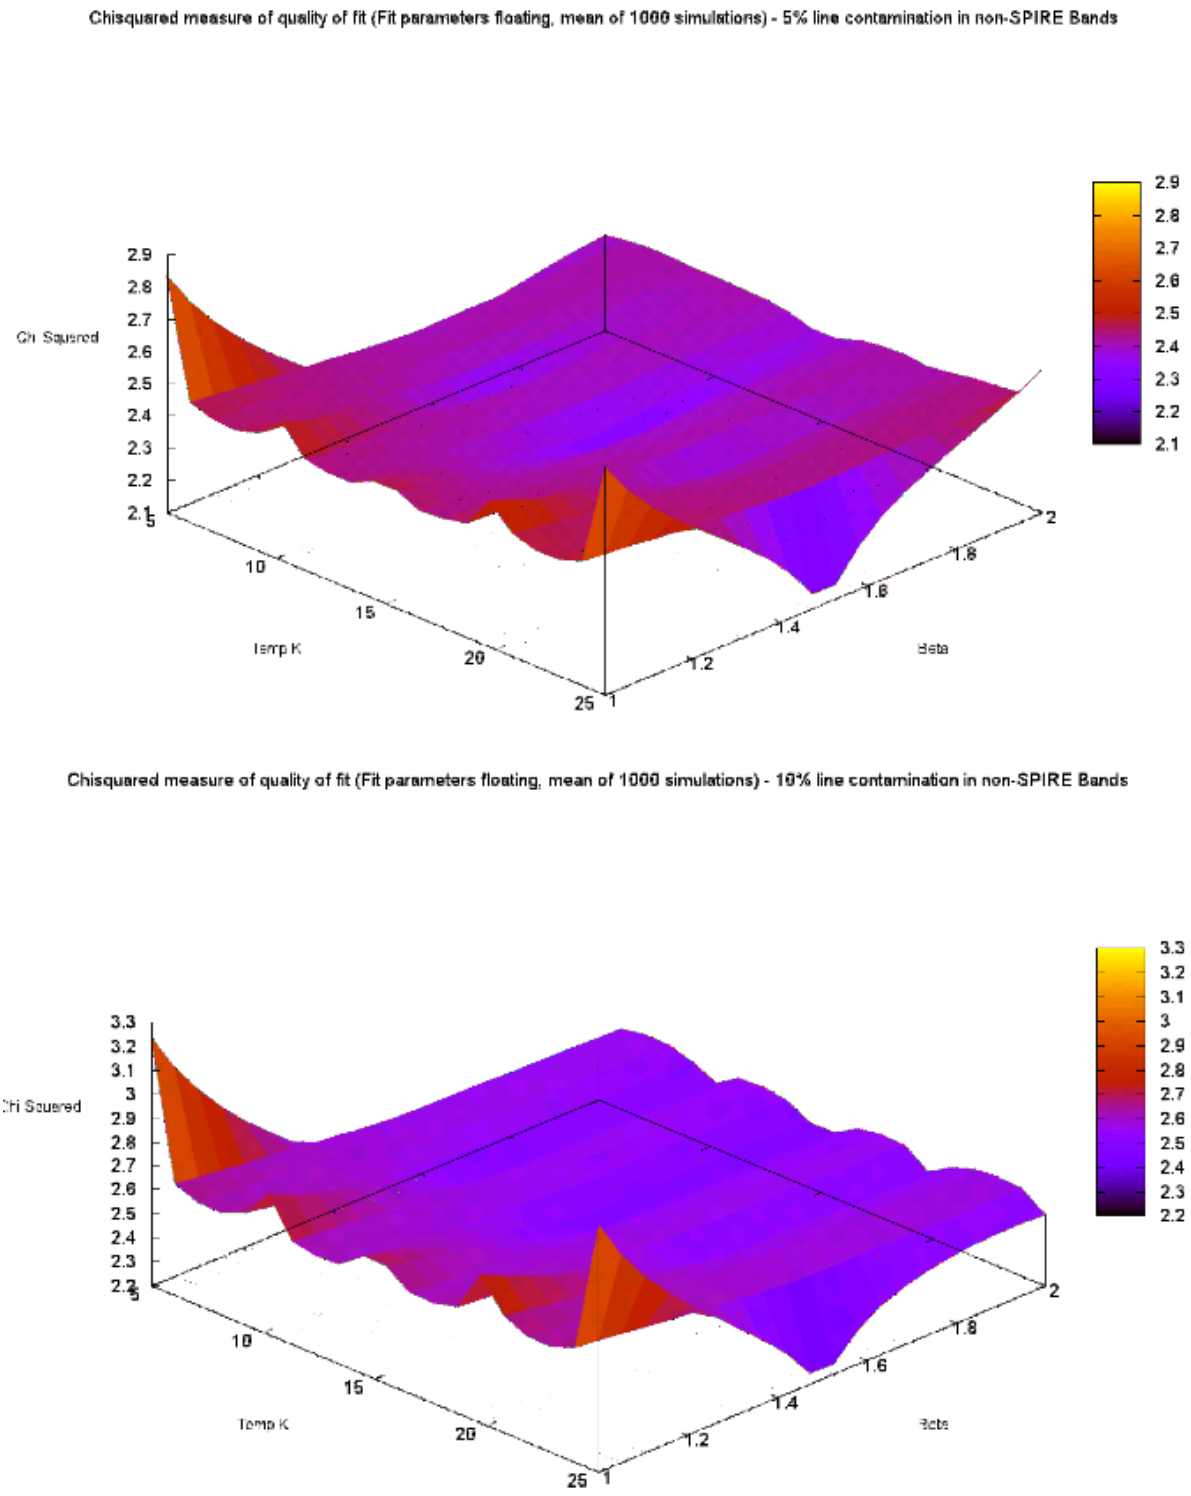

Figure 5.33: Line contamination with free fit parameters across non-SPIRE bands χ, 5% and 10% line contamination percentage.

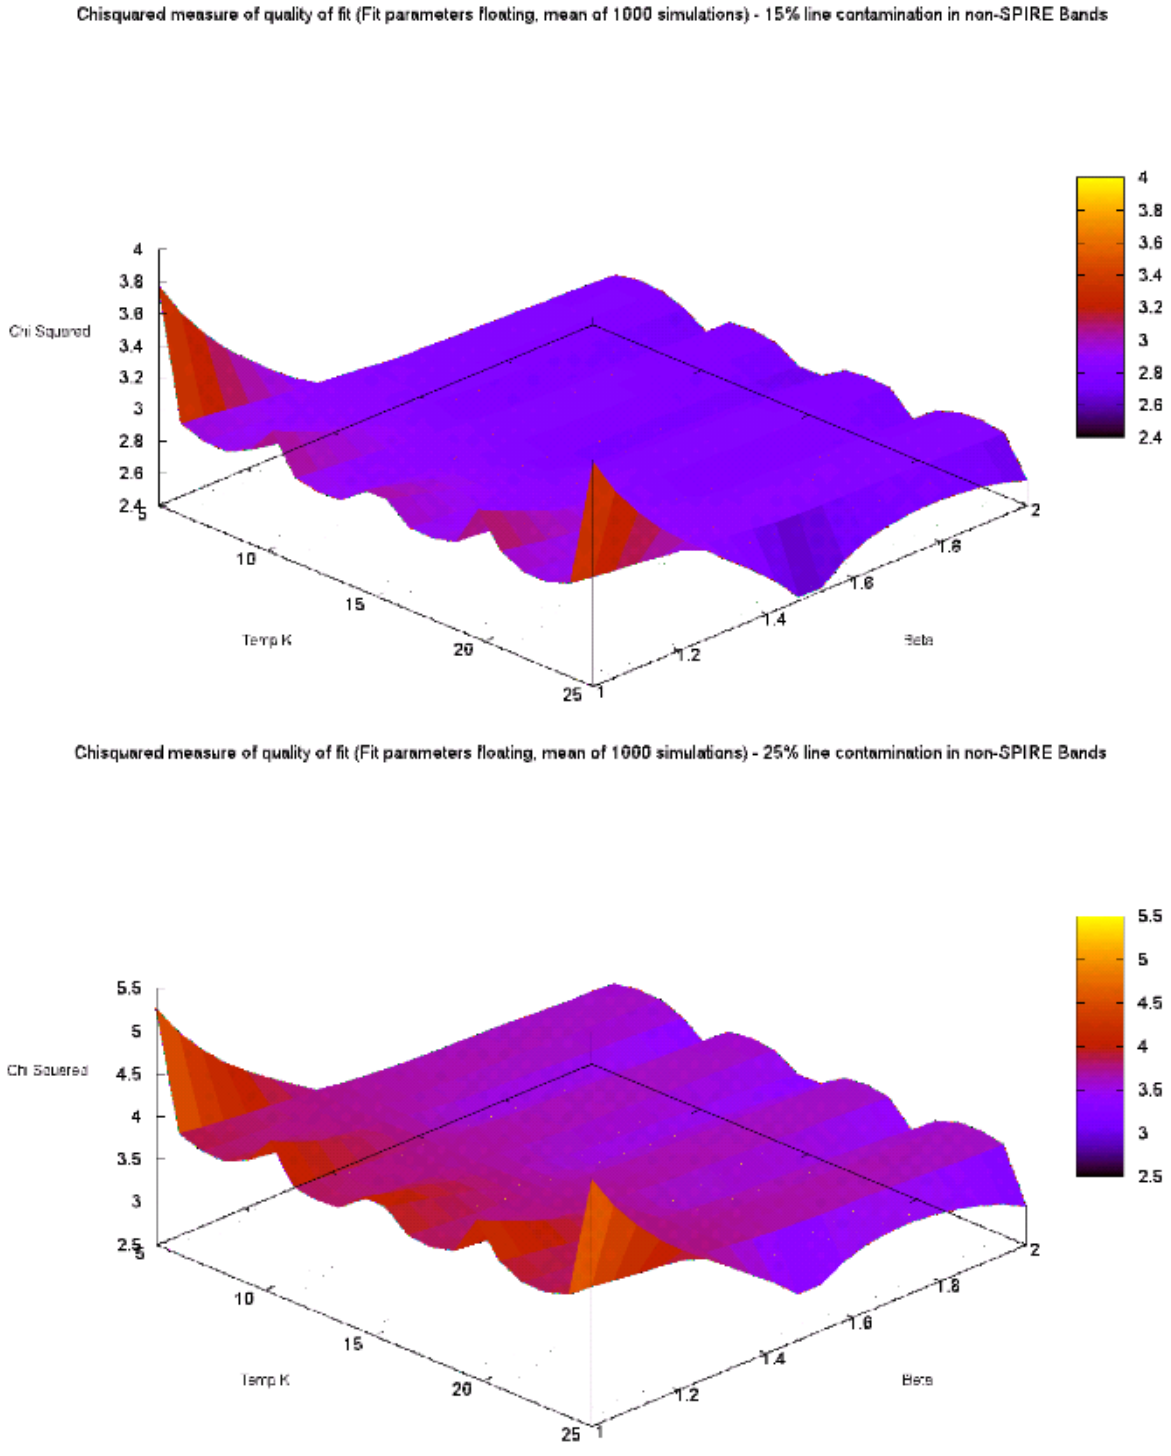

Figure 5.34: Line contamination with free fit parameters across non-SPIRE bands χ, 15% and 25% line contamination percentage.

It is shown that for so few datapoints, we need to restrict the number of free parameters to obtain a clearer idea of the effects of line contamination. As discussed at the beginning of this section, this leads to the consideration of two specific cases.

1. Fixing β_{rec} to be 2 (a common procedure in greybody fitting), and considering how the recovered temperature is affected.

2. Fixing the temperature of the greybody T_{rec} and considering the effect on recovering β_{rec}.

The second case we shall split into two, by considering how recovery is affected when the greybody peak wavelength is within, or outside, the range we are considering. β has a more pronounced effect at longer wavelengths. This means that in the following cases, the most interesting ones are where β_{rec} is fixed and we examine the effect of T_{rec} when line contamination is added to the SPIRE bands. We also fix the temperature and then examine the effect of line contamination in non-SPIRE bands in recovering β_{rec}. For the first case, we can examine how fixing β_{rec} affects recovery of temperature when there is line contamination in the SPIRE bands. Here we are contaminating the shorter wavelength bands and so we can consider β_{rec} as fixed, as β affects the longer wavelengths more. Conversely, by adding line contamination to the longer wavelength bands we can examine the effect on recovering β_{rec}, and so the temperature is fixed. This case we then split into two, considering when we fix the temperature such that the greybody peak is either inside or outside our range. As a result, the following simulations will be global, i.e applying line contamination to a number of wavebands simultaneously.

5.6.2 Greybody fitting - fixing β

5.6.2.1 Contamination in SPIRE bands

Figures 5.35 and 5.36 show the effect of line contamination on the SPIRE bands. Fixing β_{rec} in this way has affected the recovered fit temperature in different ways for different initial values of β_{model}. Where $\beta_{model} = 2$, line contamination in the SPIRE bands has resulted in the recovered temperature being around 5K higher. For other values of β_{model} the effect is to produce a lower greybody temperature of up to around 15K, with this effect being more pronounced for lower values of β_{model}.

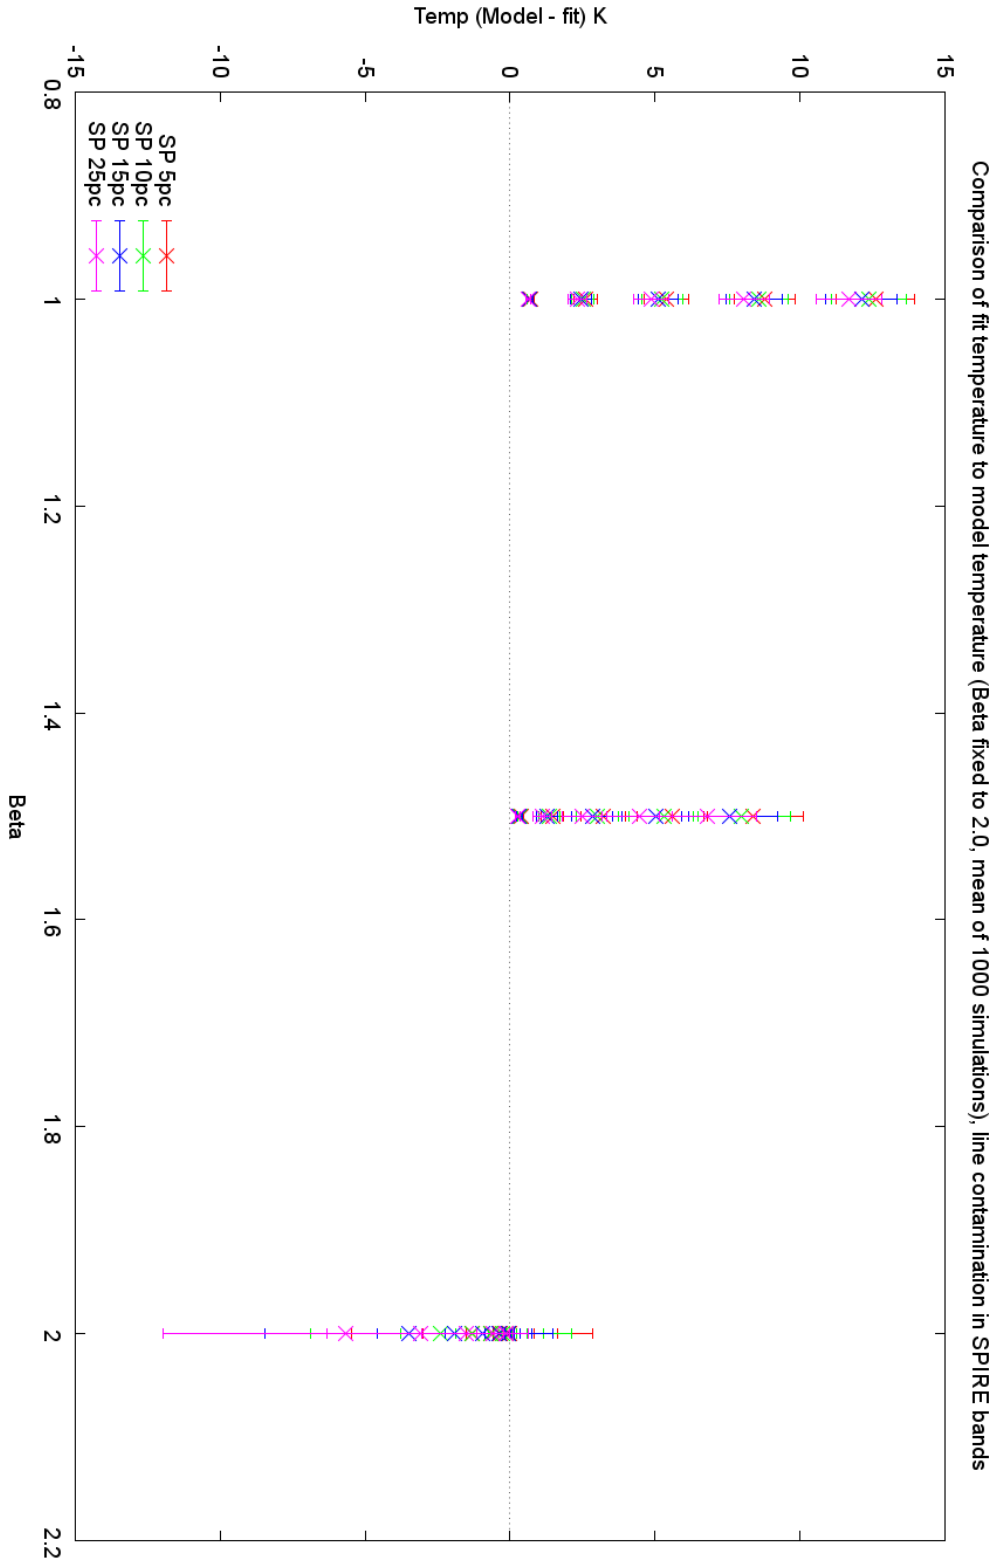

Figure 5.35: Line contamination with β_{rec} fixed across SPIRE bands, effect on T_{rec}, grouped by β_{model}.

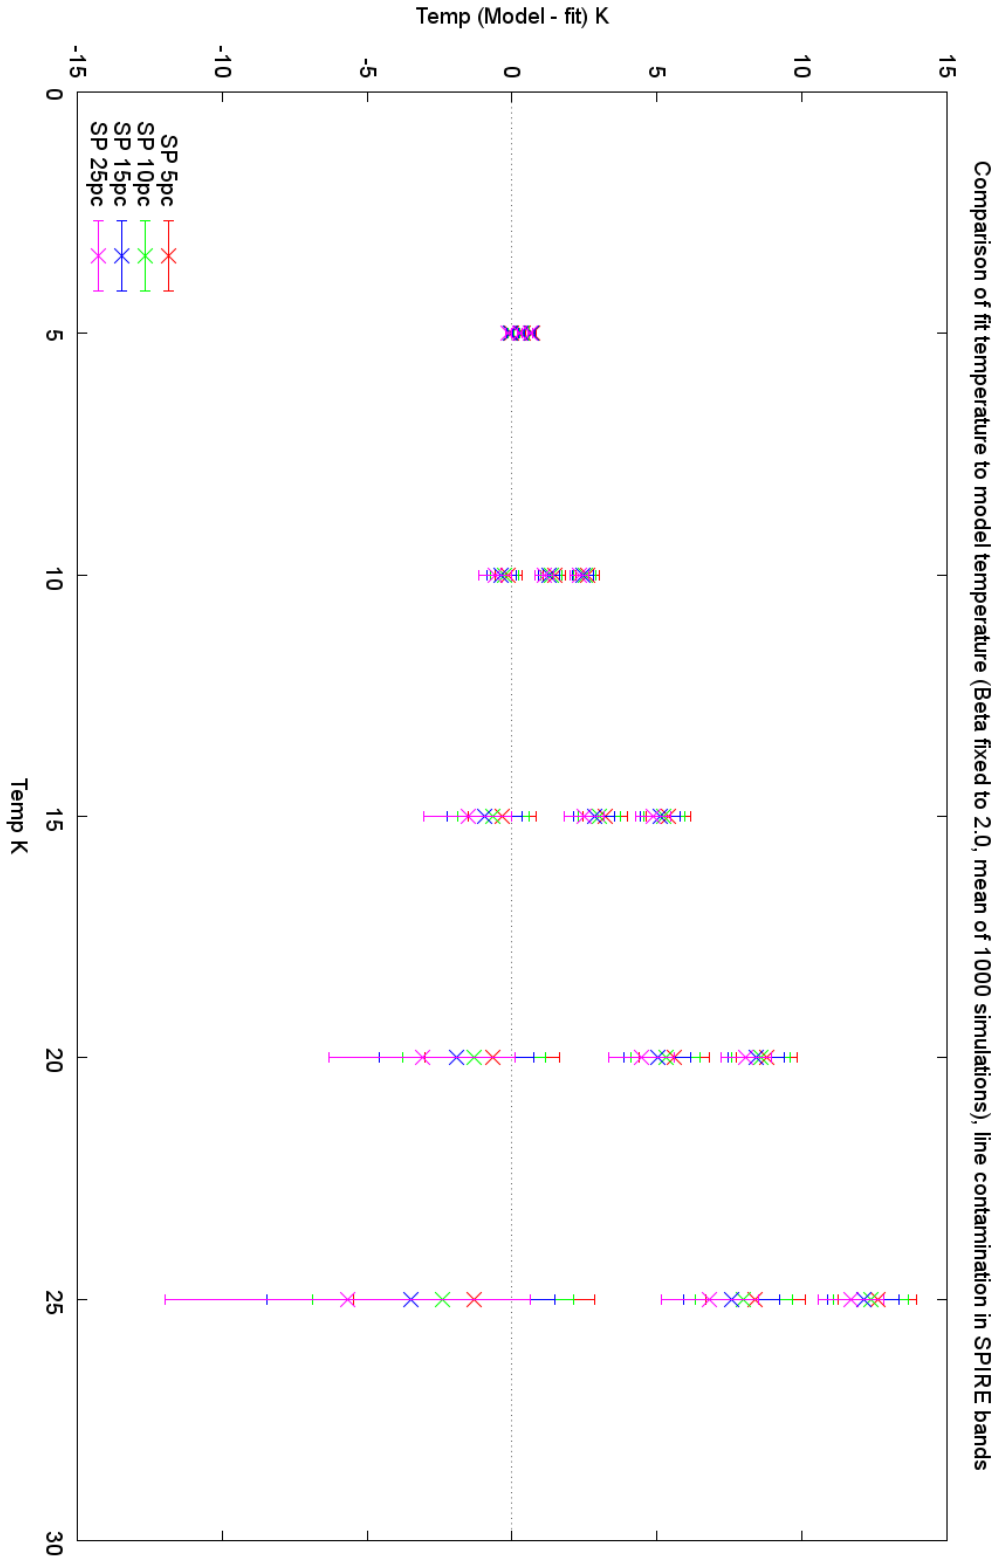

Figure 5.36: Line contamination with β_{rec} fixed across SPIRE bands, effect on T_{rec}.

Figure 5.36 shows that as the greybody temperature increases, on the whole the recovered temperature decreases. The cases where it actually increases are those where the initial $\beta_{model} = 2$. The degree of deviation in temperature increases with the initial model temperature as well. This plot shows the cases where the initial $\beta_{model} = 2$, as the grouping on the lower part of each temperature group. It can be noted that these groupings show a larger range of errors as temperature increases. As the temperature increases, the peak is moved further out of the bandwidth of our sample. As a result, we are seeing the line contamination steepening the slope of the shorter wavelengths. When the initial $\beta_{model} = 2$, then the line contamination effect on the slope leads the fit to produce a higher temperature. For other initial values of β_{model}, we see the fixing of the emissivity drives the temperature down. For lower temperatures, the line contamination does not move the peak position by a great deal and so the recovered temperature is within $\pm 5K$ until the peak moves out of the band. In this case, the recovered temperature is almost always lower than the original greybody temperature. Where the initial $\beta_{model} = 2$, the difference in the two is smaller. As before, the degree in difference between the two increases with the original greybody temperature.

5.6.2.2 Contamination in non-SPIRE bands

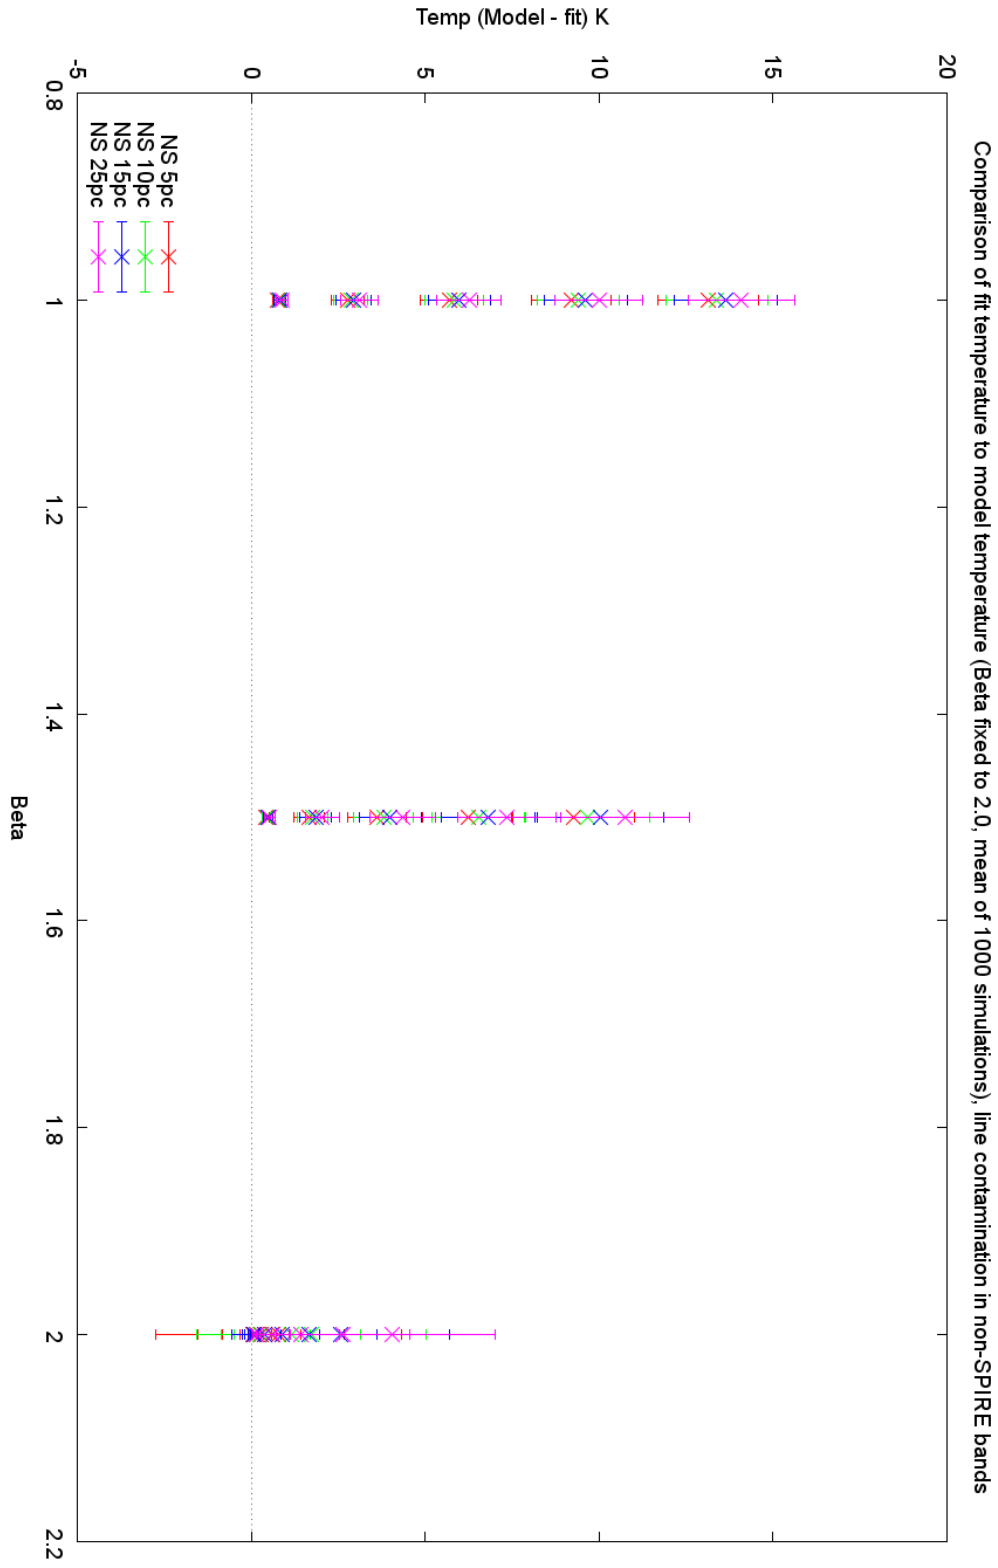

Figure 5.37: Line contamination with β_{rec} fixed across non-SPIRE bands, effect on T_{rec}, grouped by β_{model}.

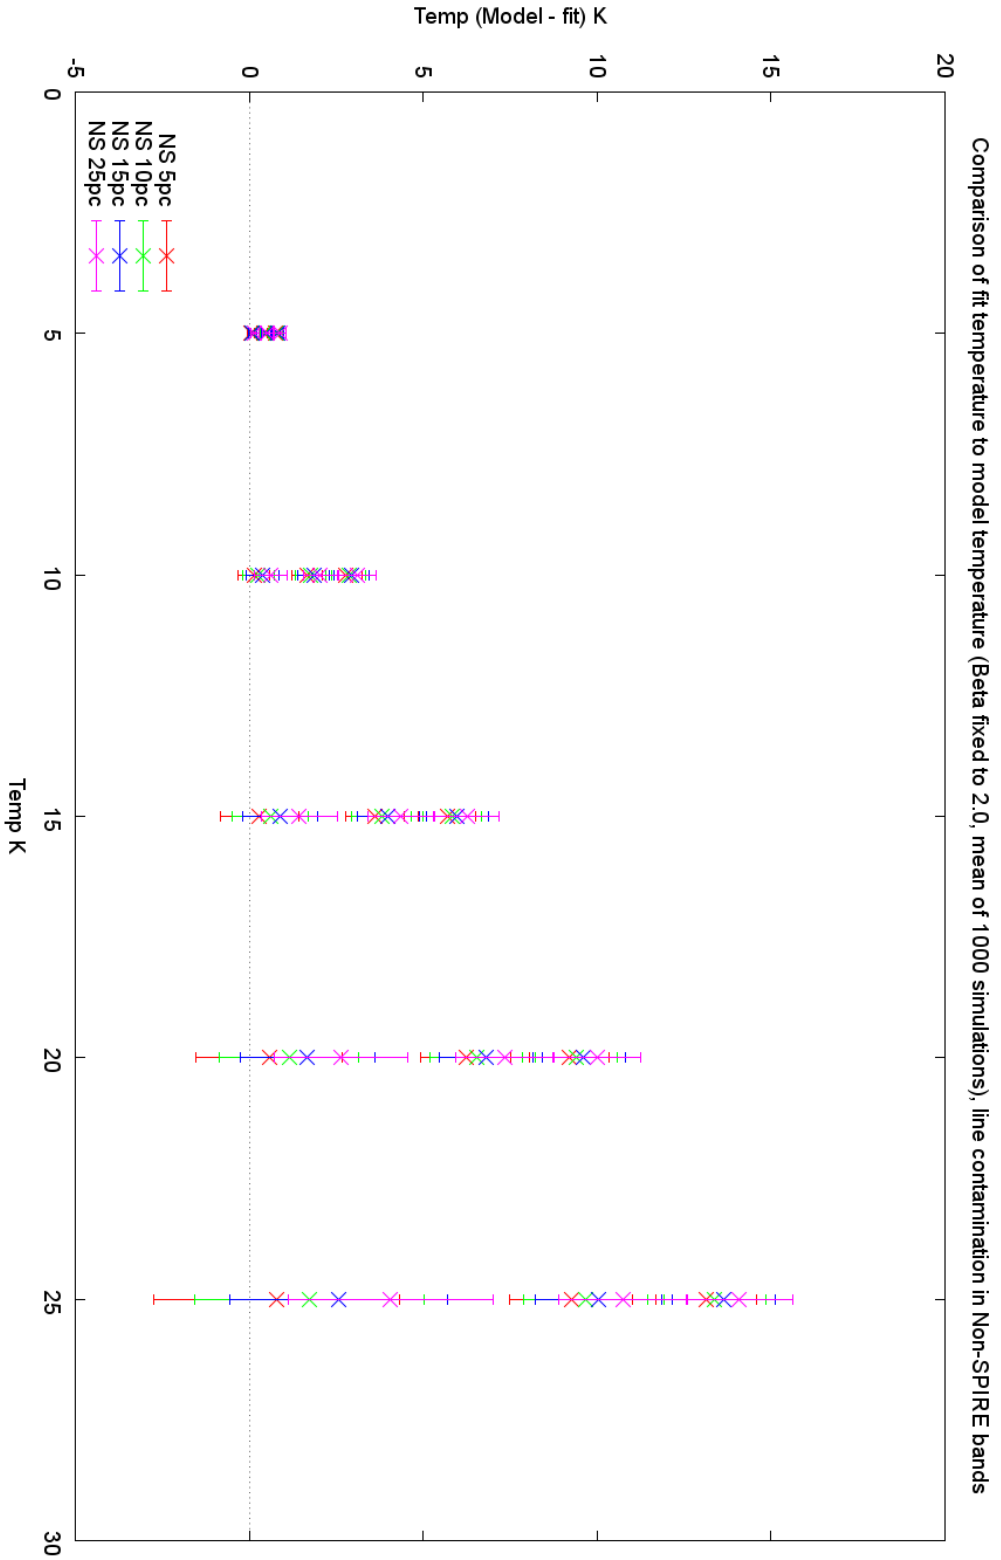

Figure 5.38: Line contamination with β_{rec} fixed across non-SPIRE bands, effect on T_{rec}.

5.6.2.3 Fixed β case - a closer examination

For global SPIRE contamination, we can examine the effect of increasing line contamination on temperature recovery (Figures 5.39 and 5.40) in closer detail. We can then examine the effect on recovery on non-SPIRE waveband line contamination (Figures 5.41 and 5.42). Figures 5.43 to 5.46 use χ as a metric for quality of recovery, where we use the metric to give a measure of how well the overall fit is.

From hereon we will often write terms as xx(XX)/yy(YY) where for either T or β, xx is the deviation from the original parameter of XX and yy is the deviation from the original parameter at YY. Figures 5.39 and 5.40 outline how well we recover temperature when β_{rec} is fixed at 2.0 for SPIRE band line contamination. Here we see a trend develop. Where the initial greybody β_{model} is also 2.0, the temperature recovered is driven high. Over the temperature range, as the initial temperature of the greybody increases, the recovered temperature shows a tendency to be driven higher. For 5% line contamination this range is 0.02K (initial temperature of 5K) to 1.32K (25K) and for 25% line contamination this becomes 0.12K (5K) to 5.68K (25K).

For the lower values of $\beta_{model} = 1.0$ and 1.5 we see the temperature is driven lower, and the degree of this effect is magnified by the difference in β_{model} from the fixed β_{rec}. For $\beta_{model} = 1.0$ we see 5% line contamination driving the temperature low by 0.7K(5K)/12.6K(25K). However, for 25% this becomes 0.62K(5K)/11.7K(25K). Where $\beta_{model} = 1.5$ the recovered temperature is closer to the initial greybody temperature. For 5% line contamination the deviation is 0.37K(5K)/8.4K(25K) and for 25% it is 0.27K(5K)/6.83K(25K)

This would suggest that as the peak of the greybody moves out of the band

we are examining, the fit is strengthened in quality by the loss of this feature. This would seem to be the opposite of what we would expect, in that the peak should help constrain the fit. However it seems in this case the inclusion of line contamination to the spectra in these cases does not affect the recovery adversely.

In conclusion, for low initial β_{model} the temperature is driven lower with increasing strength as the temperature increases, and for $\beta_{model} = 2$ the temperature is driven higher with increasing strength. Figures 5.41 and 5.42 show the effect of line contamination in the non-SPIRE bands. Here the picture is quite different. In all cases, the temperature is driven low, although the degree of deviation from the initial greybody temperature decreases significantly as the initial temperature increases. For $\beta_{model} = 1.0$ we see for 5% line contamination the temperature is driven low by 0.76K(5K)/13.14K(25K) and for 25% this becomes 0.87K(5K)/14.10K(25K) and for $\beta_{model} = 1.5$ this becomes 0.42K(5K)/9.25K(25K) at 5% and 0.53K(5K)/10.73K(25K) at 25%. This shows that the degree of temperature deviation increases as β_{model} decreases. The degree of deviation also increases as the initial greybody temperature increases.

Where $\beta_{model} = 2.0$, the same as the fixed limit, then the degree of deviation at 5% becomes 0.03K(5K)/0.78K(25K) and at 25% is 0.15K(5K)/4.05K(25K). It is unsurprising to show that low amounts of line contamination effect the temperature recovery adversely, although we see at 25% line contamination the effect is still much lower than the case for lower line contamination percentages, with β_{model} lower than the fixed limit. Here we see that fixing β_{rec} has far more effect on the recovery of the temperature than the line contamination in this case.

Examining the quality of fit using χ as a metric, Figures 5.43 and 5.44 show the case for 5% line contamination in SPIRE bands. It shows that the quality of

fit is better in the cases where the greybody β_{model} is the same as the fixed limit. As the temperature increases, for the lowest value of β_{model} the fit degrades slightly more than the others, although the value of χ shows little dynamic range across the graph. Overall, there is a tendency for the fit to deteriorate as temperature increases. As the line contamination increases, it seems that the fit quality becomes slightly better across the board, but eventually at 25% we see a case where the best fit results are for $\beta_{model} = 1.5$, the middle value. Again the range of χ is low. As temperature increases, all cases show a deterioration of the fit quality. For $\beta_{model} = 1.5$ the increase in χ is far less than for the outlying cases. A slope develops as at the lower β_{model} values the fit degrades at a greater rate than the $\beta_{model} = 2$ simulations.

Figures 5.45 and 5.46 show the quality of fit in the non-SPIRE case, and here we do not see the same effect at high line contamination percentages. The range of χ is higher overall, denoting an overall lower quality of recovery. Across the range of line contamination percentages, the better quality fits are those where β_{model} is the same as the fixed fit value, and as β_{model} decreases the fit degrades accordingly. The fit also degrades with increasing line contamination, with the fit degrading most where $\beta_{model} = 1.0$ and as such is the furthest value from the fixed fit value. As temperature increases the χ value increases.

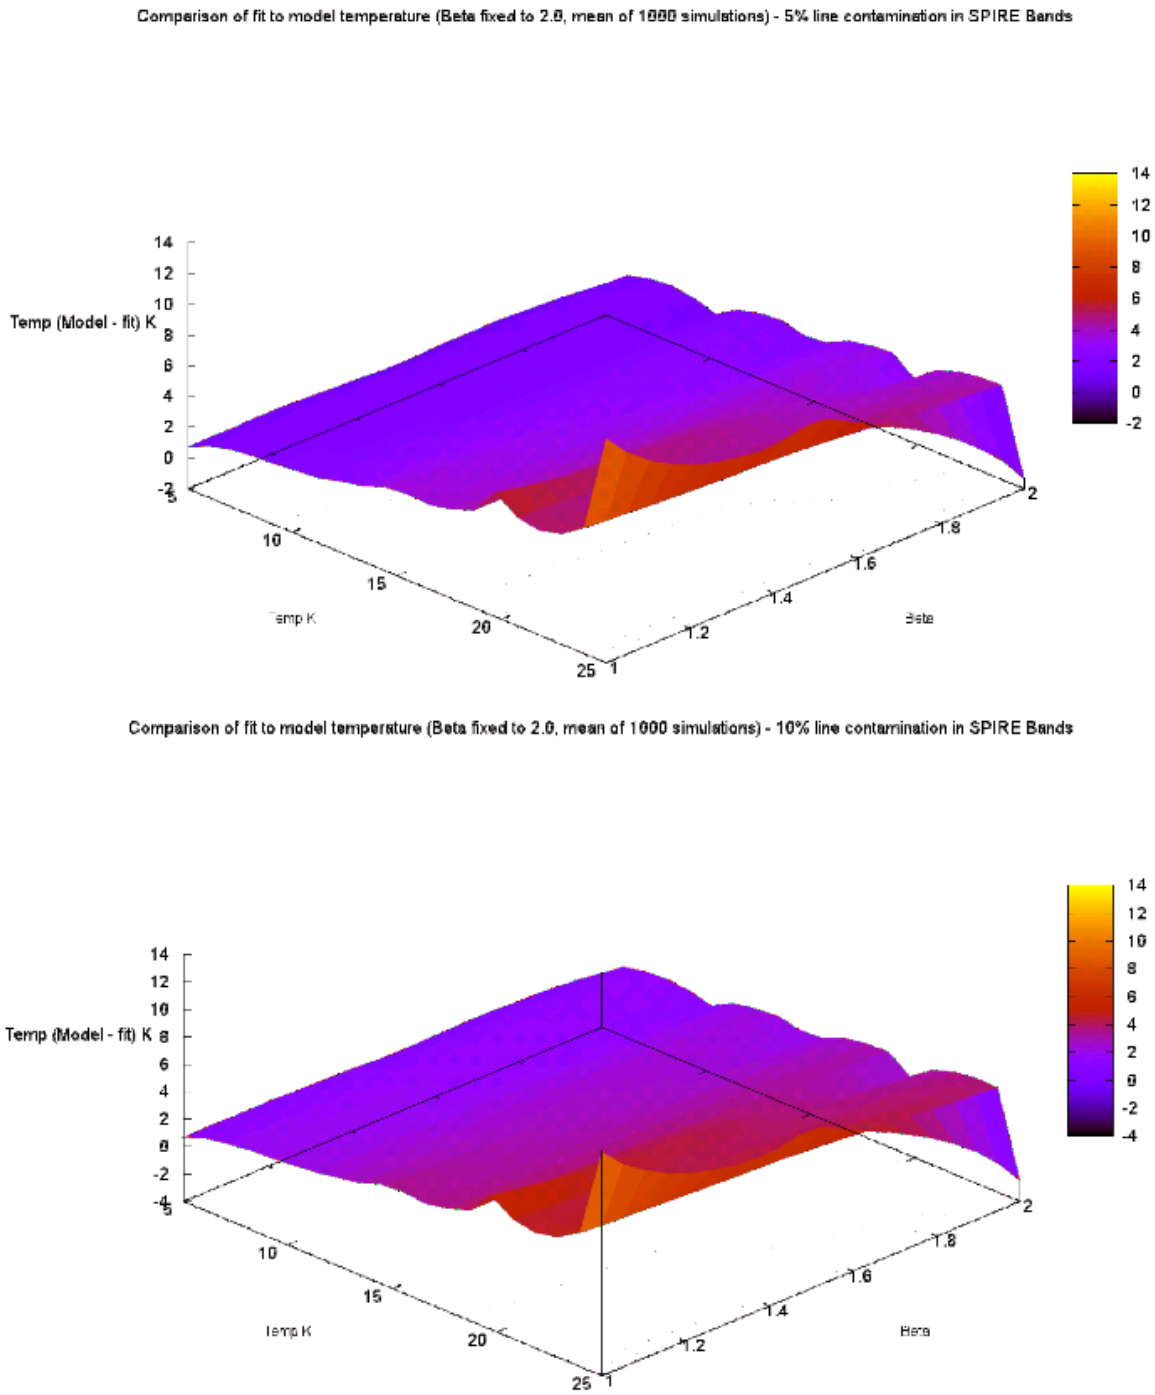

Figure 5.39: Line contamination with β_{rec} fixed across SPIRE bands, effect on T_{rec}, 5% and 10% line contamination percentage.

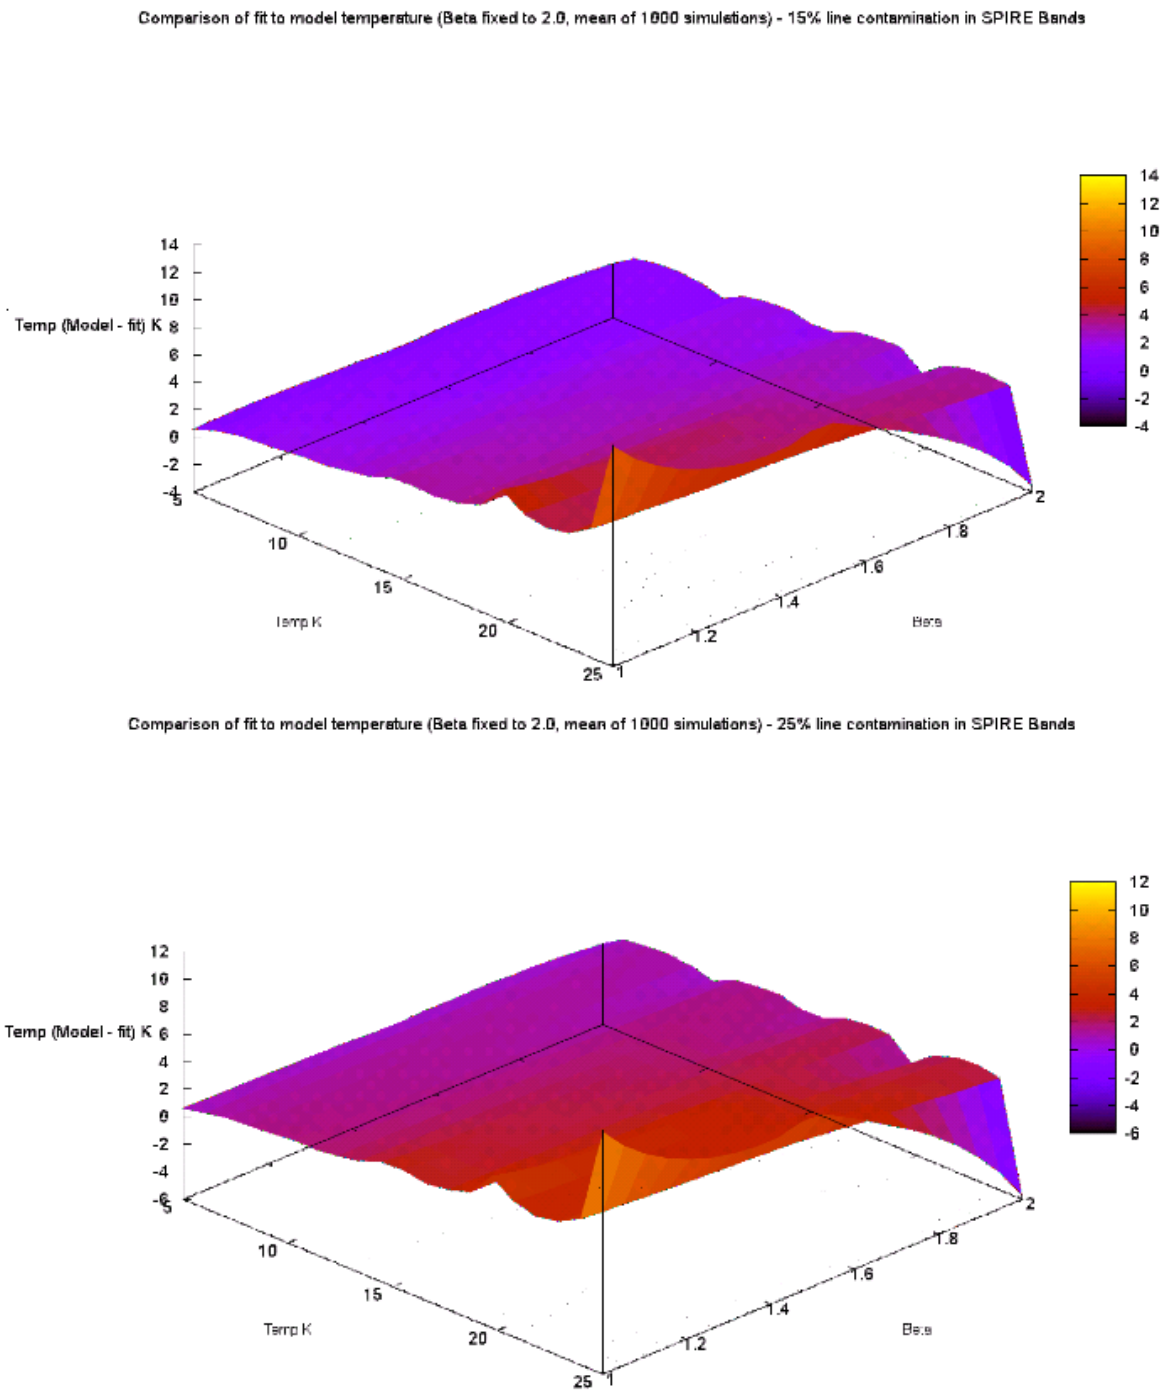

Figure 5.40: Line contamination with β_{rec} fixed across SPIRE bands, effect on T_{rec}, 15% and 25% line contamination percentage.

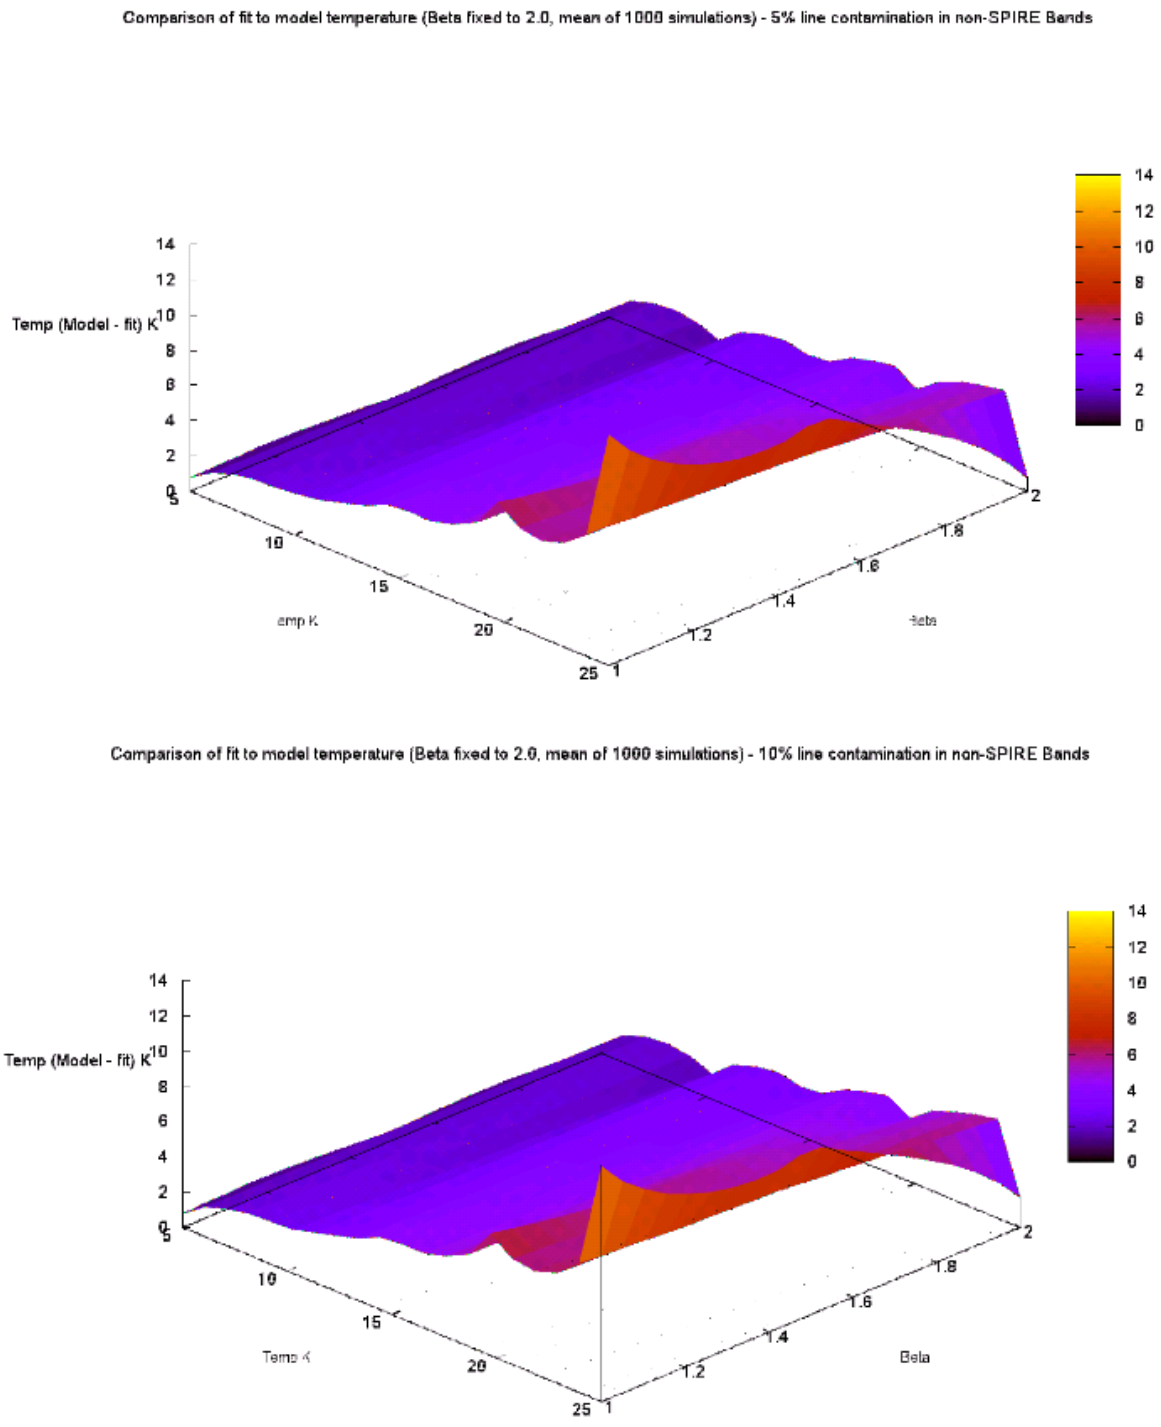

Figure 5.41: Line contamination with β_{rec} fixed across non-SPIRE bands, effect on T_{rec}, 5% and 10% line contamination percentage.

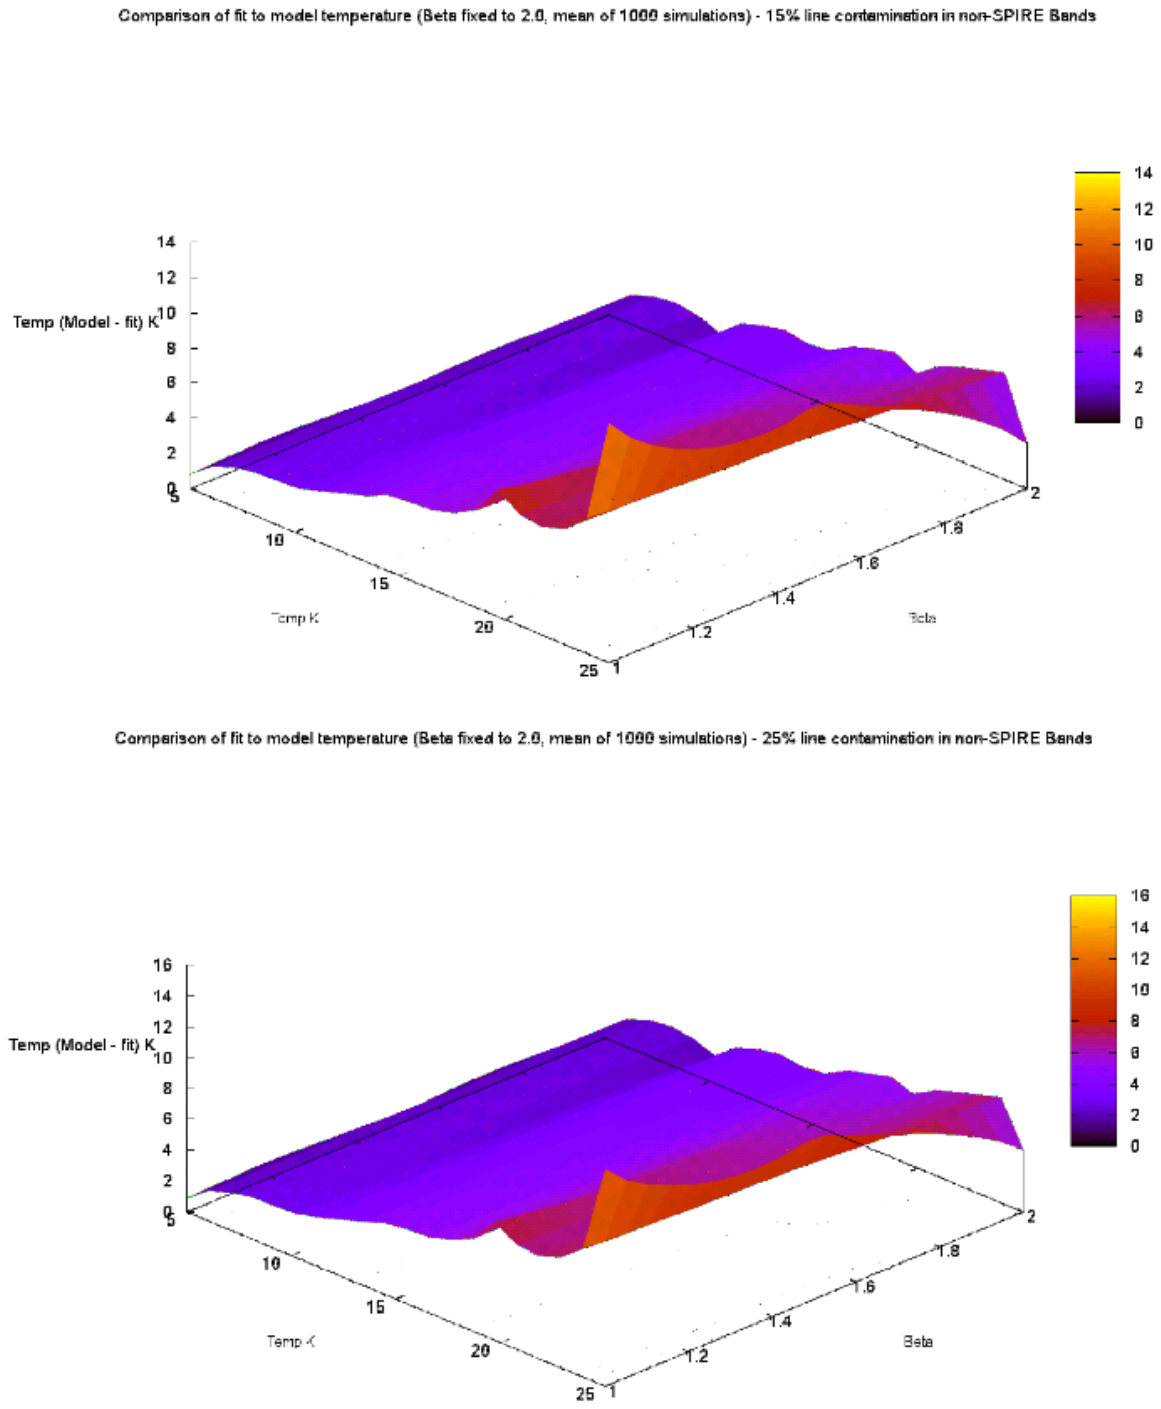

Figure 5.42: Line contamination with β_{rec} fixed across non-SPIRE bands, effect on T_{rec}, 15% and 25% line contamination percentage.

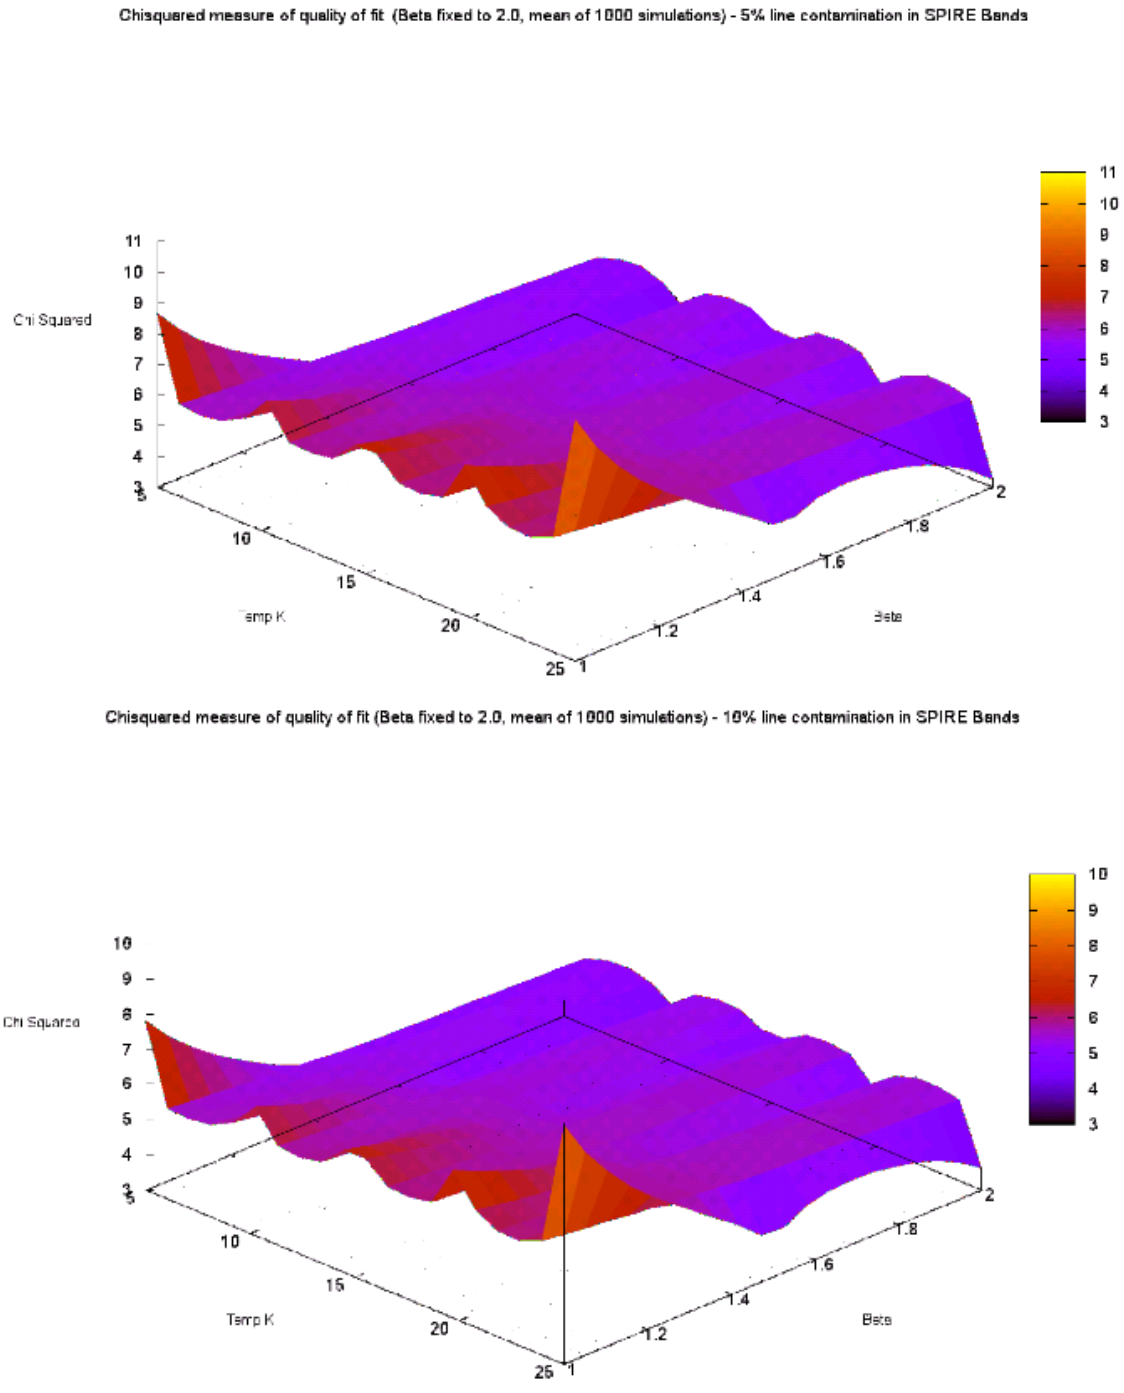

Figure 5.43: Line contamination with β_{rec} fixed across SPIRE bands χ, 5% and 10% line contamination percentage.

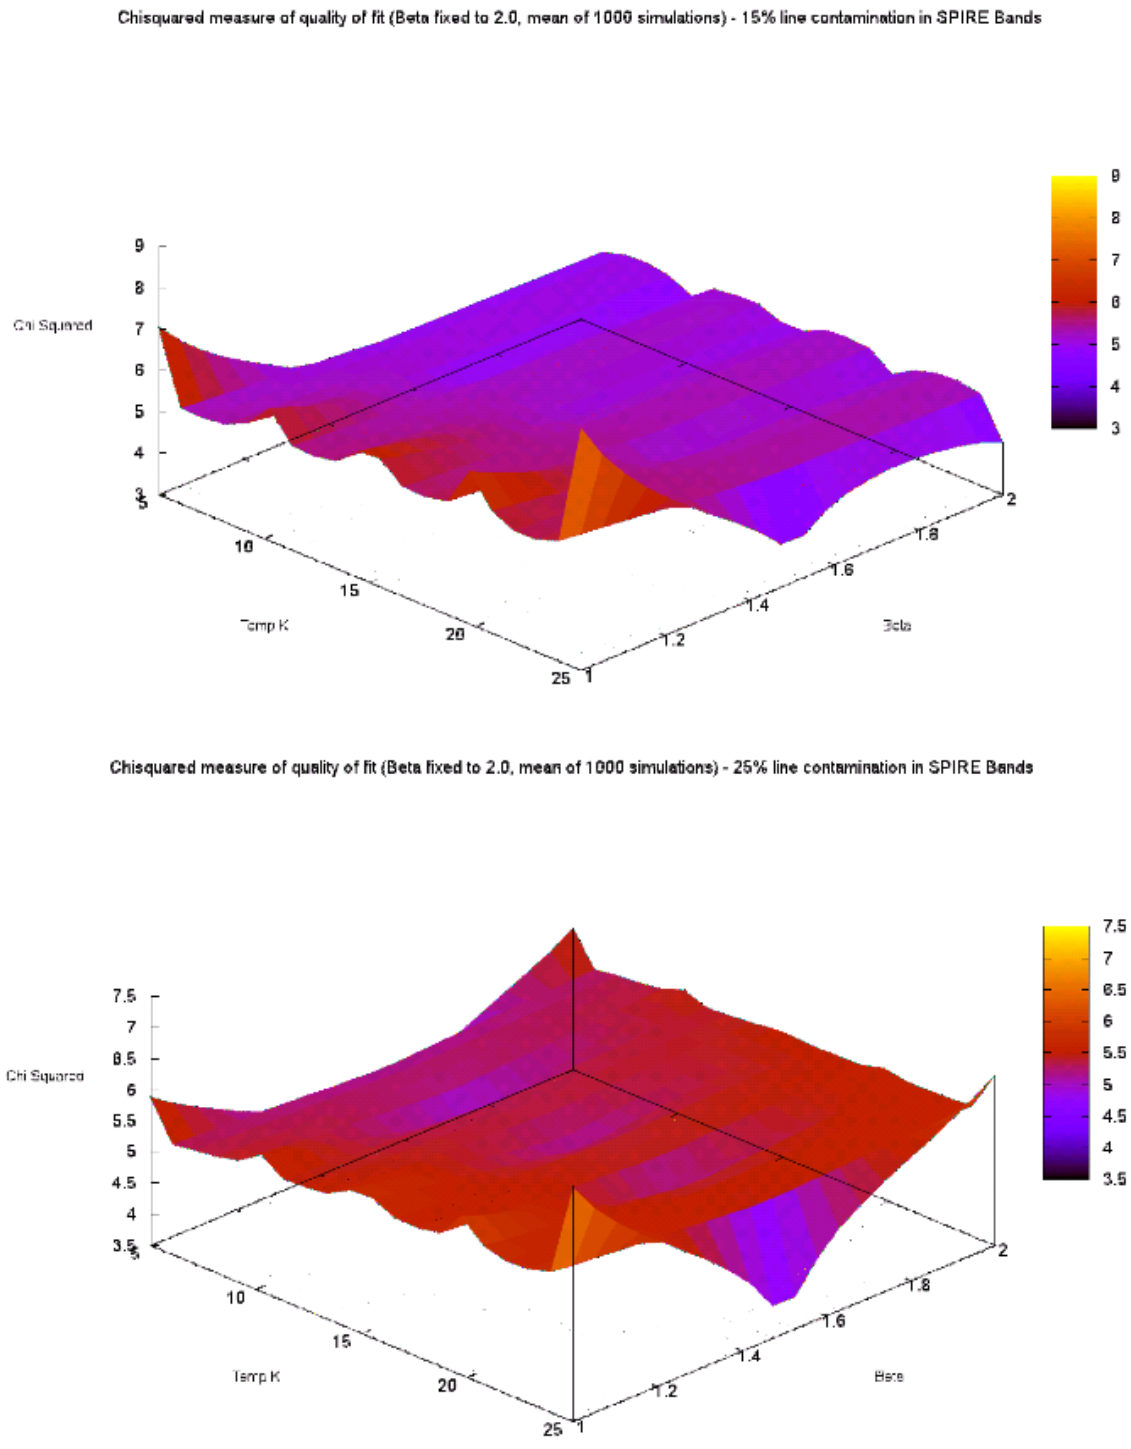

Figure 5.44: Line contamination with β_{rec} fixed across SPIRE bands χ, 15% and 25% line contamination percentage.

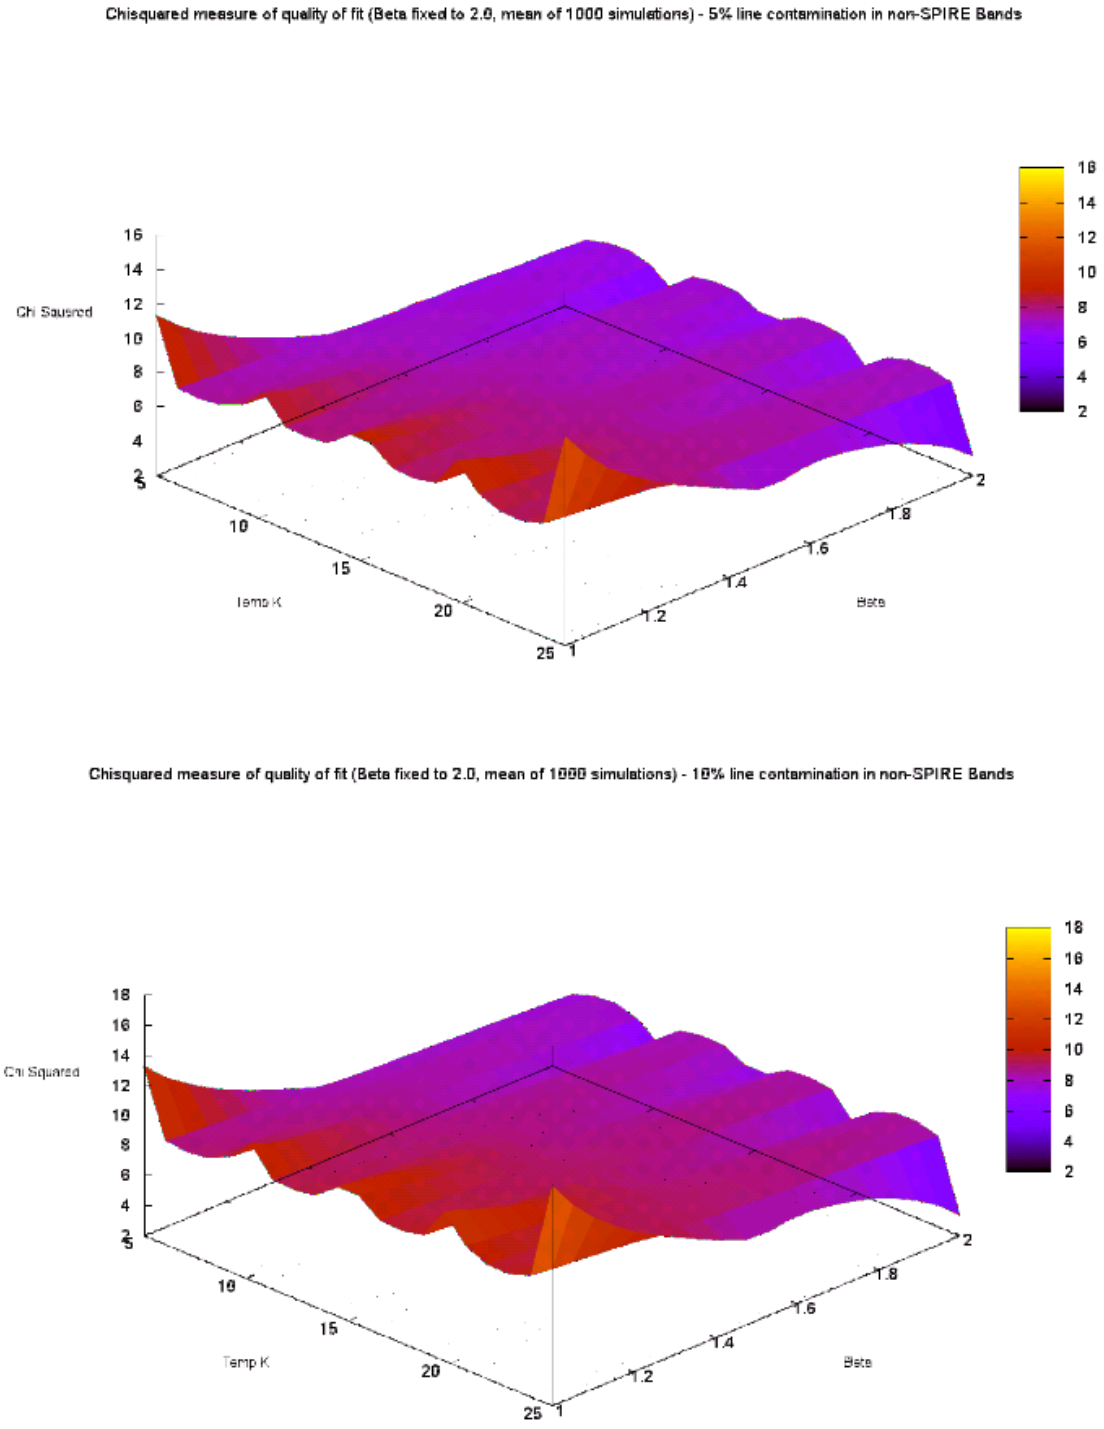

Figure 5.45: Line contamination with β_{rec} fixed across non-SPIRE bands χ, 5% and 10% line contamination percentage.

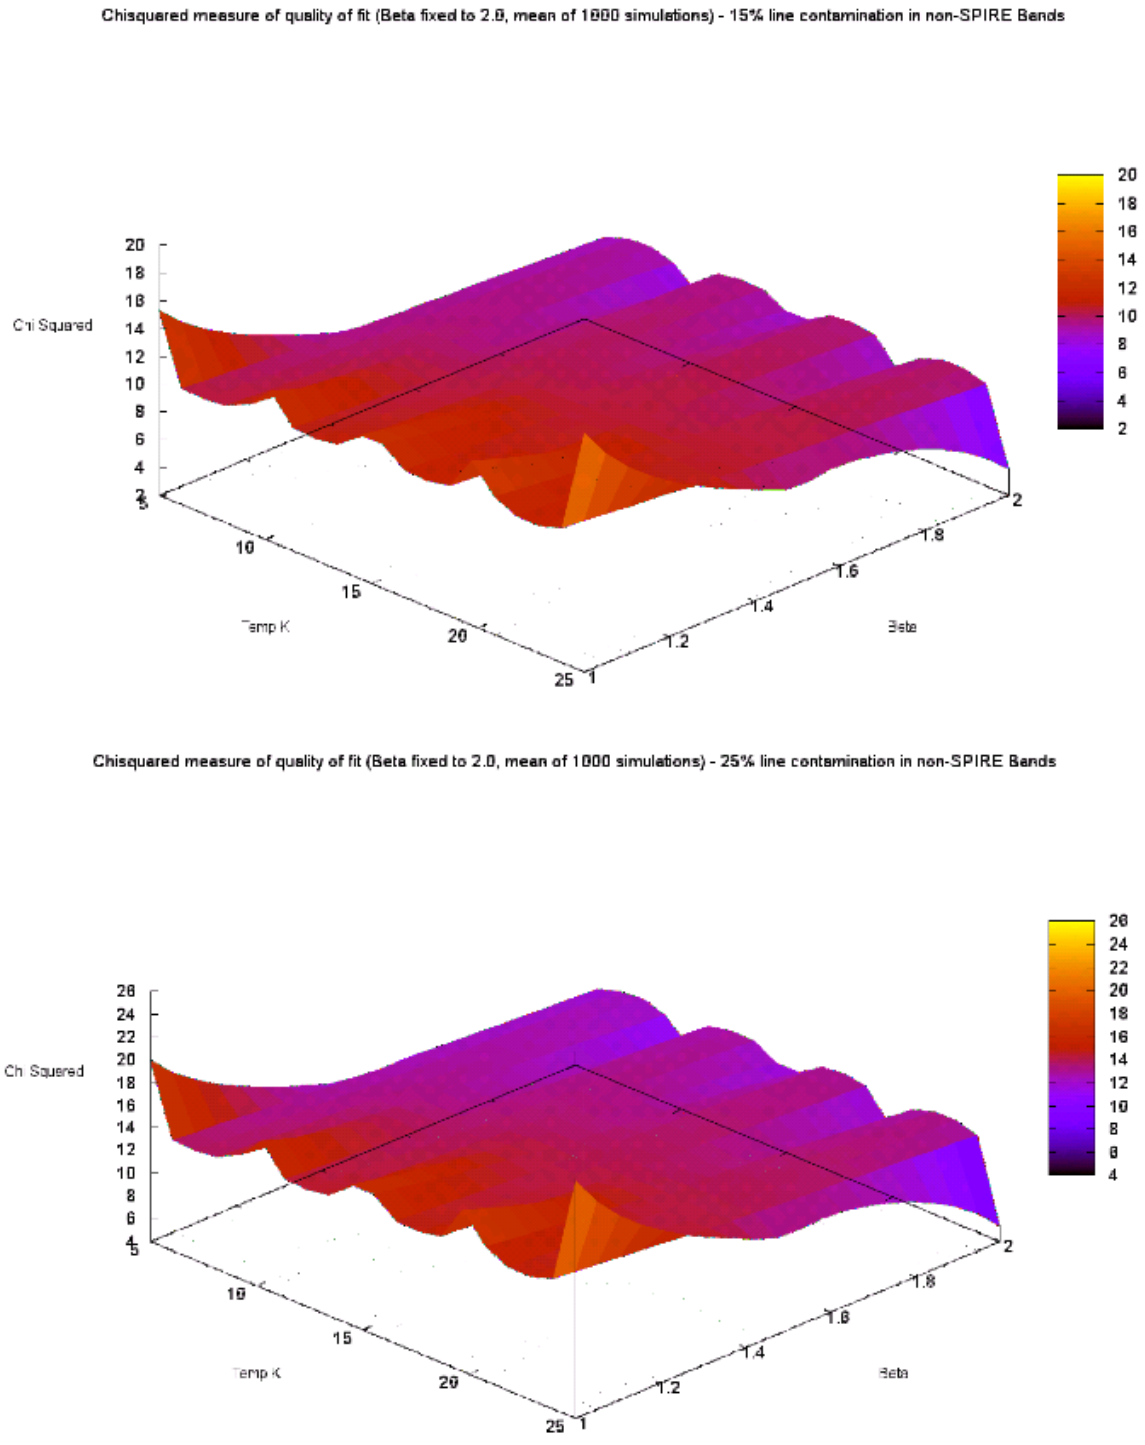

Figure 5.46: Line contamination with β_{rec} fixed across non-SPIRE bands χ, 15% and 25% line contamination percentage.

5.6.3 Greybody fitting - fixing temperature

As explained above, in this case we can consider the different results we obtain when we fix the fit temperature. As before, we generate a range of greybodies between 5K and 25K. We then fix the fit temperature to either 10K (where the peak is within our band) or 25K (where it is outside the band) and examine the consequences.

5.6.3.1 Contamination in SPIRE bands

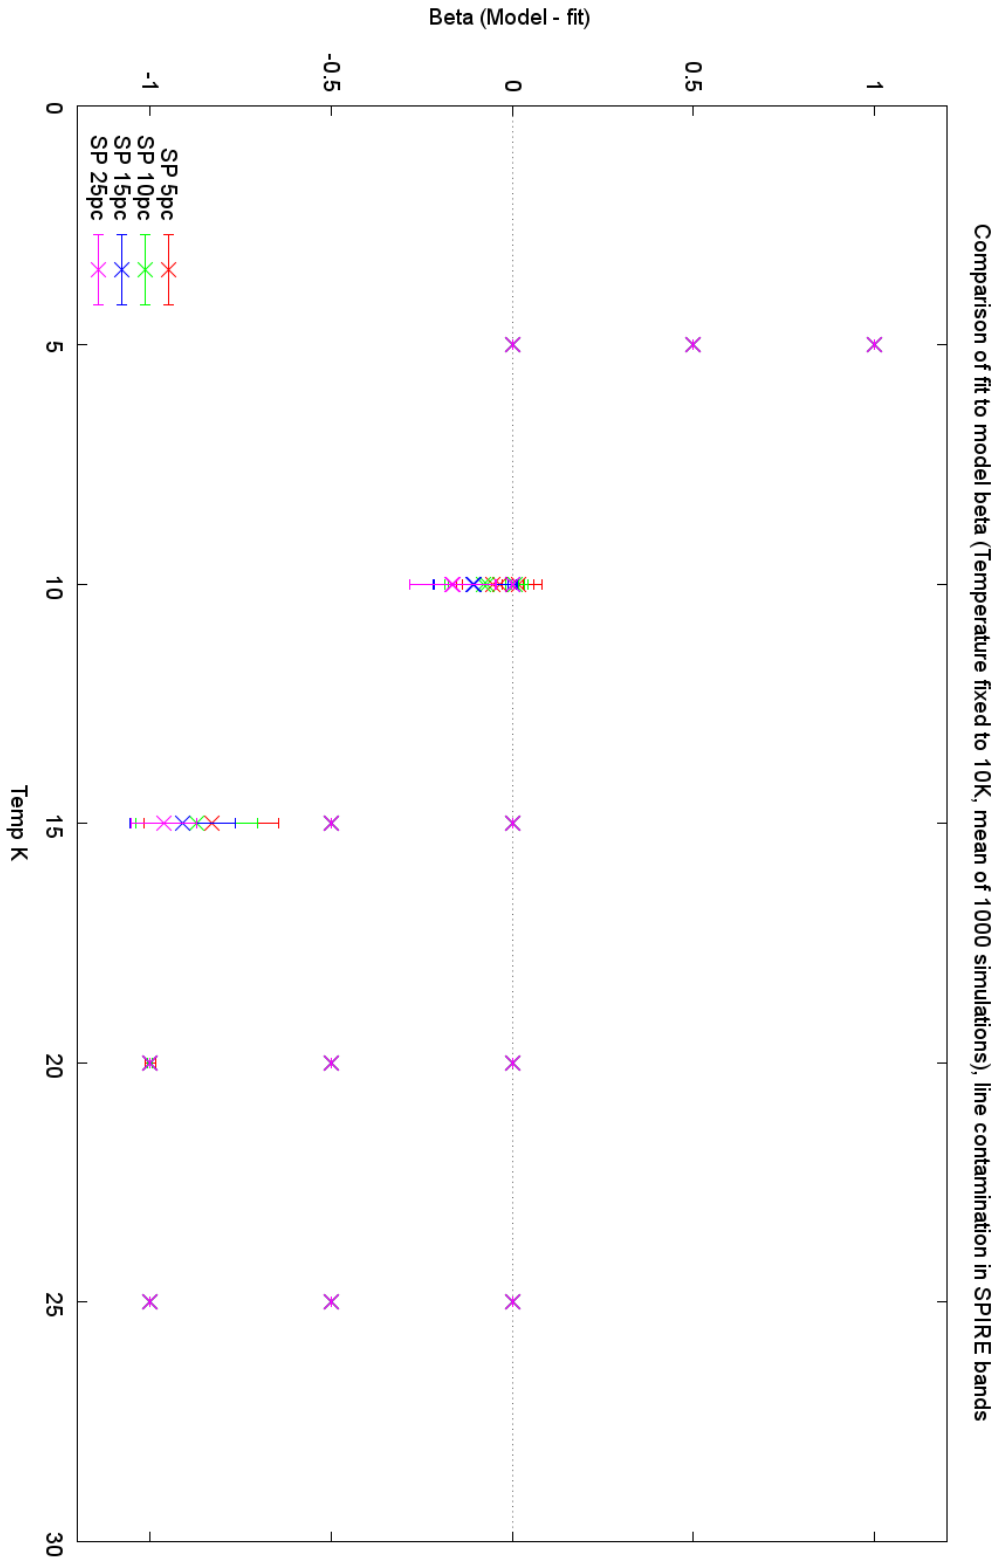

Figure 5.47: Line contamination with $T_{rec} = 10K$ fixed across SPIRE bands, effect on β_{rec}, grouped by T_{model}.

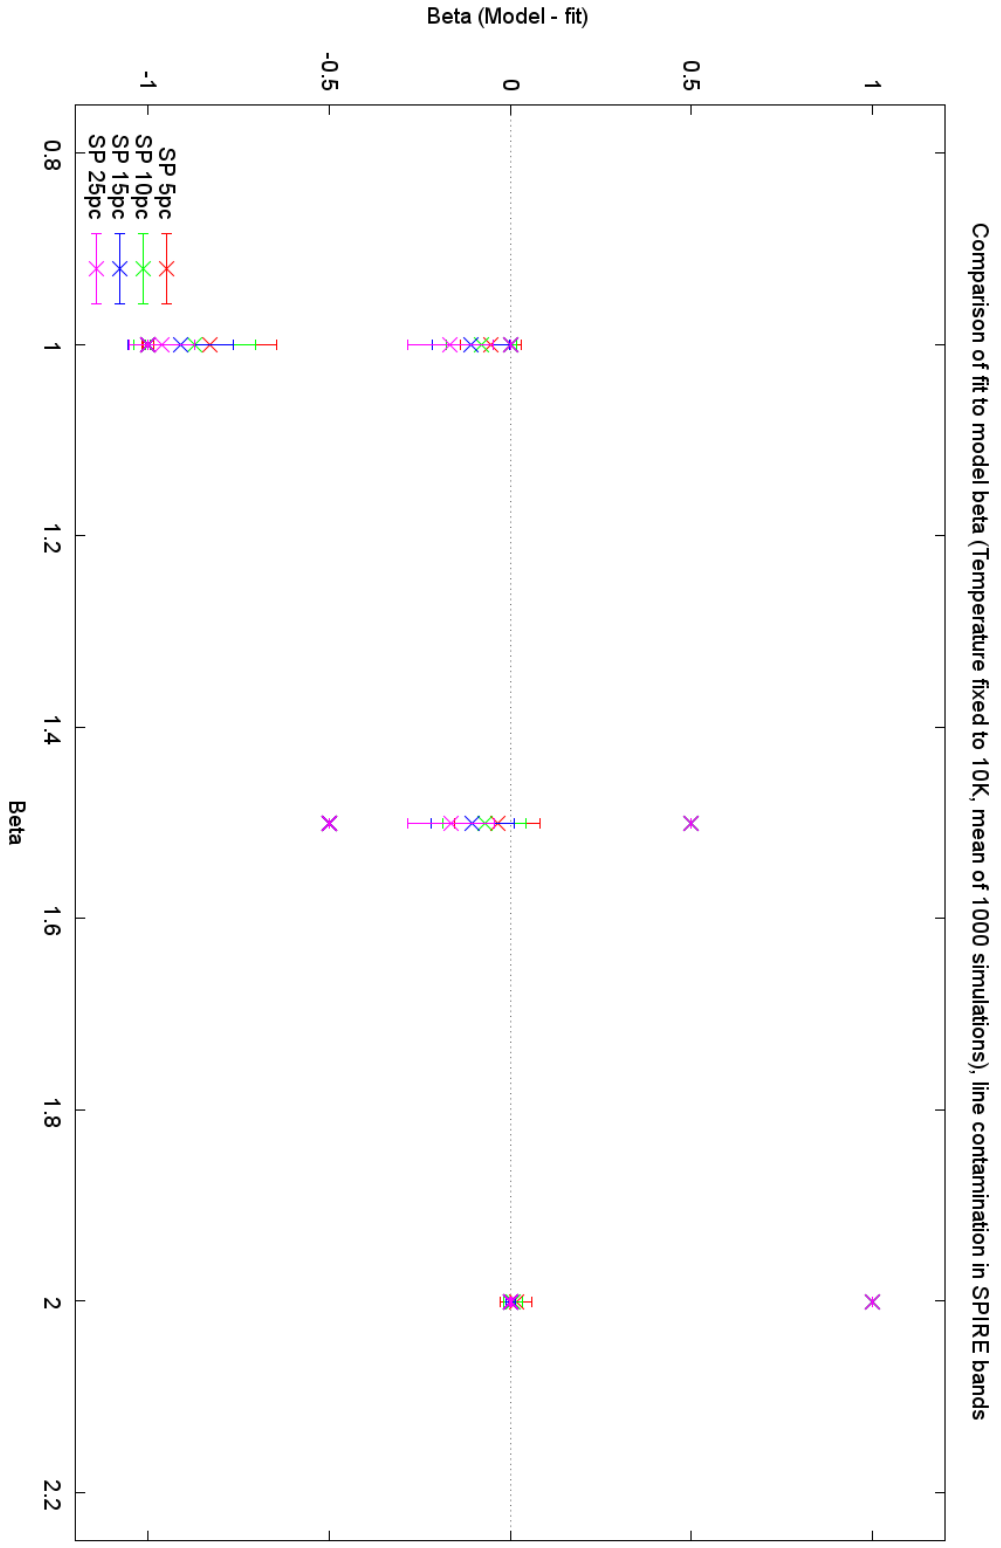

Figure 5.48: Line contamination with $T_{rec} = 10K$ fixed across SPIRE bands, effect on β_{rec}.

For the fixed 10K fits, we obtain Figures 5.47 and 5.48. Figure 5.47 shows that for greybody temperatures less than our fixed temperature, the β_{rec} we obtain from the fit is driven low, and vice versa for higher greybody temperatures. The distinctive pattern we see is where the fit limits of β_{rec} are met. For colder greybodies, β_{rec} is consistently the lower limit of 1, and for warmer greybodies the higher limit of 2 is recovered.

For 15K greybodies, we see that in some cases β_{rec} is recovered high, but not at the limit. The spread of the one set of points is where greybodies with a β_{model} of 1 have a recovered β_{rec} between 1.5 and 2. For greybody temperatures the same as the fixed fit temperature, β_{rec} is slightly higher than β_{model} with a spread of around 0.4. Figure 5.48 shows that when the results are grouped by β_{model}, we obtain the best recovery of this parameter when β_{model} is 2. Only the higher 25% line contamination is strong enough to affect recovery. For other values of β_{model} we see a spread of recovered values.

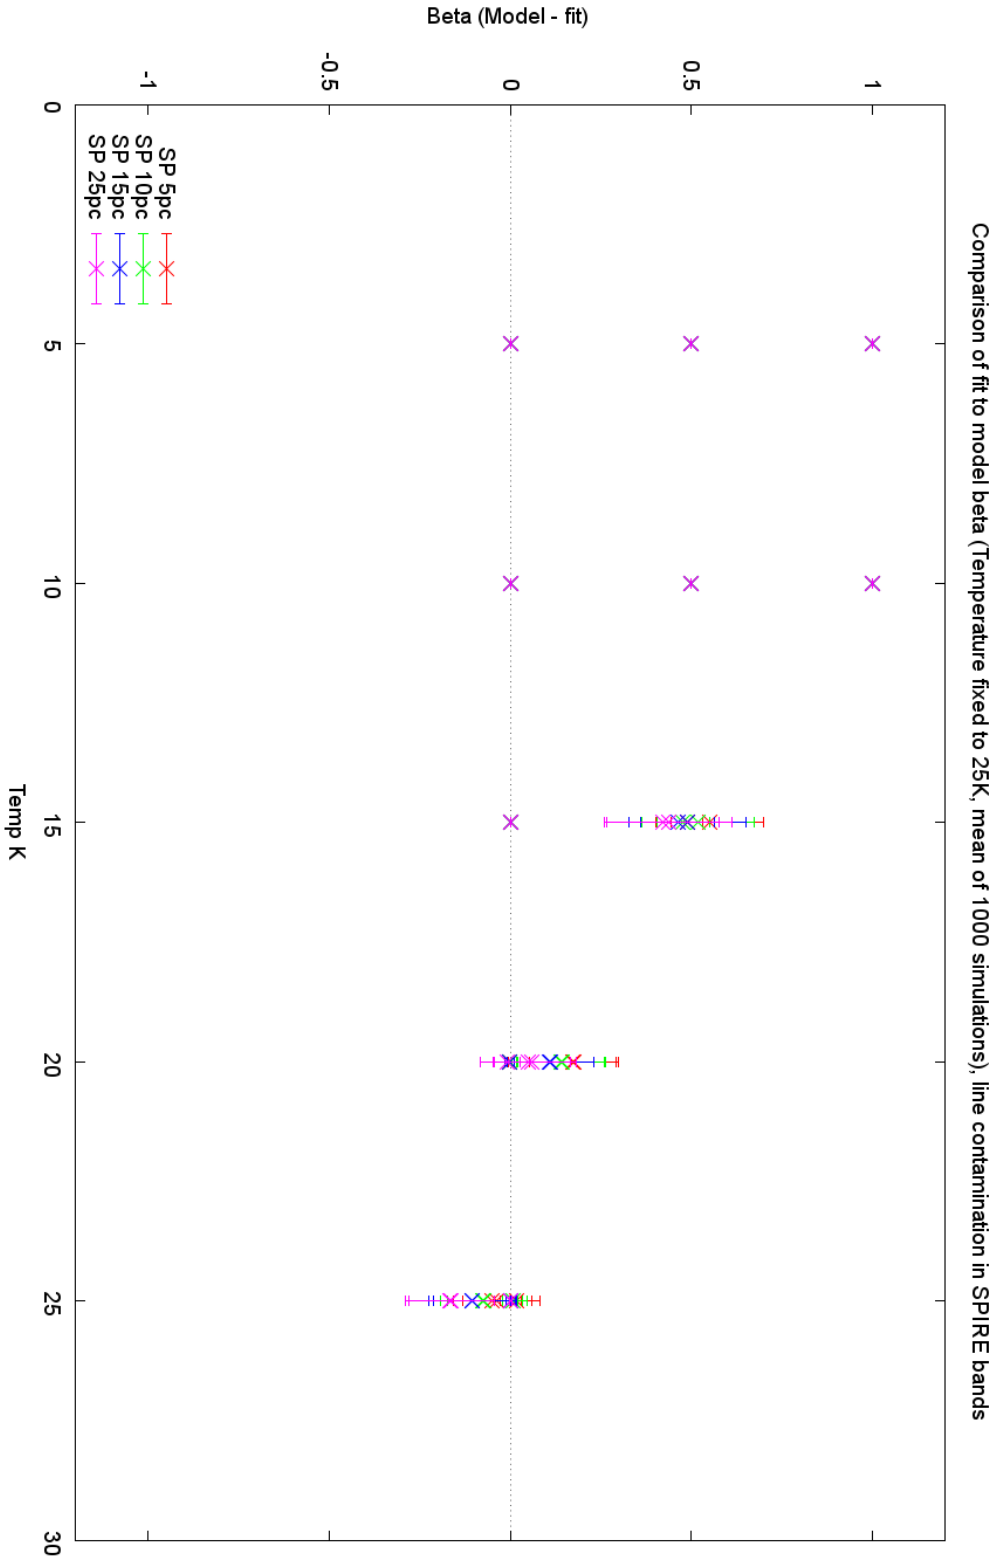

Figure 5.49: Line contamination with $T_{rec} = 25K$ fixed across SPIRE bands, effect on β_{rec}, grouped by T_{model}.

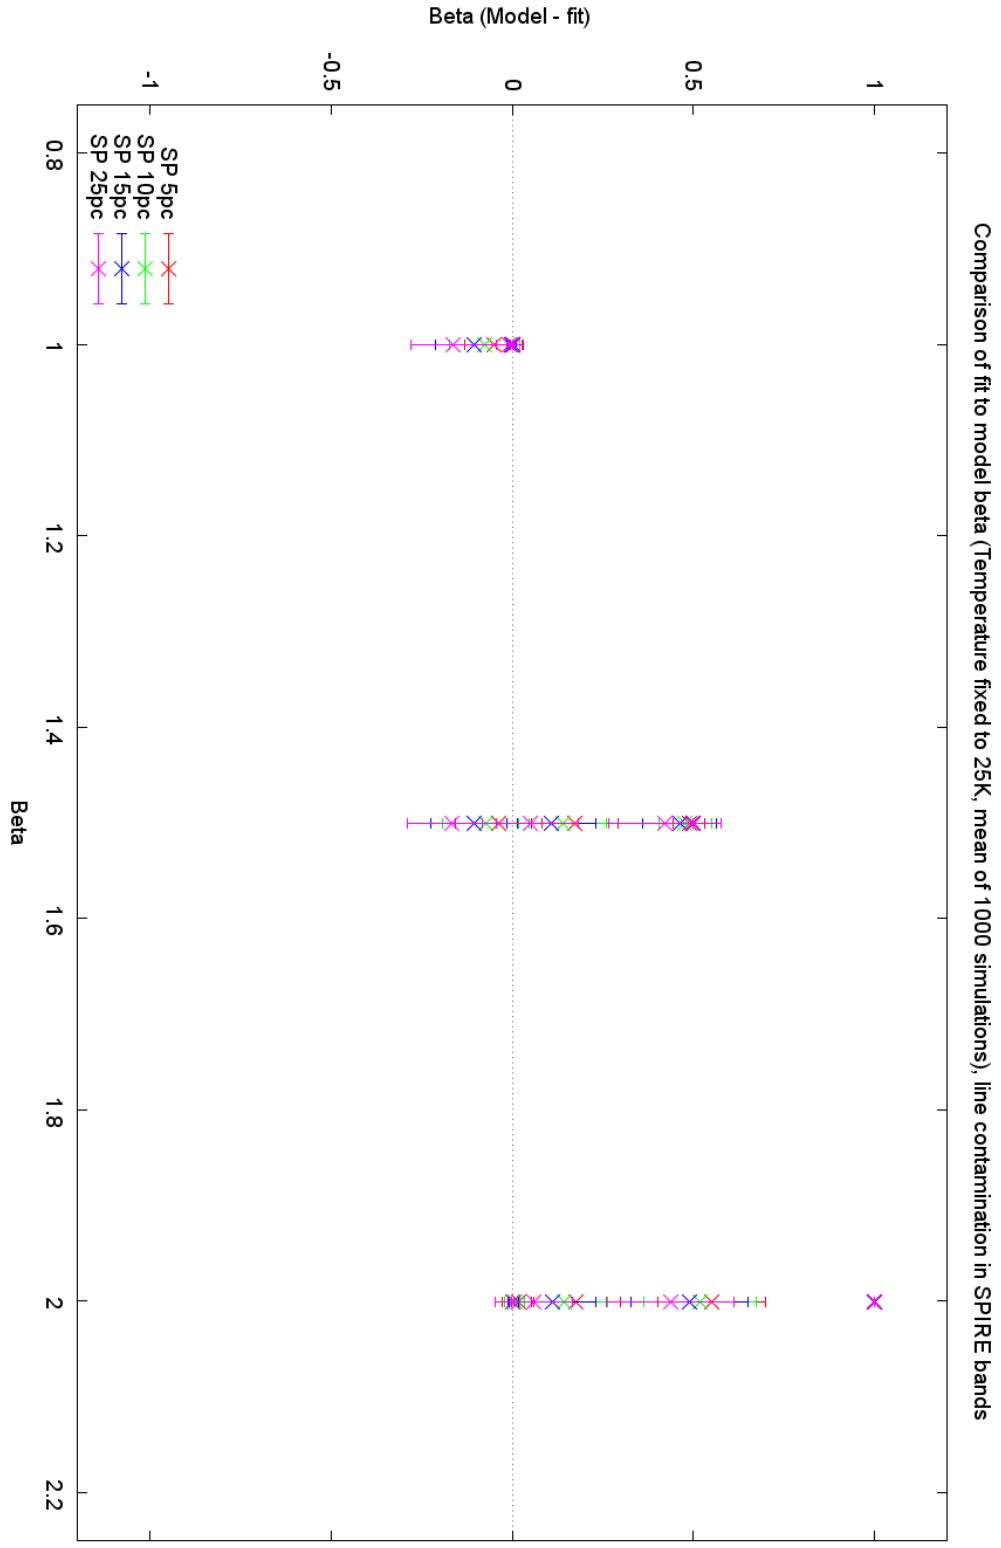

Figure 5.50: Line contamination with $T_{rec} = 25K$ fixed across SPIRE bands, effect on β_{rec}.

For the fixed 25K fits, we obtain Figures 5.49 and 5.50. Figure 5.49 yields a similar result to the 10K diagram of Figure 5.47. In this case however, we see that the recovered β_{rec} is not limited until 10K below our fixed fit temperature. For 5K and 10K greybodies, β_{rec} is consistently recovered at the lower limit. As the temperature approaches the fixed fit temperature, the quality of β_{rec} recovery improves. Figure 5.50 is in some ways an opposite of the 10K equivalent (Figure 5.48). Here, we obtain the best recovery of β_{rec} when it is at the lower limit, although here we find the tendency is for β_{rec} to be consistently higher than the greybody actual value. For middle values of β_{model}, where there is room for the fit to go either way, there is more of a tendency for β_{rec} to decrease.

5.6.3.2 Contamination in non-SPIRE bands

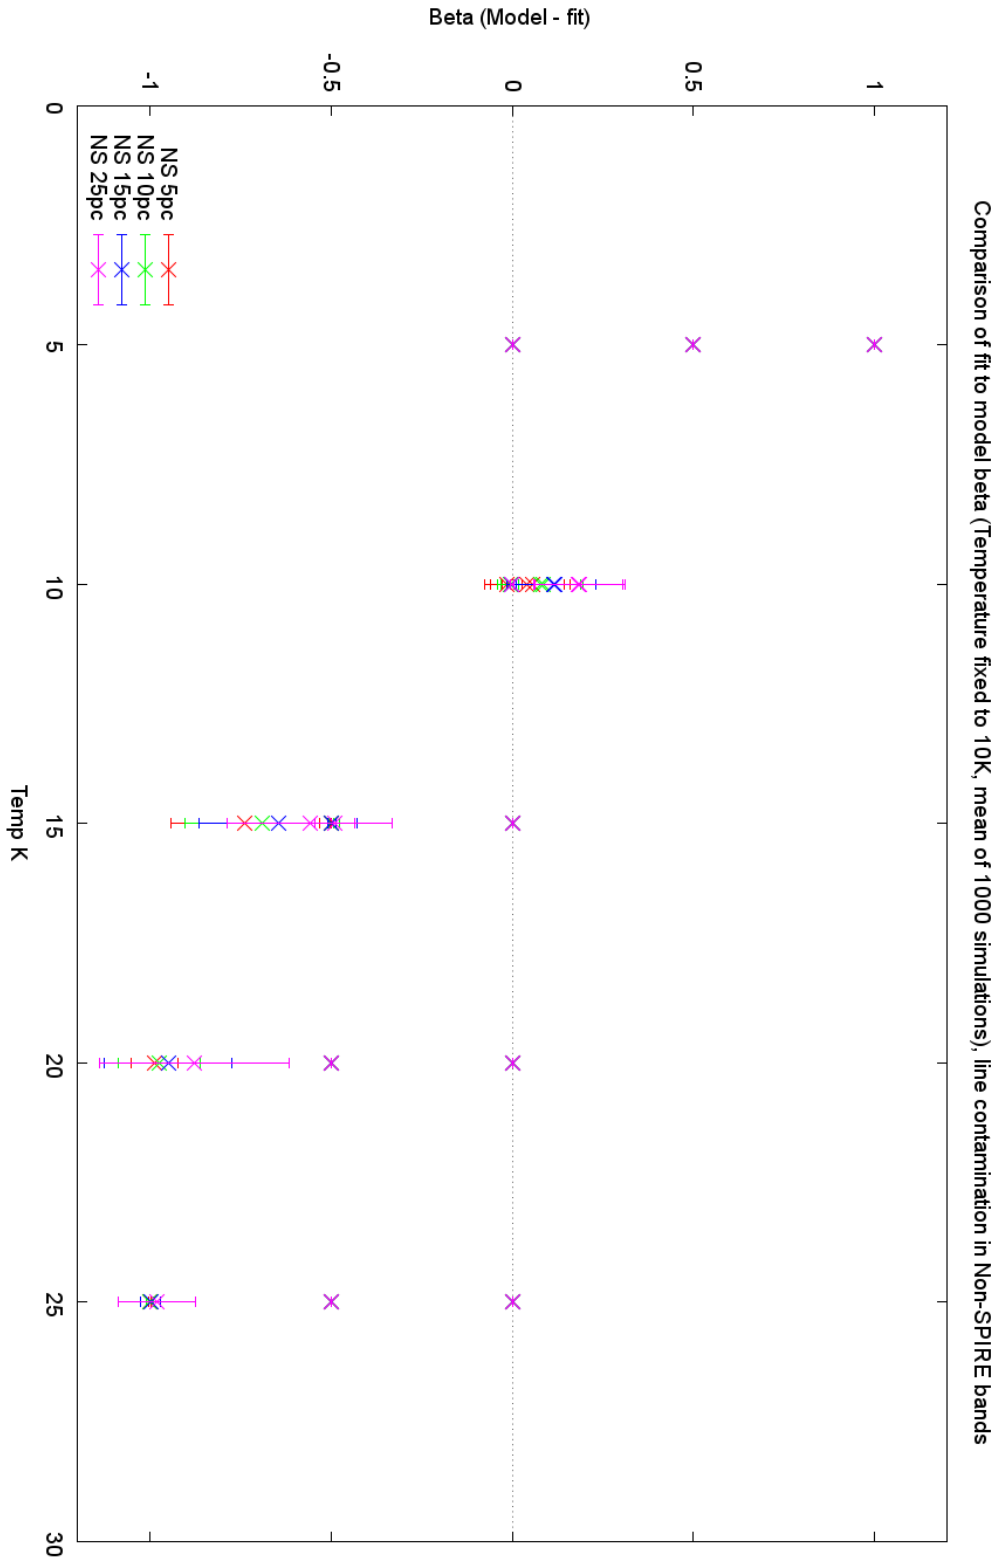

Figure 5.51: Line contamination with $T_{rec} = 10K$ fixed across non-SPIRE bands, effect on β_{rec}, grouped by T_{model}.

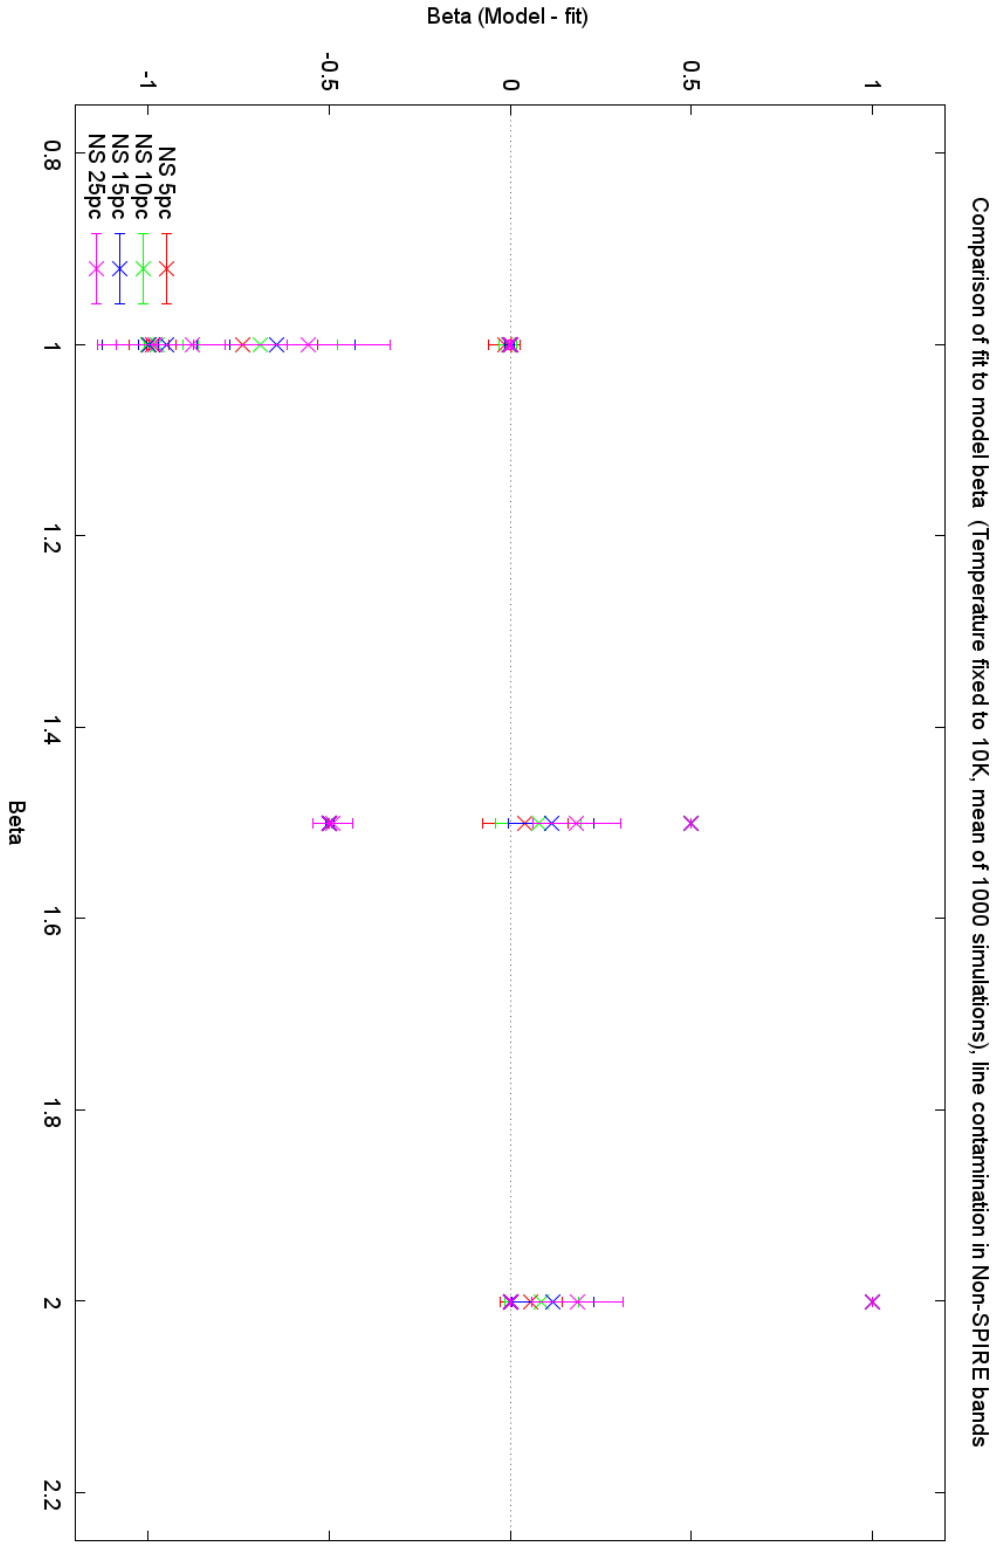

Figure 5.52: Line contamination with $T_{rec} = 10K$ fixed across non-SPIRE bands, effect on β_{rec}.

For the fixed 10K fits, we obtain Figures 5.51 and 5.52. Comparing Figure 5.51 to 5.47, we see that fixing the temperature in this case has had less effect on the recovery of β_{rec}. The tendency of β_{rec} to be recovered at the fit limits is reduced, with the 15K greybodies showing improved recovery of β_{rec} compared to the case where the line contamination is in the SPIRE bands. When grouped by β_{rec}, there is little to distinguish the SPIRE and non-SPIRE line contamination cases. The recovery of β_{rec} at the higher end of its range is adversely affected, although the degree of the change is small.

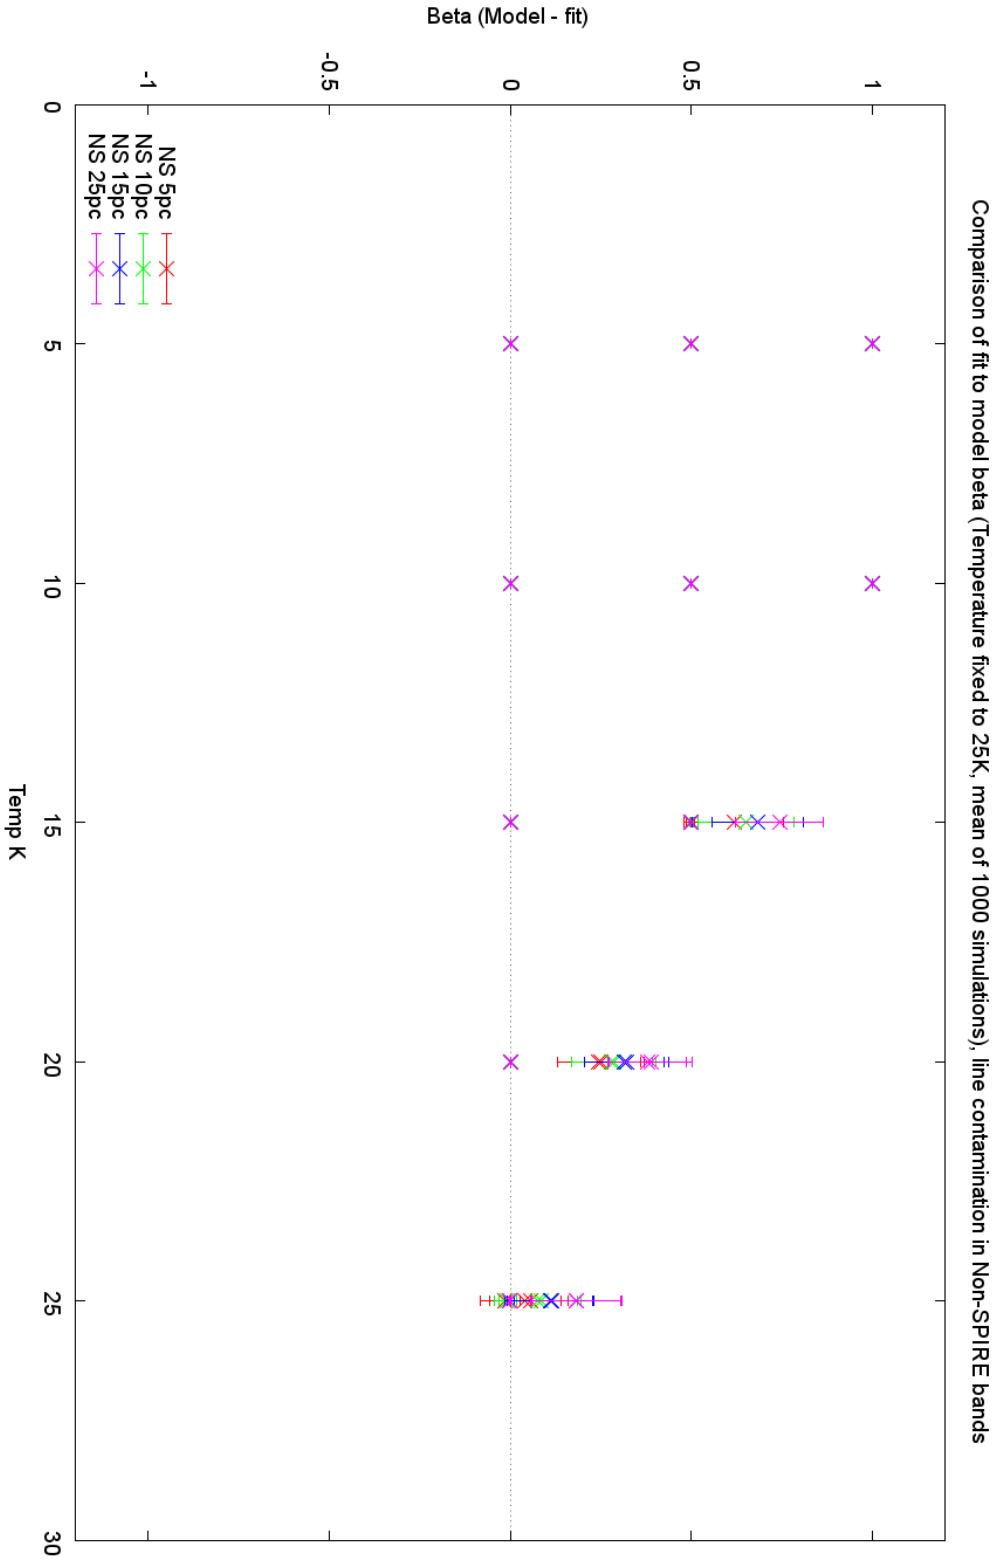

Figure 5.53: Line contamination with $T_{rec} = 25K$ fixed across non-SPIRE bands, effect on β_{rec}, grouped by T_{model}.

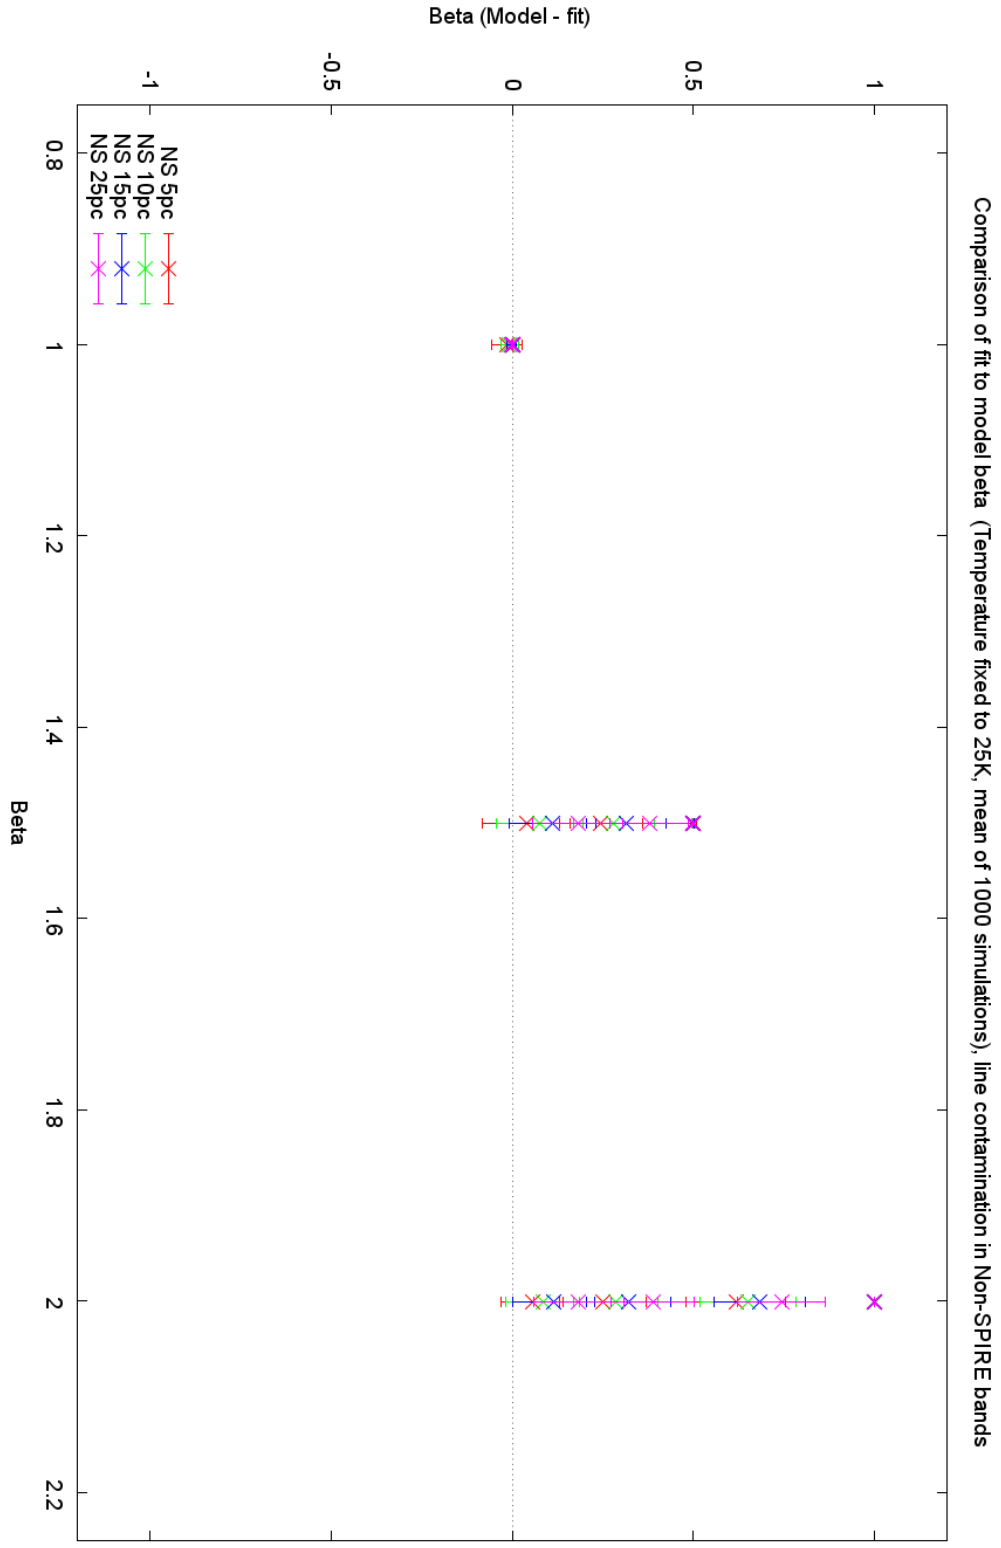

Figure 5.54: Line contamination with $T_{rec} = 25K$ fixed across non-SPIRE bands, effect on β_{rec}.

For the fixed 25K fits, we obtain Figures 5.53 and 5.54. Comparing Figure 5.53 to 5.49, again we see that β_{rec} is recovered low for temperatures below 15K, which is close to the lower fit limit. After 15K we see an increasing tendency for β_{rec} to be correctly recovered, with higher β_{model} greybodies having less tendency for the fit limit to be reached. At 25K, the recovered β_{rec} is still low, but now the fit limit is no longer reached.

Figure 5.54 shows that in this case, it is the lower β_{model} value greybodies that are successfully recovered. As β_{model} increases, the value of β_{rec} recovered tends to lower values. Compared to the equivalent figure for the SPIRE-band line contamination (Figure 5.50) we see an improved recovery of β_{rec}.

5.6.3.3 Fixed Temperature cases - a closer examination

For global SPIRE contamination, we can examine the effect of increasing line contamination on β_{rec} recovery (Figures 5.55 to 5.58) in closer detail. We can then examine the effect on recovery on non-SPIRE waveband line contamination (Figures 5.59 to 5.62).

Figures 5.63 to 5.70 use χ as a metric for quality of recovery, where we use the metric to give a measure of how good the overall fit is. We start by examining how β_{rec} recovery is affected by line contamination when the fit temperature is 10K. This is a case where the greybody peak lies within our spectral range.

Figures 5.55 and 5.56 show the effect of SPIRE-band contamination. In all line contamination percentages we see β_{rec} is driven to the lower fit limit of 1.0 for the 5K cases. At 10K, the same as the fixed fit limit, recovery of β_{rec} is hardly affected by line contamination. Low β_{model} simulations return a β_{rec} slightly higher

than the original and vice versa. For 5% the effect on β_{rec} is 0.05(1.0)/0.02(2.0) and for 25% it is 0.17(1.0)/0.01(2.0). Thus as line contamination increases, lower β_{model} simulations at this temperature show β_{rec} driven slightly higher – until we are at the upper β_{rec} level where it cannot increase any more.

As the temperature increases beyond the fixed fit limit, we see β_{rec} consistently at the upper level, with only slight variations. For 25% line contamination it is always at the limit, though for 5% the lower β_{model} value of 1.0 shows a small deviation from the upper limit at 15K and 20K although this is only 0.18(15K) and 0.01(20K).

Figures 5.59 and 5.60 show the effect of non-SPIRE band line contamination. Largely, the results mirror that of SPIRE band contamination. For 5K greybodies all line contamination percentages drive β_{rec} to the lower limit, and again for 10K we see recovery of β_{rec} values close to β_{model}. Here β_{rec} is driven slightly high where it is at 1.0 (0.2 for 5% and 0.0002 for 25%) and low where it at 2.0 (0.05 at 5% and 0.19 at 25%). Here, as expected increasing line contamination shows an increasing deviation in the recovered β_{rec} from the initial simulation value. Compared to the SPIRE case, temperatures above 15K do not show β_{rec} immediately driven to the upper limit when recovered. For 5% line contamination, β_{rec} is driven to 2.0 when β is originally 1.5 and 2.0 from 15K onwards. For the lower value of β_{model}, it is consistently recovered close to 2.0 but not quite at the fit limit, with a deviation from this limit of 0.24 (15K). At 25K there is no deviation.

As line contamination increases, we see at 15K β_{rec} is driven high (for example where $\beta_{model} = 1.0$, $\beta_{rec} = 1.56$ for 25% contamination). The degree by which it is driven high is lower for the 25% simulations than the 5%. At 5K above our fit limit, we see increasing line contamination does not give answers that are easily

interpreted. At 20K and 25K the results are similar to the SPIRE case, though again for $\beta_{rec} = 1.0$ we see the degree to which β is driven high is lower for increasing line contamination.

Figures 5.63 and 5.64 show the quality of fit for SPIRE line contamination. The large slope leading to the 10K data shows the poor quality of fit across the board for the 5K modelling. χ for the 5K data is around a hundred times the value of the equivalent 10K models. As the temperature increases, the fit quality from 15K to 25K degrades but even for the 25K modelling χ is less than half the 5K value.

The data above 10K show that the better fits are for the low β_{model} simulations, and it is the lower line contamination models that show a better fit than the higher line contamination ones. This slope is seen to a lesser extent in the 5K data. It is only at 10K that $\chi \leq 10$, the value that χ has for the earlier free-parameter and fixed-beta models. Fitting a greybody with a peak in the spectral range to data without a peak is always going to result in a poorer fit. However, the very poor quality of the 5K fit compared to using a 10K model is more surprising as both have a peak in the spectral range. The fact the 15K doesn't show such a poor quality is likely because at 15K the peak is only just outside our spectral range ($200\mu m$ is the peak of a blackbody at 15K) and as such between 10K and 15K the fit improves simply because with only a few datapoints we cannot constrain the peak too strongly.

Figures 5.67 and 5.68 show the effect of non-SPIRE band contamination and mirrors the previous SPIRE band line contamination figure for χ. There are some slight differences to be seen however. Compared to the SPIRE band figures, the value for χ is slightly higher for non-SPIRE contamination than for SPIRE contam-

ination. This is not a significant increase however. At 5K and 10K, the non-SPIRE contamination shows a slope such that the lower β_{model} models have a slightly better fit than the higher β_{model} models. The SPIRE-band models agree for 5K but at 10K it is the other way around. For temperatures of 15K and above, the fit degrades but overall χ increases at a slower rate than in the SPIRE models. This implies that for this case non-SPIRE line contamination leads to fewer issues with fitting. At 25% line contamination for 15K and above, χ is roughly half that of the SPIRE band values and for 5% it is roughly 75%. One sole exception is the 25% line contamination for the 15K and $\beta_{model} = 1$ model where the non-SPIRE case is worse than the SPIRE case.

We can now examine cases where the original greybody temperature is 25K, and hence the greybody peak is outside our spectral range. This is interesting as this means we lose the peak as a feature to fit. Here we expect increasing line contamination to show increasing degradation in the fit. Figures 5.57 and 5.58 show how β_{rec} behaves for this case for SPIRE band line contamination. Across the line contamination range, we see for the lower temperature greybodies, 5K and 10K, $\beta_{rec} = 1.0$ regardless of the simulation value. For 15K, at 5% line contamination the recovery of β_{rec} is towards the low end of its range (recovered as 1.0 for an initial 1.0 and 1.45 for an initial 2.0) and for 25% this hardly changes (recovered as 1.0 for an initial 1.0 and 1.57 for an initial 2.0).

For 20K we see β_{rec} recovery improve, and the degree of improvement itself improves as the line contamination percentage increases. For 20K, 5% line contamination recovered β low by about 0.18, except for the $\beta_{model} = 1.0$ simulation as this is a lower limit and the increase in β_{rec} is marginal. At 25% we see the amount by which it is recovered low is only 0.06. At 25K the opposite happens. β_{rec} is recovered more accurately at low line contamination percentages (as is to

be expected). What this shows is that in the 25K case, the behaviour of β_{rec} is difficult to categorise for lower temperatures with SPIRE band contamination.

Figures 5.61 and 5.62 show what happens with non-SPIRE band line contamination. Here 5K and 10K simulations all show β_{rec} is at the lower fit value for all values of line contamination, and at 15K it is recovered at the level for all but the $\beta_{model} = 2.0$ cases, where it is recovered at 1.38 (5%) and 1.26 (25%). At 20K β_{rec} is still low, 1.00/1.25/1.75 for the 1.0/1.5/2.0 simuations at 5% and 1.00/1.12/1.61 for 25%. At 25K we see the same results as in the SPIRE case, that low line contamination recovers β_{rec} reasonably well and that this degrades as line contamination increases. For non-SPIRE line contamination, the deviation from the original β_{model} is greater than for SPIRE line contamination. For 25% line contamination, the 25K cases recover β_{rec} as 1.00/1.32/1.82 for the non-SPIRE case and 1.16/1.67/1.99 for SPIRE. Hence we see SPIRE line contamination drives β_{rec} up and non-SPIRE line contamination drives it lower.

Figures 5.65 and 5.66 show the quality of fit for SPIRE line contamination where T_{rec} is fixed at 25K. As to be expected, χ reduces significantly as we approach this temperature. For 5% line contamination, we see a slope along constant temperature values develop at 10K that flattens out by 25K. This slope shows that at 10K, the χ value for the $\beta_{model} = 2$ fit is significantly less than the $\beta_{model} = 1$ case, implying that the fit is better for higher values of β_{model}. This slope flattens out by the 25K models. For larger line contamination percentages, this slope is roughly the same, although overall χ values are higher, as expected. In both cases, at the 25K models the slope actually inverts slightly – more so for the higher line contamination percentages. This implies in this case, where the fit and model temperature are the same, that the fit quality is slightly better for the lower β_{model} values.

Figures 5.69 and 5.70 show the effect of non-SPIRE band contamination, and mirrors the previous SPIRE band line contamination value for χ. There are some slight differences to be seen but overall the same conclusions are to be found, except for the 25K case. For the non-SPIRE case, the slope does not invert and so across the temperature range from 10K onwards it is the higher β_{model} values that have the better fit, and in the non-SPIRE case the $\beta_{model} = 1$ models have consistently higher χ (and corresponding worse fits) than the other cases. At 25K the slope is reduced significantly, but not eradicated.

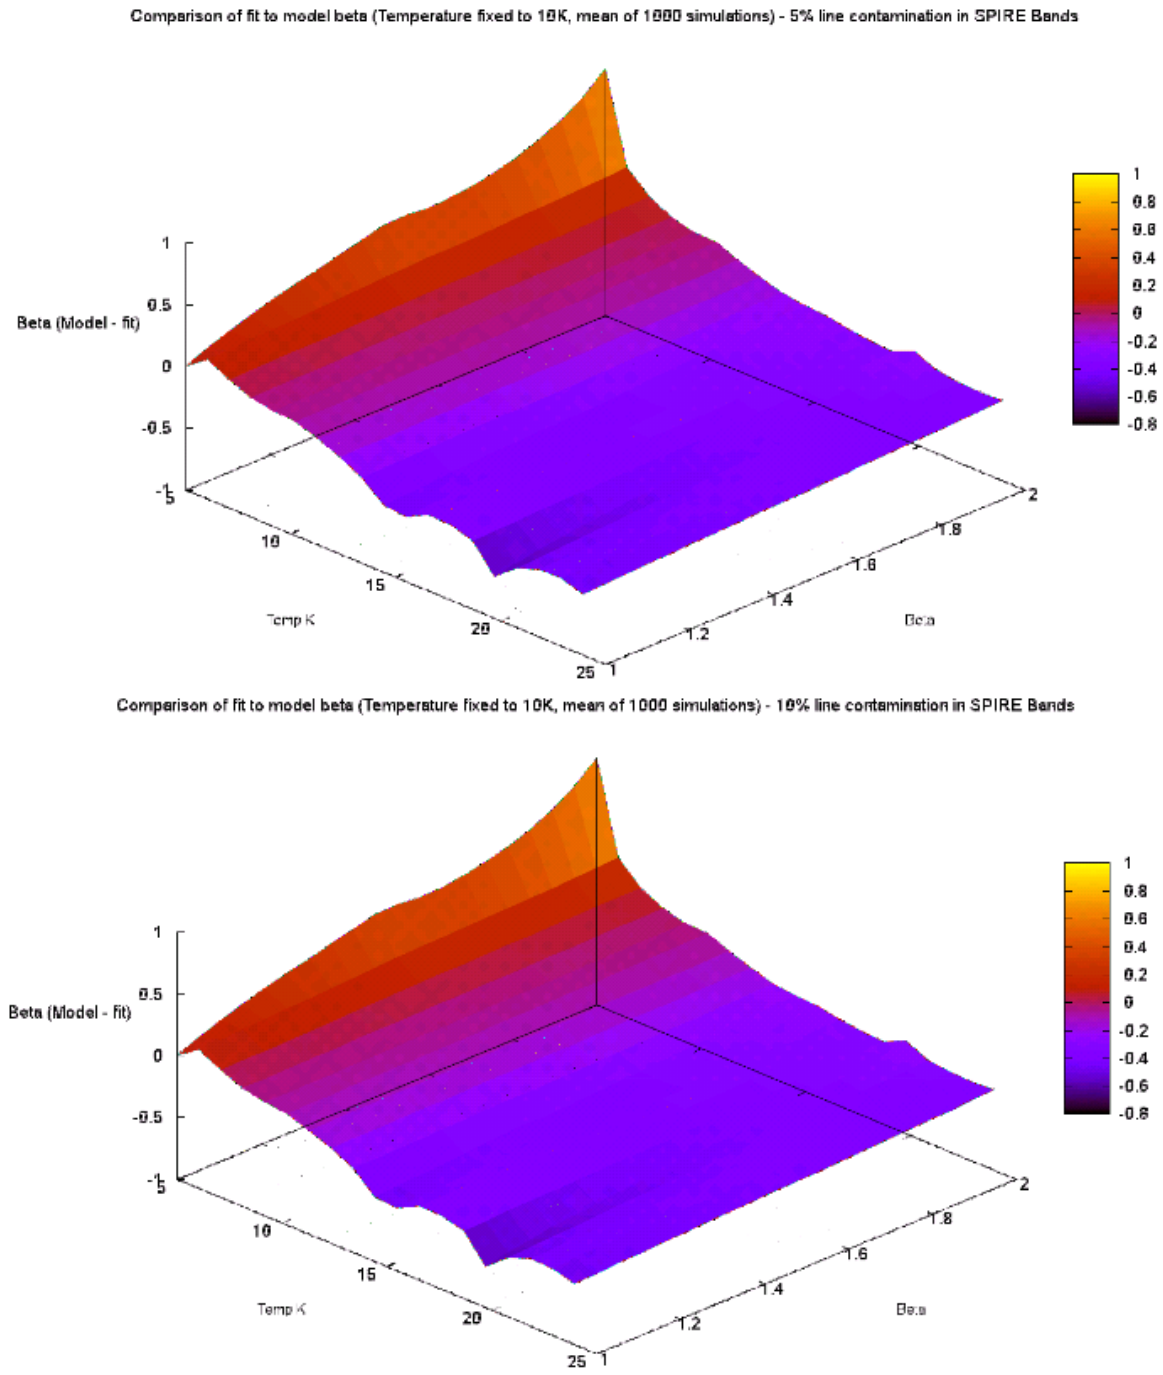

Figure 5.55: Line contamination with $T_{rec} = 10K$ fixed across SPIRE bands, effect on β_{rec}, 5% and 10% line contamination percentage.

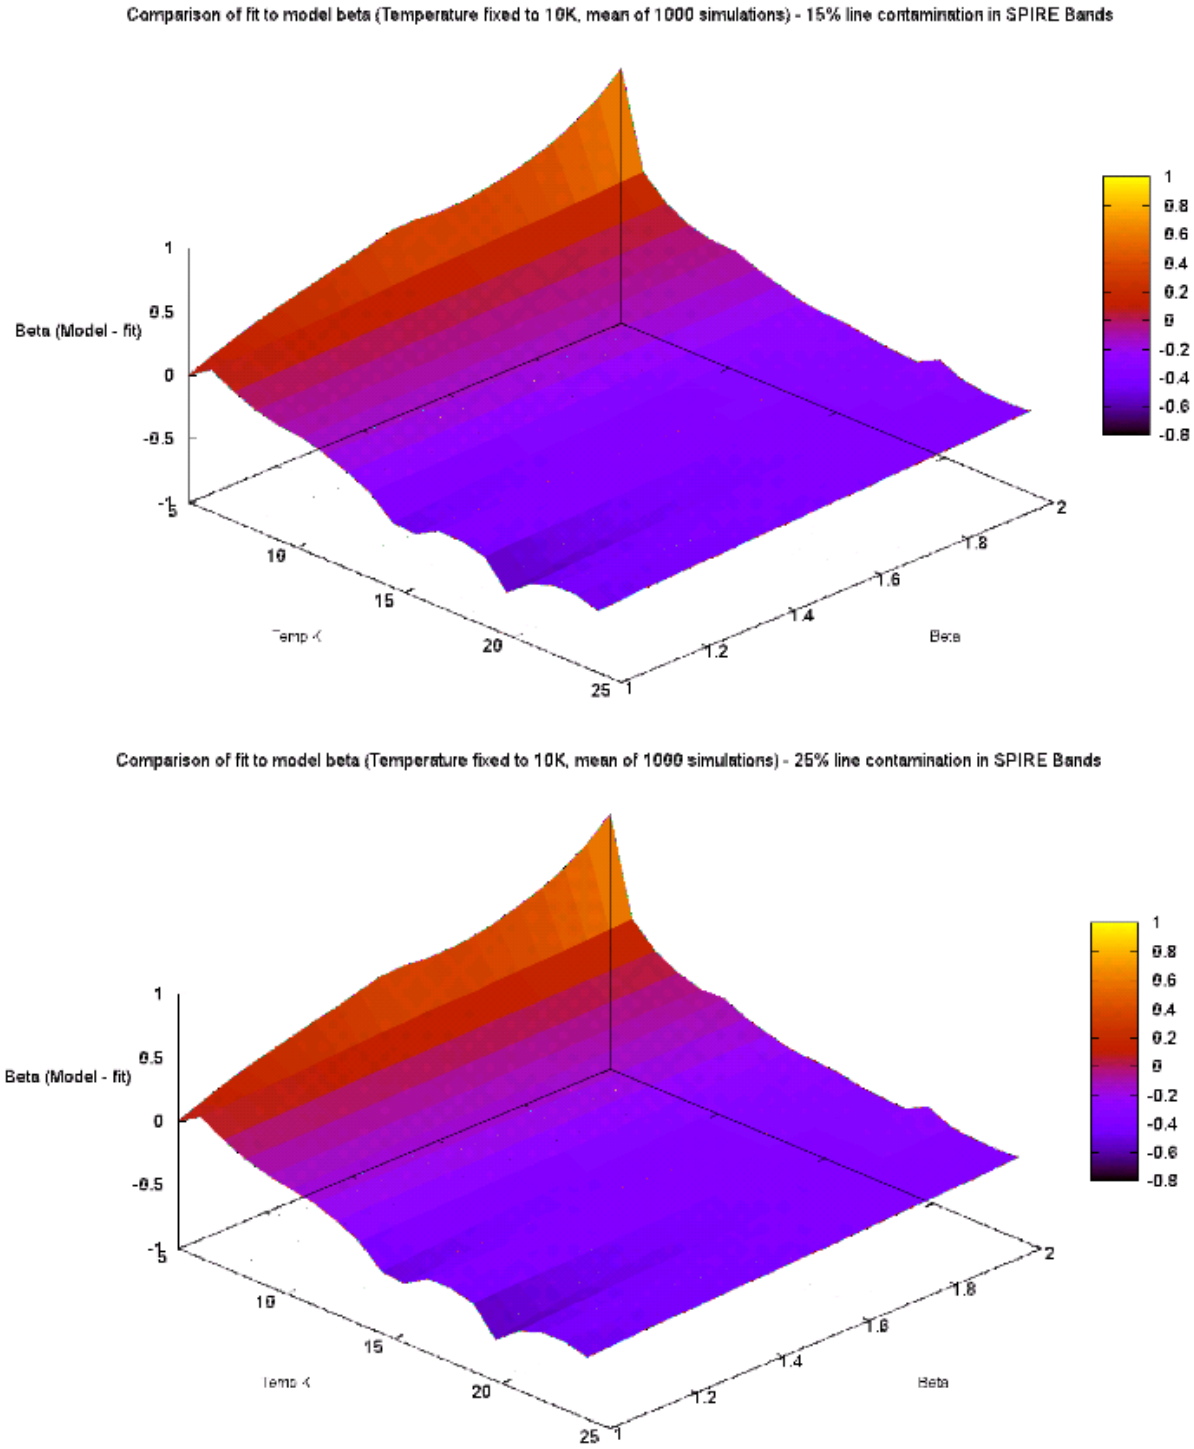

Figure 5.56: Line contamination with $T_{rec} = 10K$ fixed across SPIRE bands, effect on β_{rec}, 15% and 25% line contamination percentage.

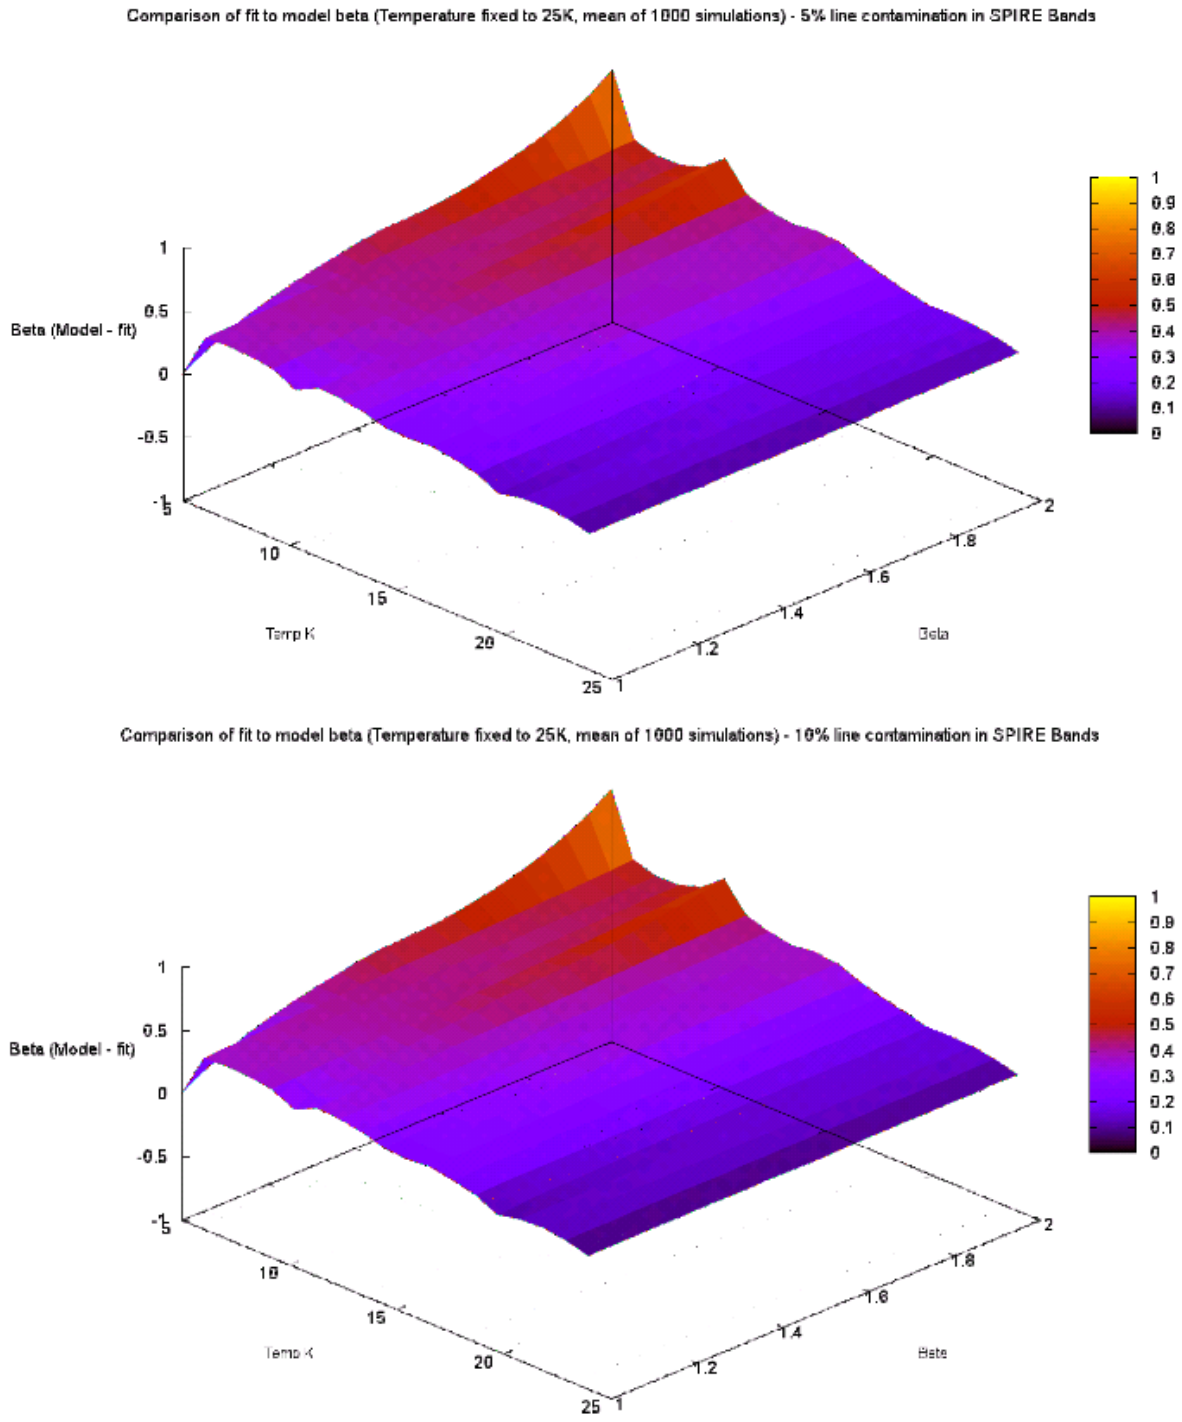

Figure 5.57: Line contamination with $T_{rec} = 25K$ fixed across SPIRE bands, effect on β_{rec}, 5% and 10% line contamination percentage.

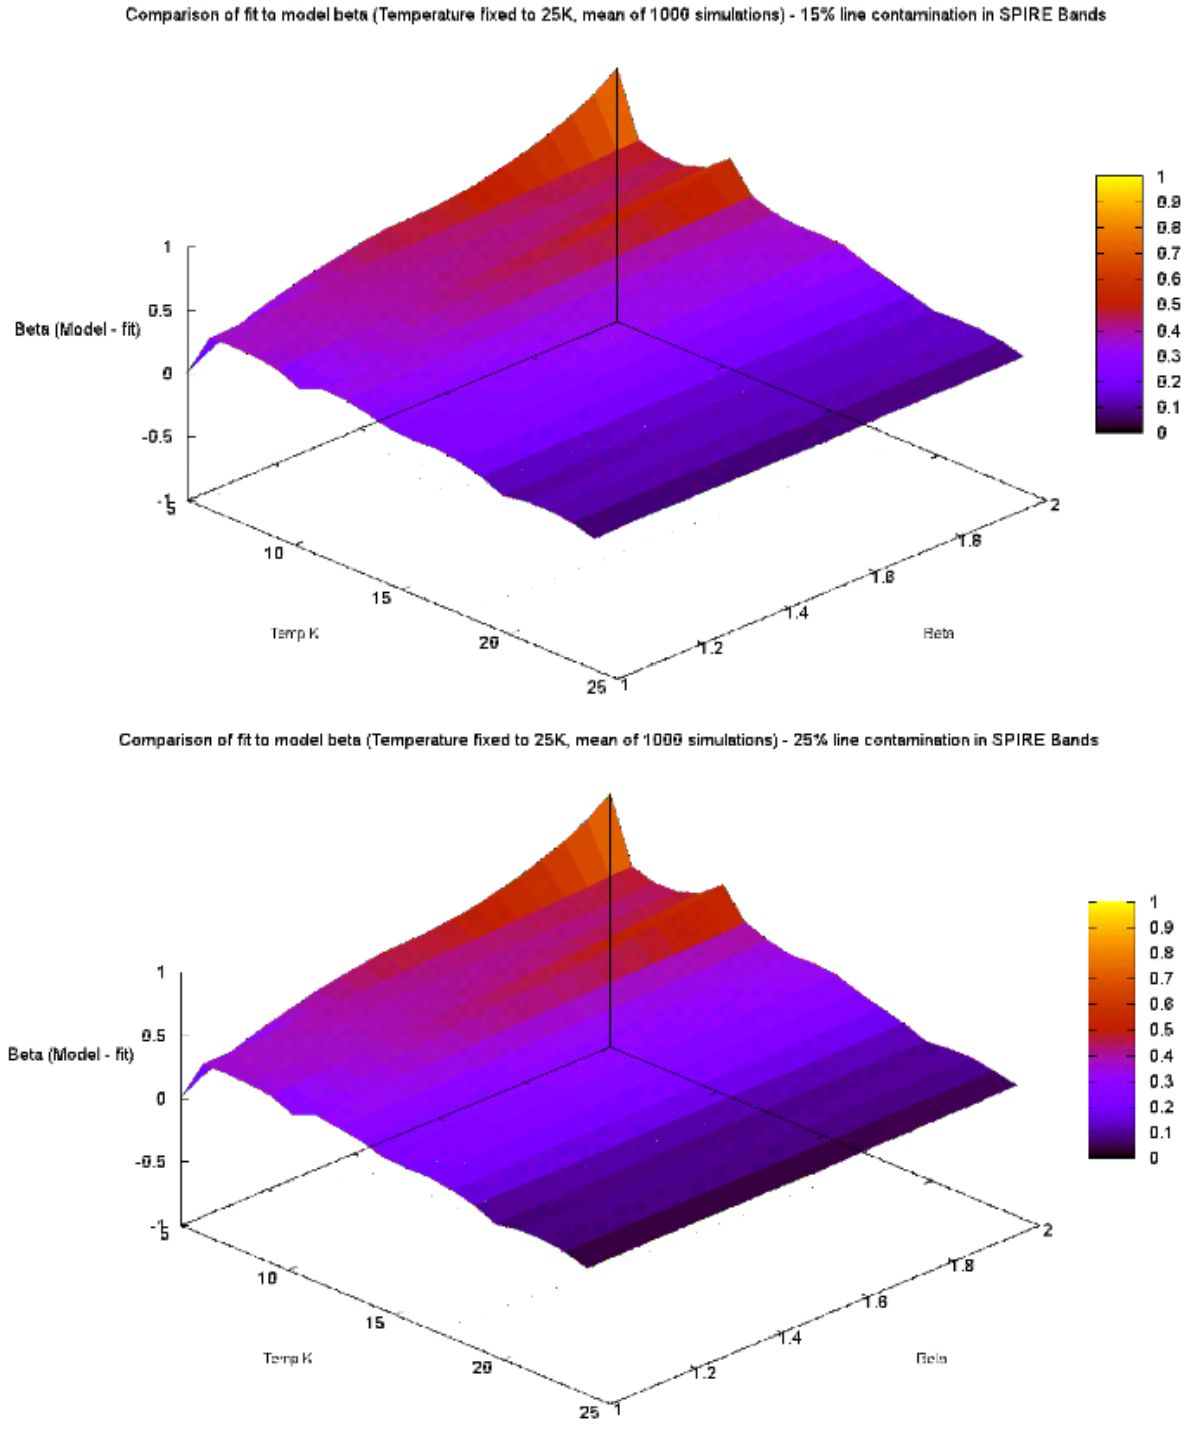

Figure 5.58: Line contamination with $T_{rec} = 25K$ fixed across SPIRE bands, effect on β_{rec}, 15% and 25% line contamination percentage.

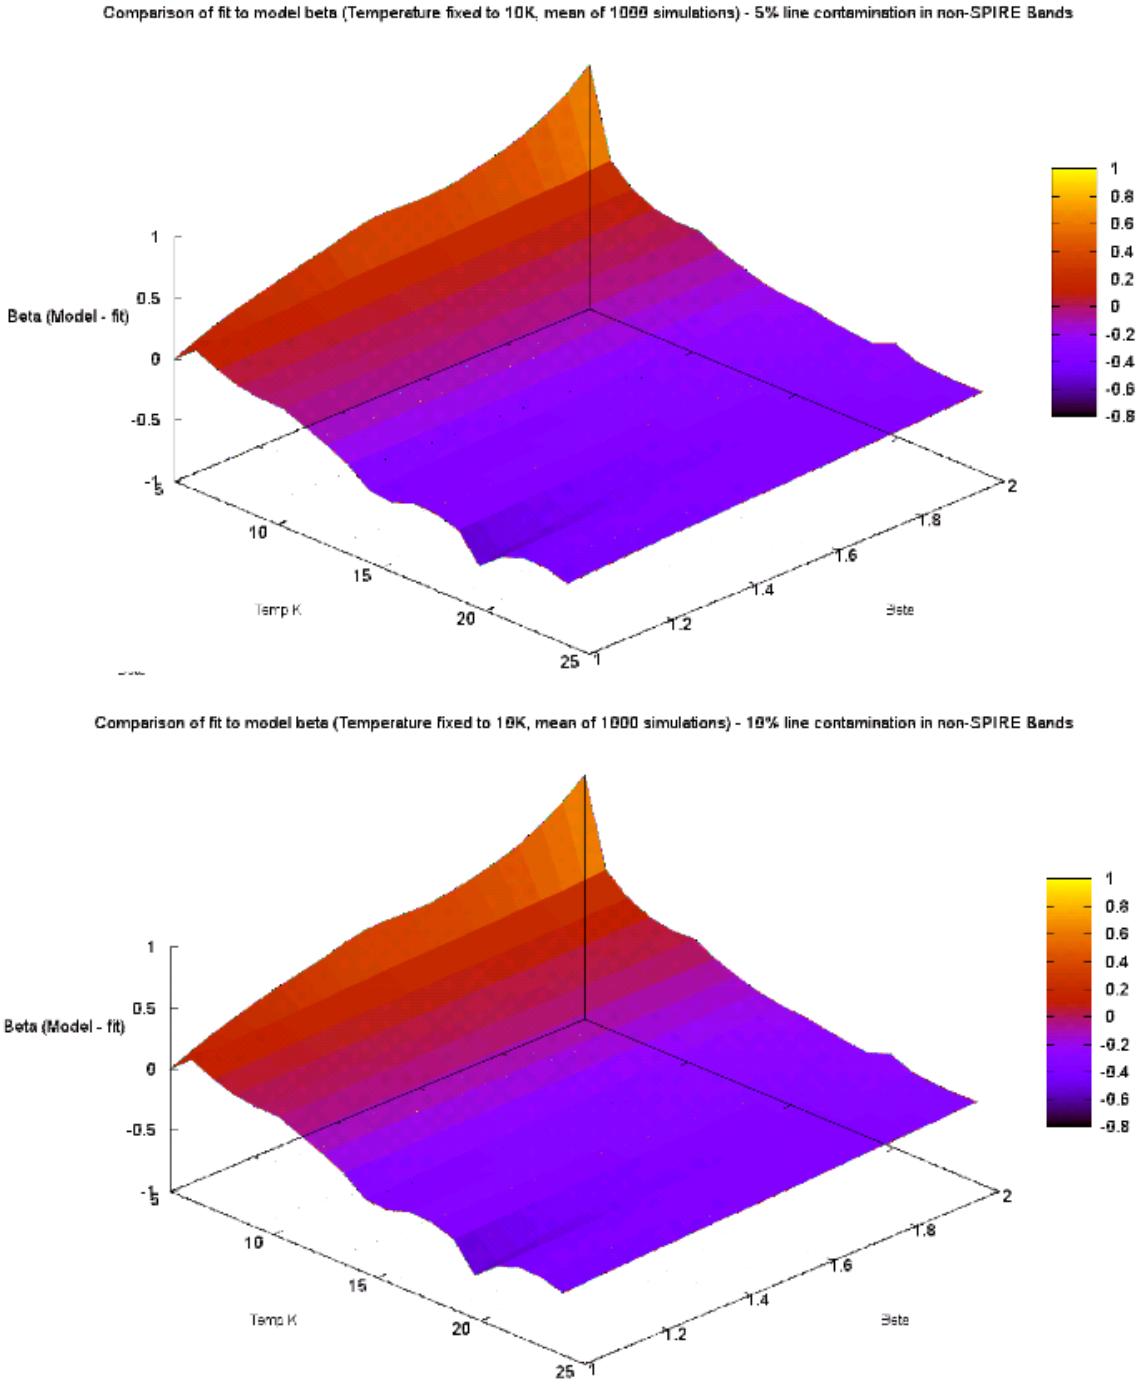

Figure 5.59: Line contamination with $T_{rec} = 10K$ fixed across non-SPIRE bands, effect on β_{rec}, 5% and 10% line contamination percentage.

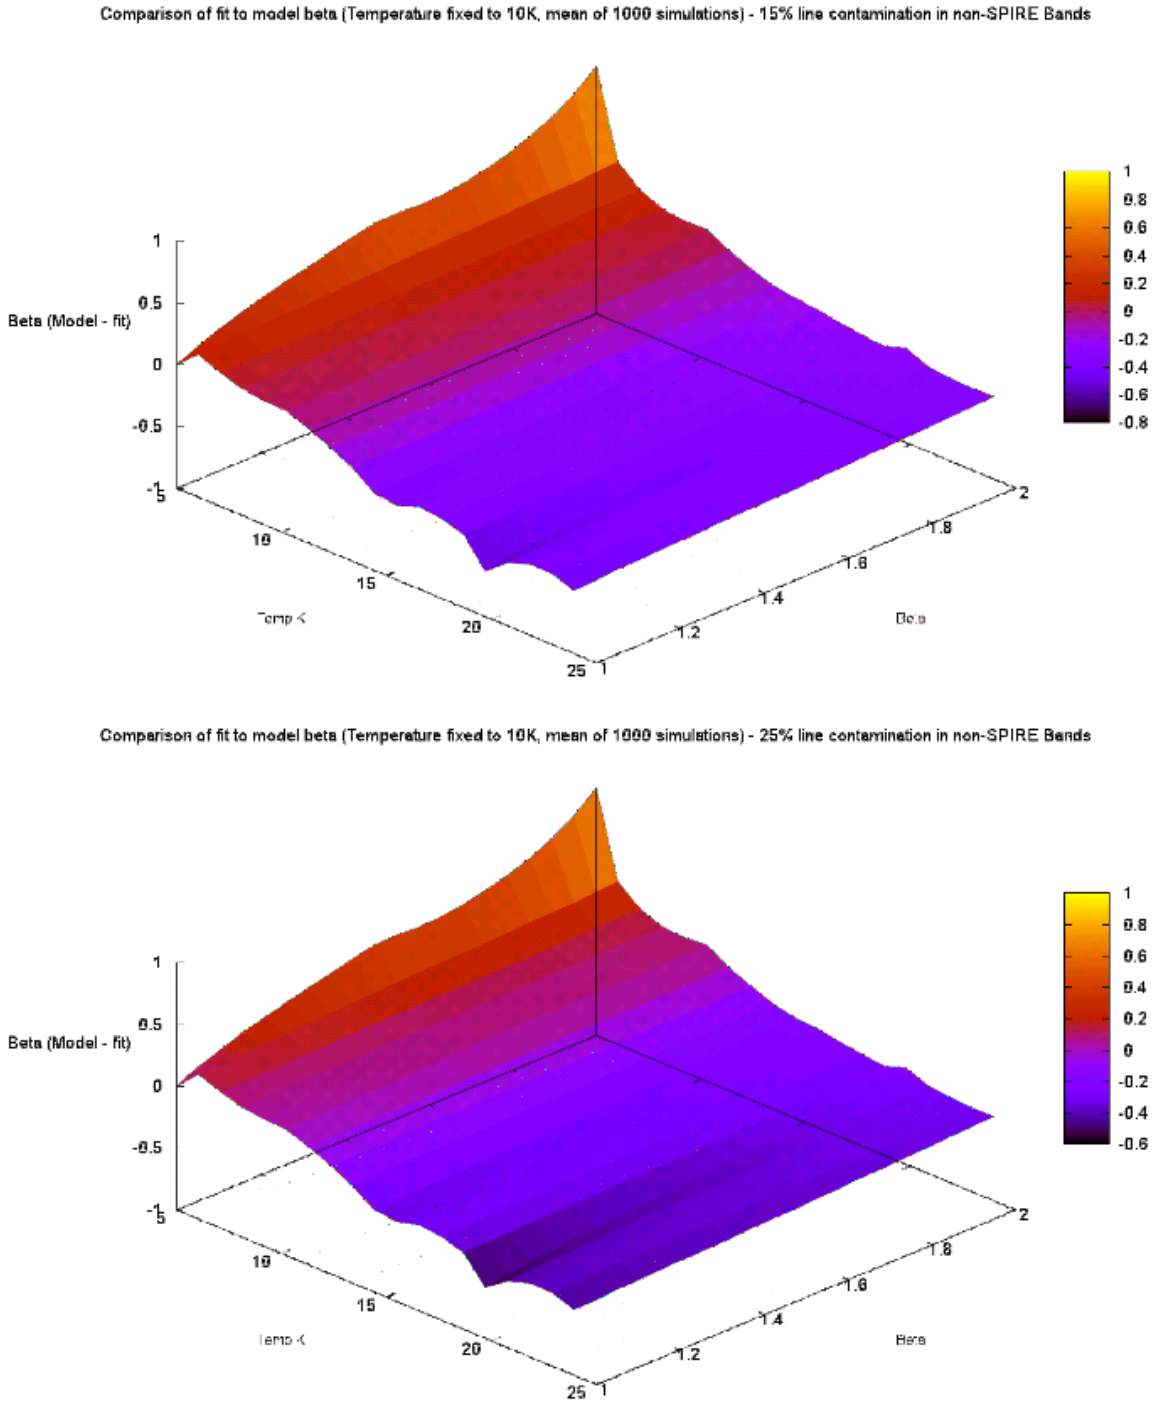

Figure 5.60: Line contamination with $T_{rec} = 10K$ fixed across non-SPIRE bands, effect on β_{rec}, 15% and 25% line contamination percentage.

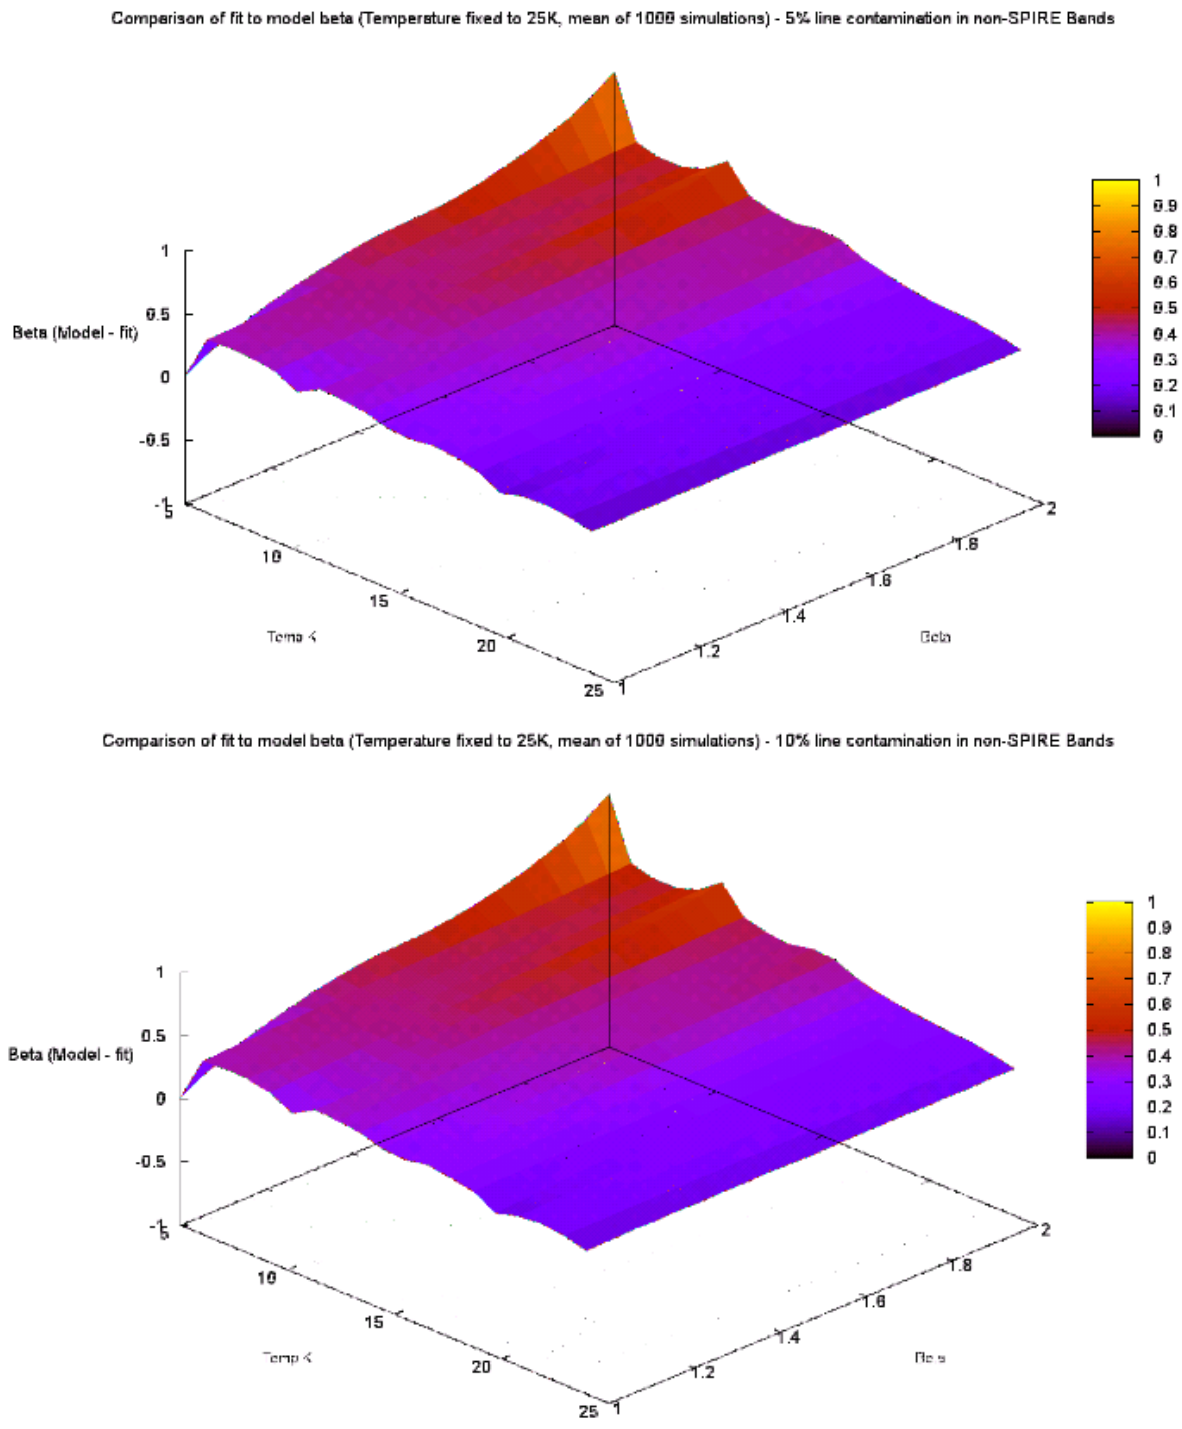

Figure 5.61: Line contamination with $T_{rec} = 25K$ fixed across non-SPIRE bands, effect on β_{rec}, 5% and 10% line contamination percentage.

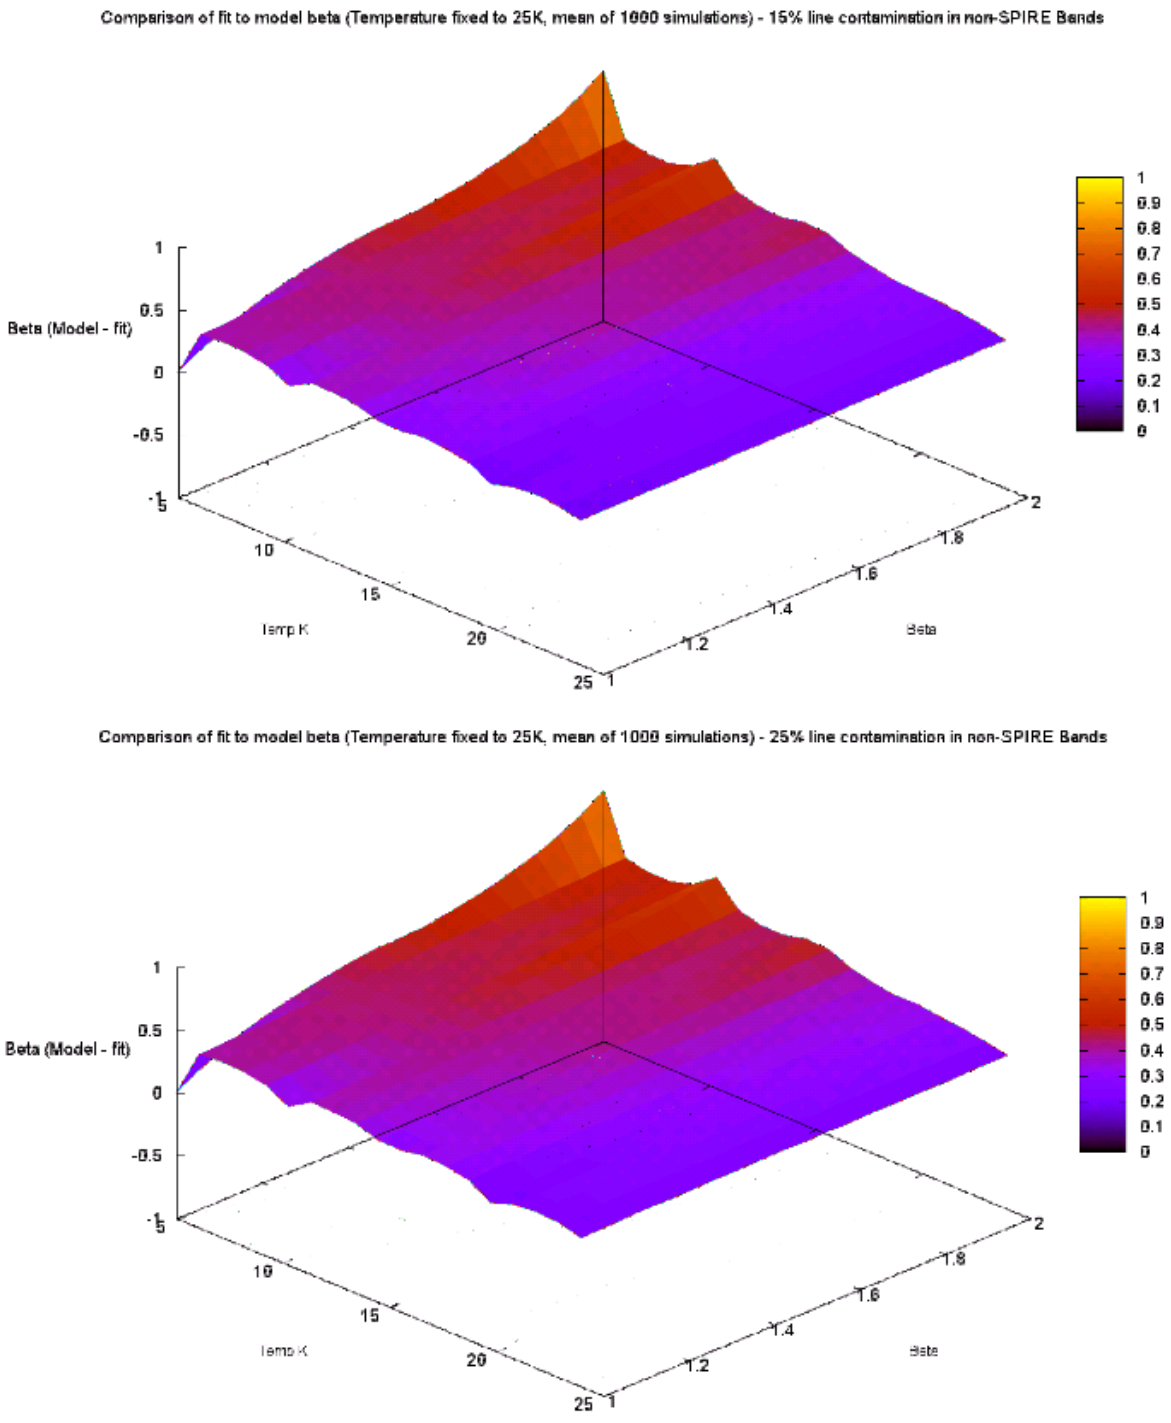

Figure 5.62: Line contamination with $T_{rec} = 25K$ fixed across non-SPIRE bands, effect on β_{rec}, 15% and 25% line contamination percentage.

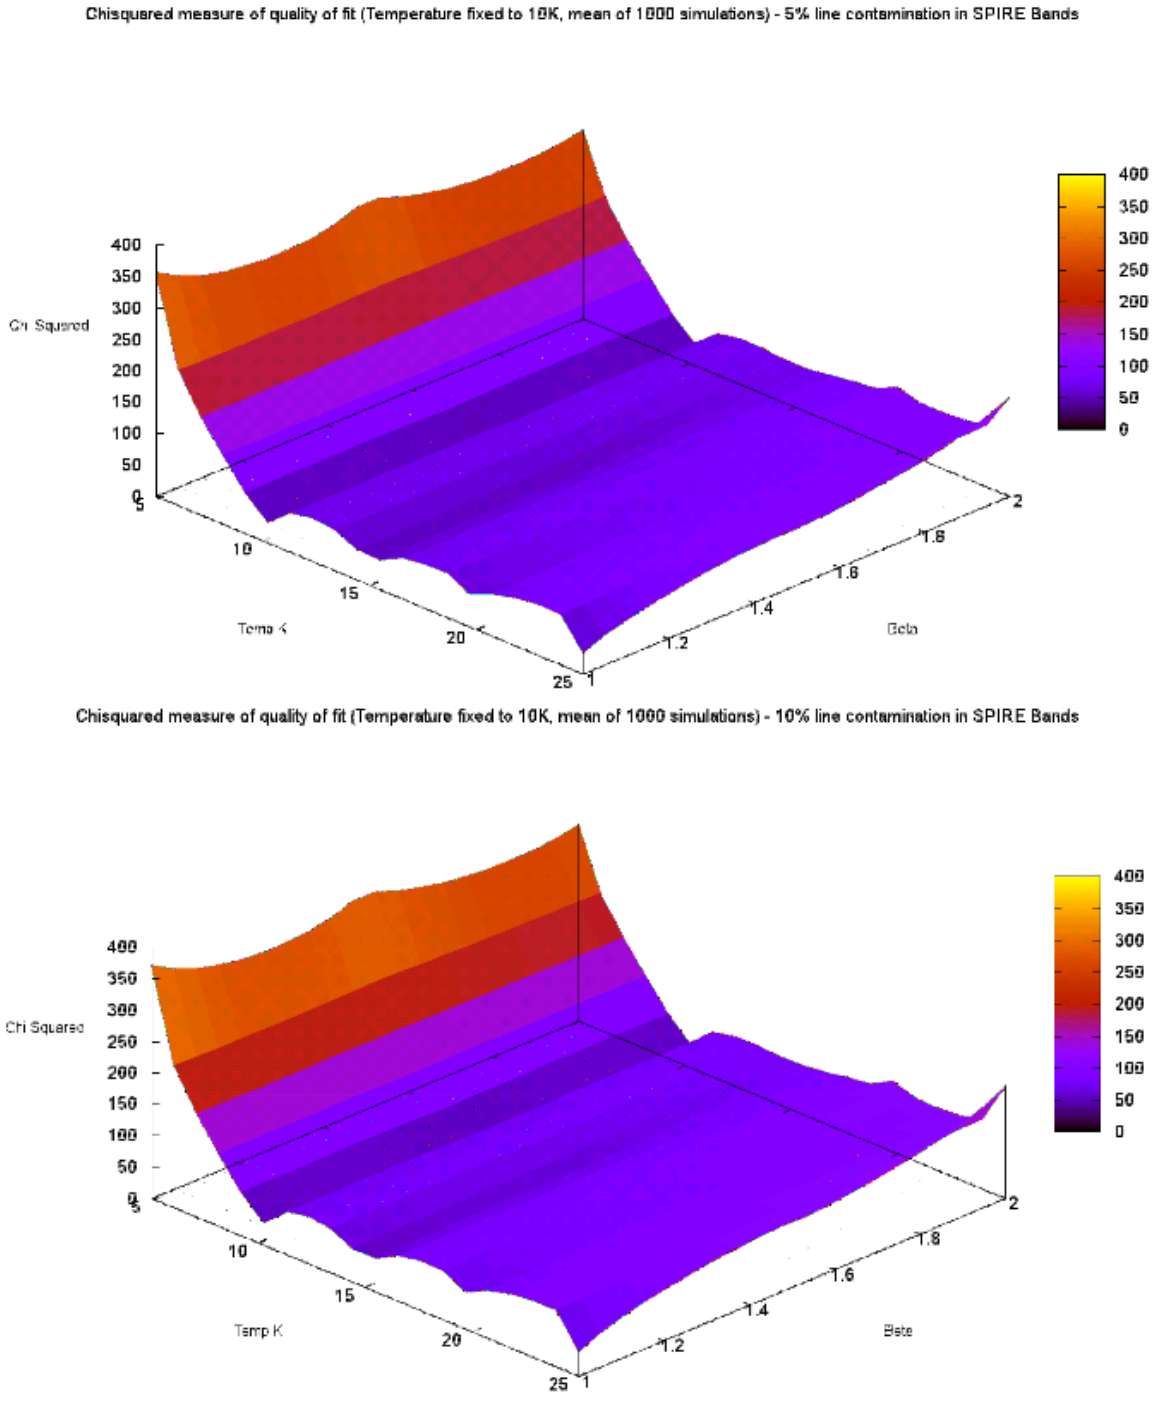

Figure 5.63: Line contamination with $T_{rec} = 10K$ fixed across SPIRE bands χ, 5% and 10% line contamination percentage.

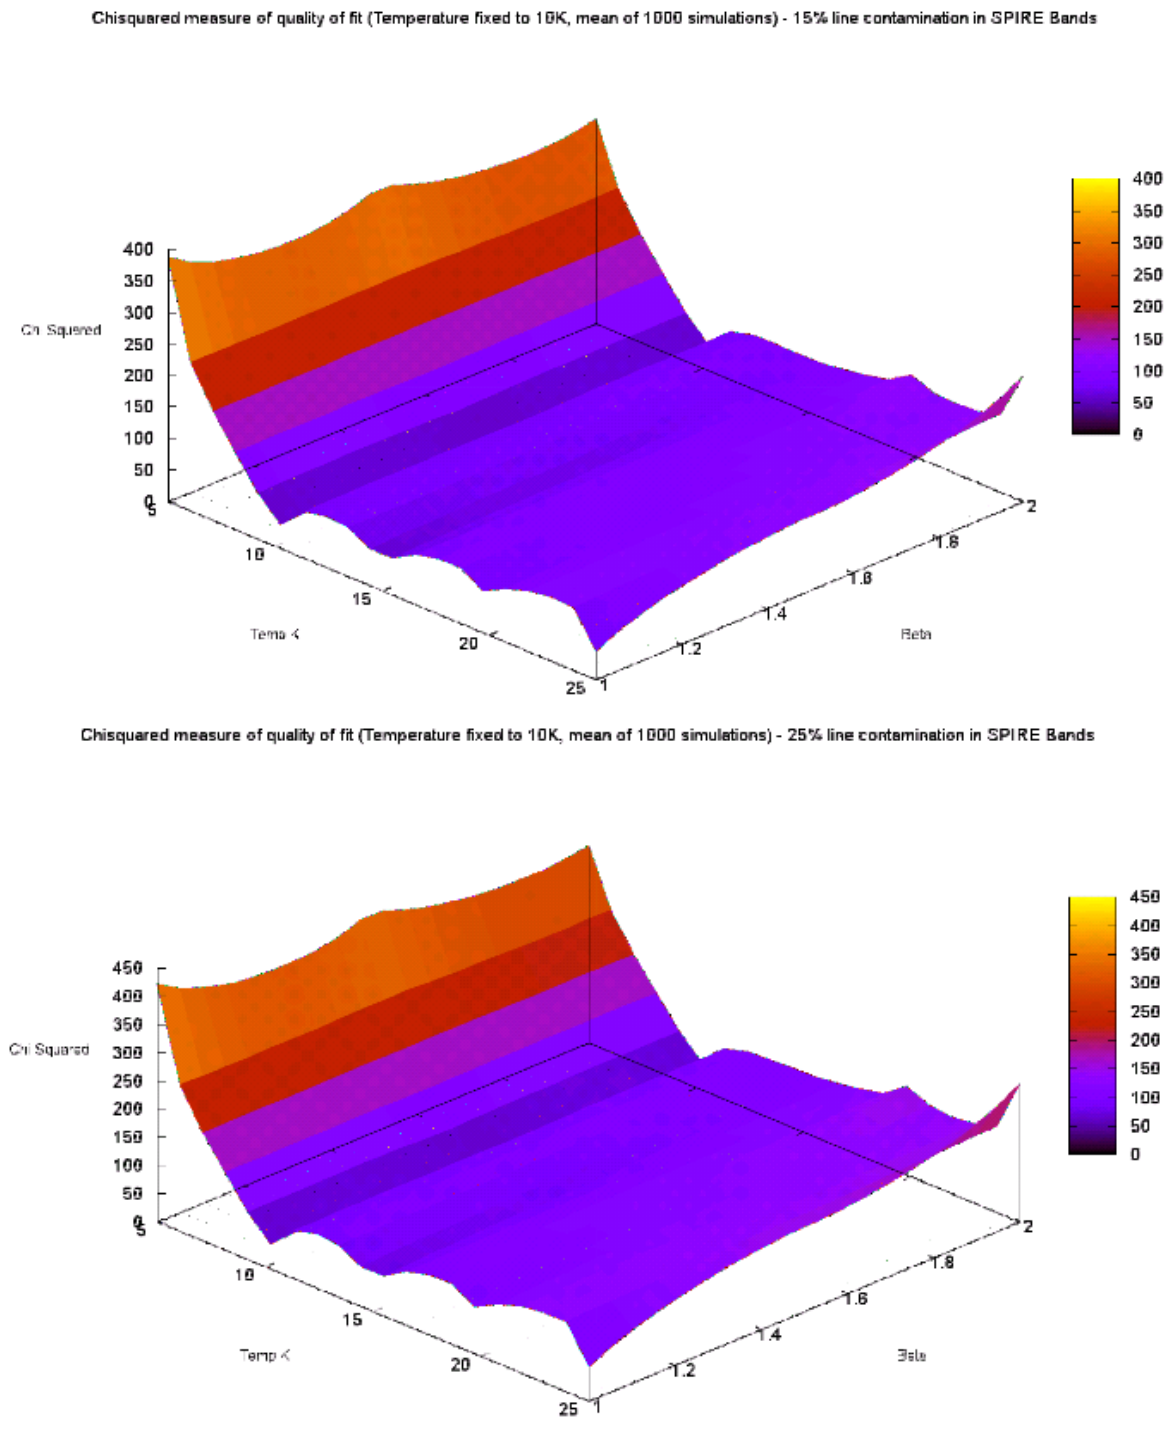

Figure 5.64: Line contamination with $T_{rec} = 10K$ fixed across SPIRE bands χ, 15% and 25% line contamination percentage.

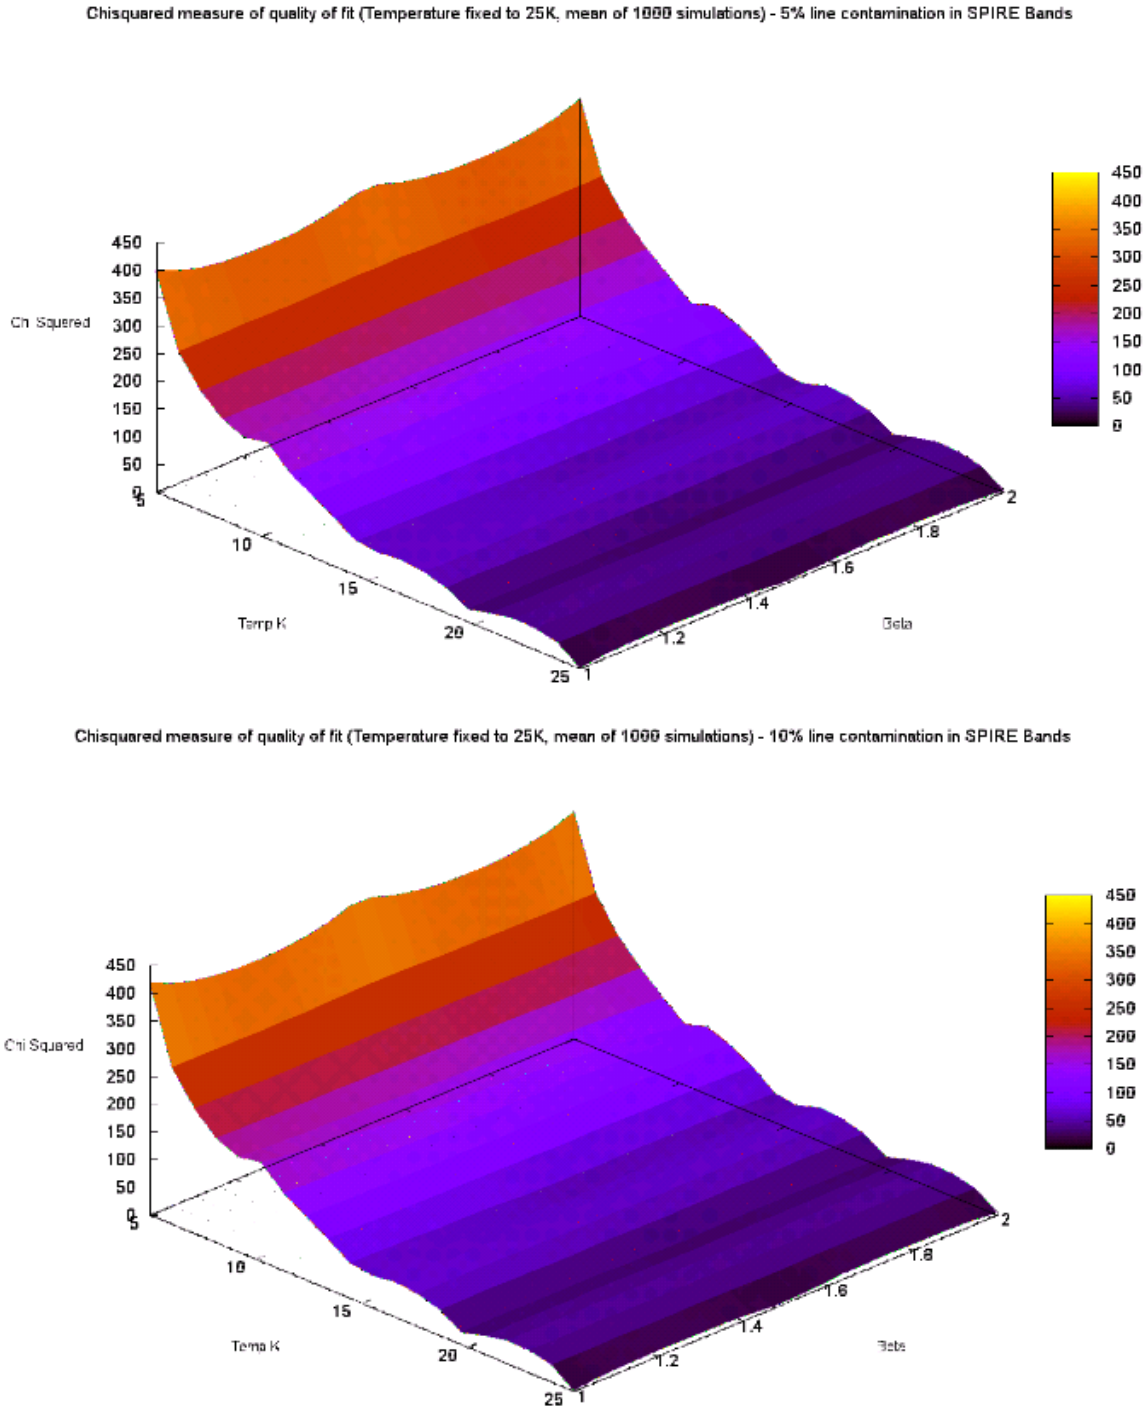

Figure 5.65: Line contamination with $T_{rec} = 25K$ fixed across SPIRE bands χ, 5% and 10% line contamination percentage.

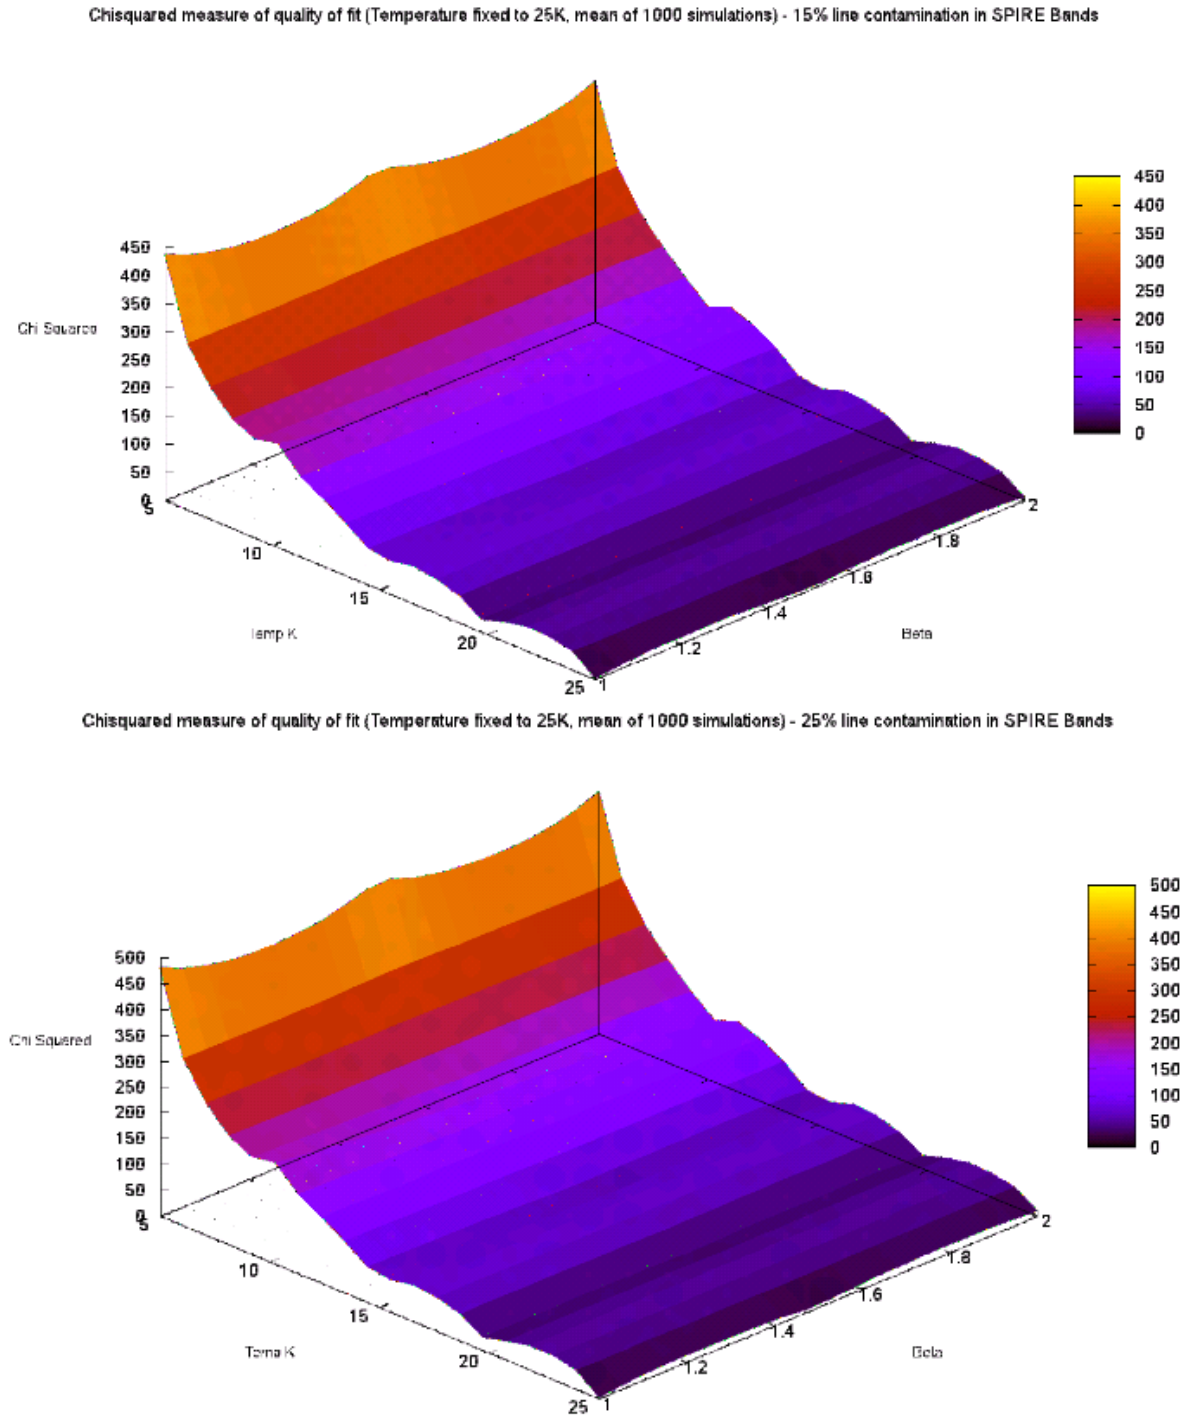

Figure 5.66: Line contamination with $T_{rec} = 25K$ fixed across SPIRE bands χ, 15% and 25% line contamination percentage.

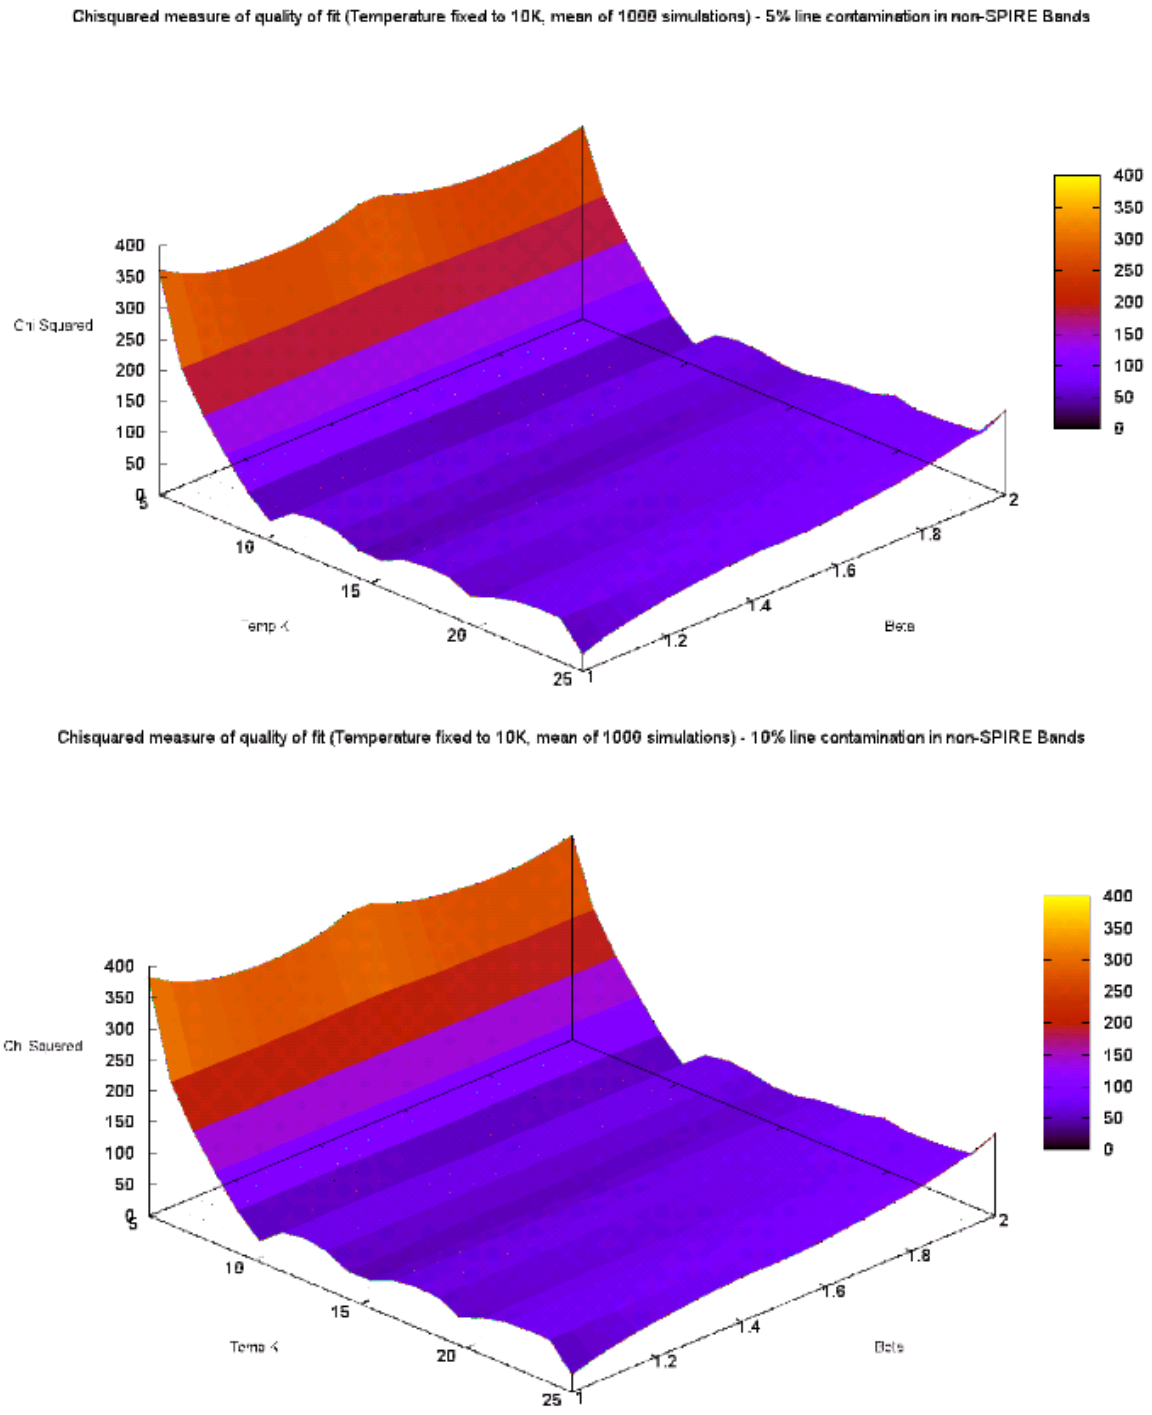

Figure 5.67: Line contamination with $T_{rec} = 10K$ fixed across non-SPIRE bands χ, 5% and 10% line contamination percentage.

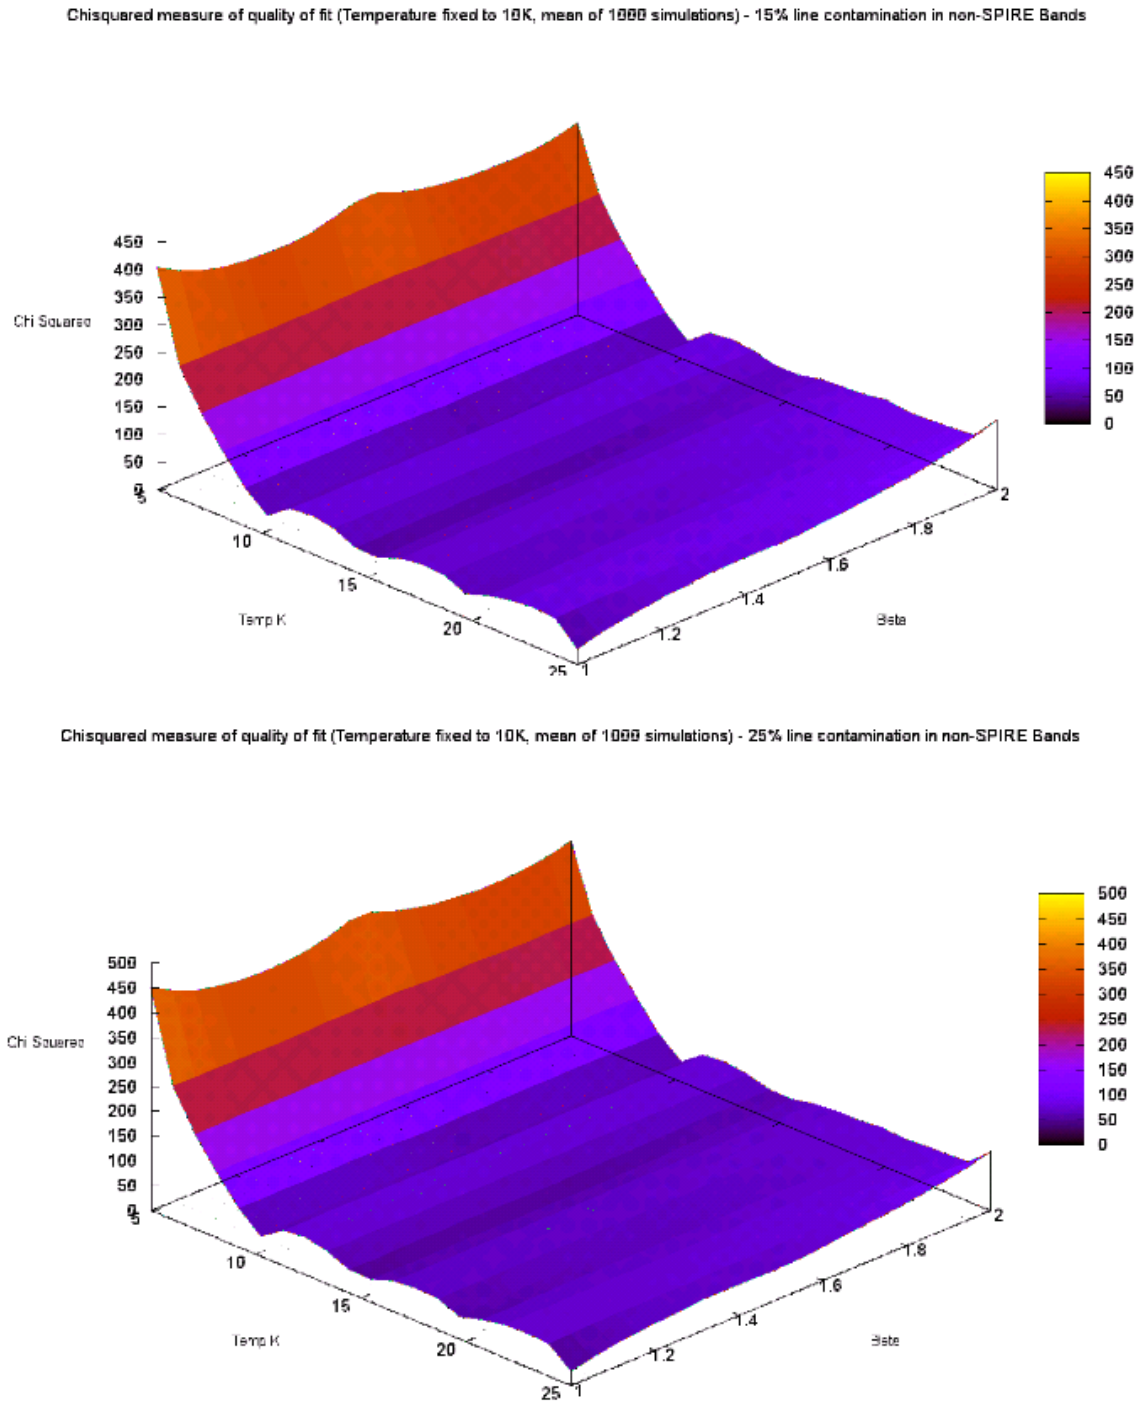

Figure 5.68: Line contamination with $T_{rec} = 10K$ fixed across non-SPIRE bands χ, 15% and 25% line contamination percentage.

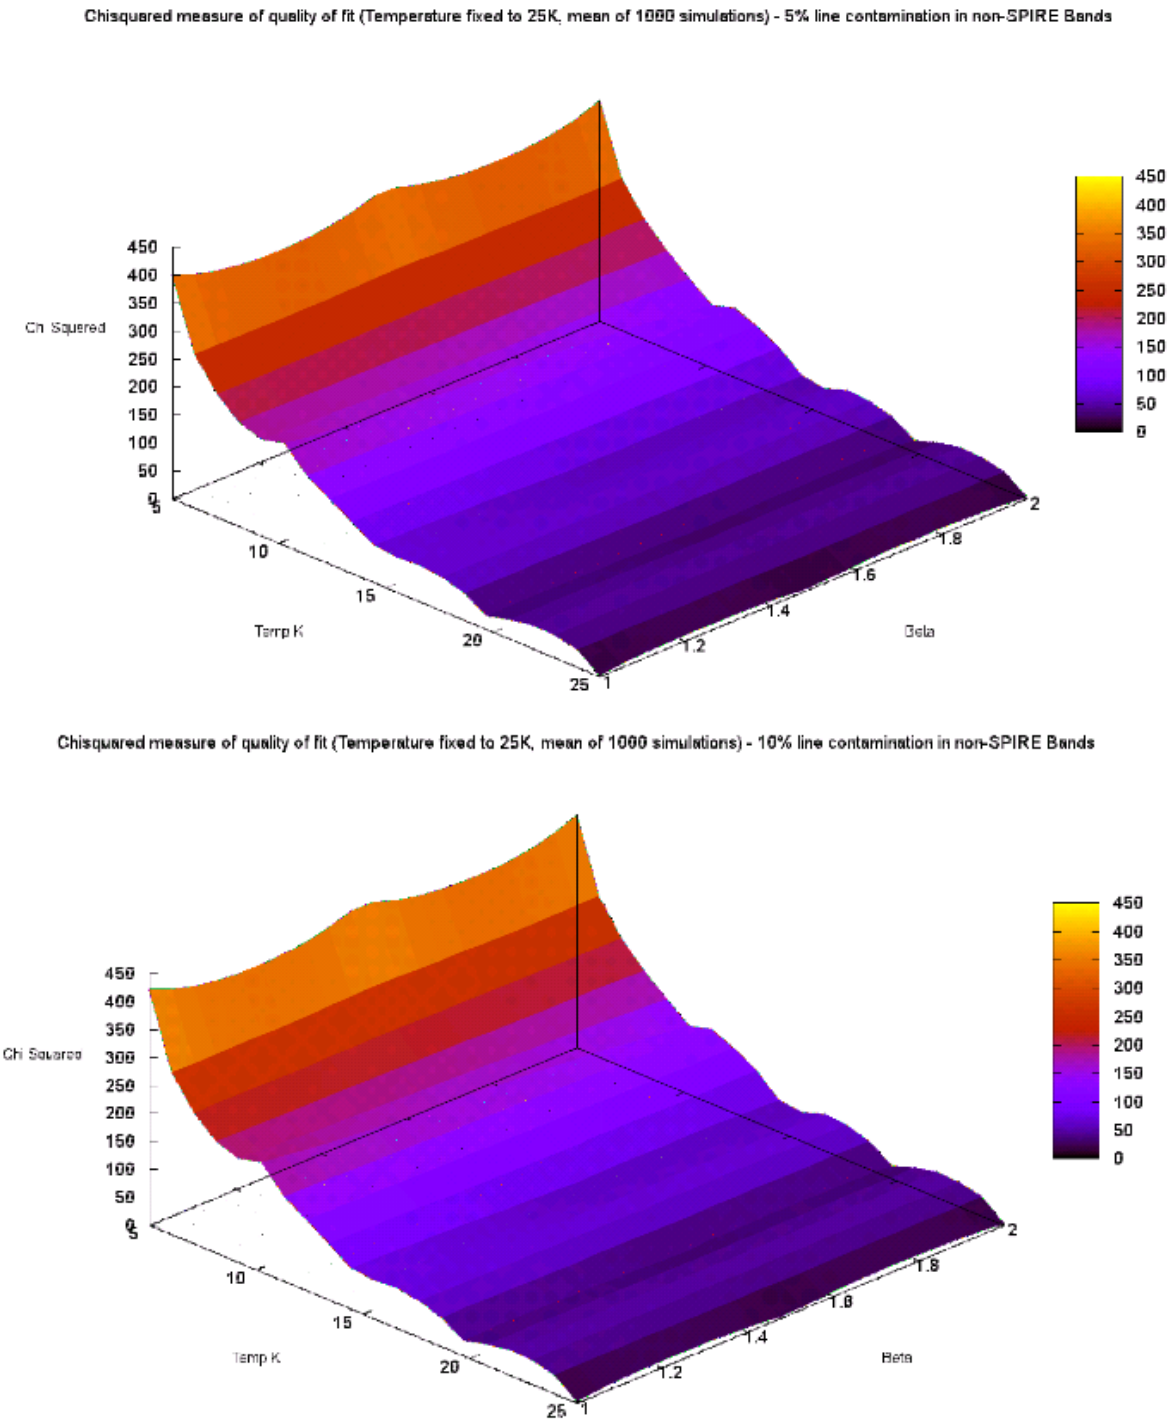

Figure 5.69: Line contamination with $T_{rec} = 25K$ fixed across non-SPIRE bands χ, 5% and 10% line contamination percentage.

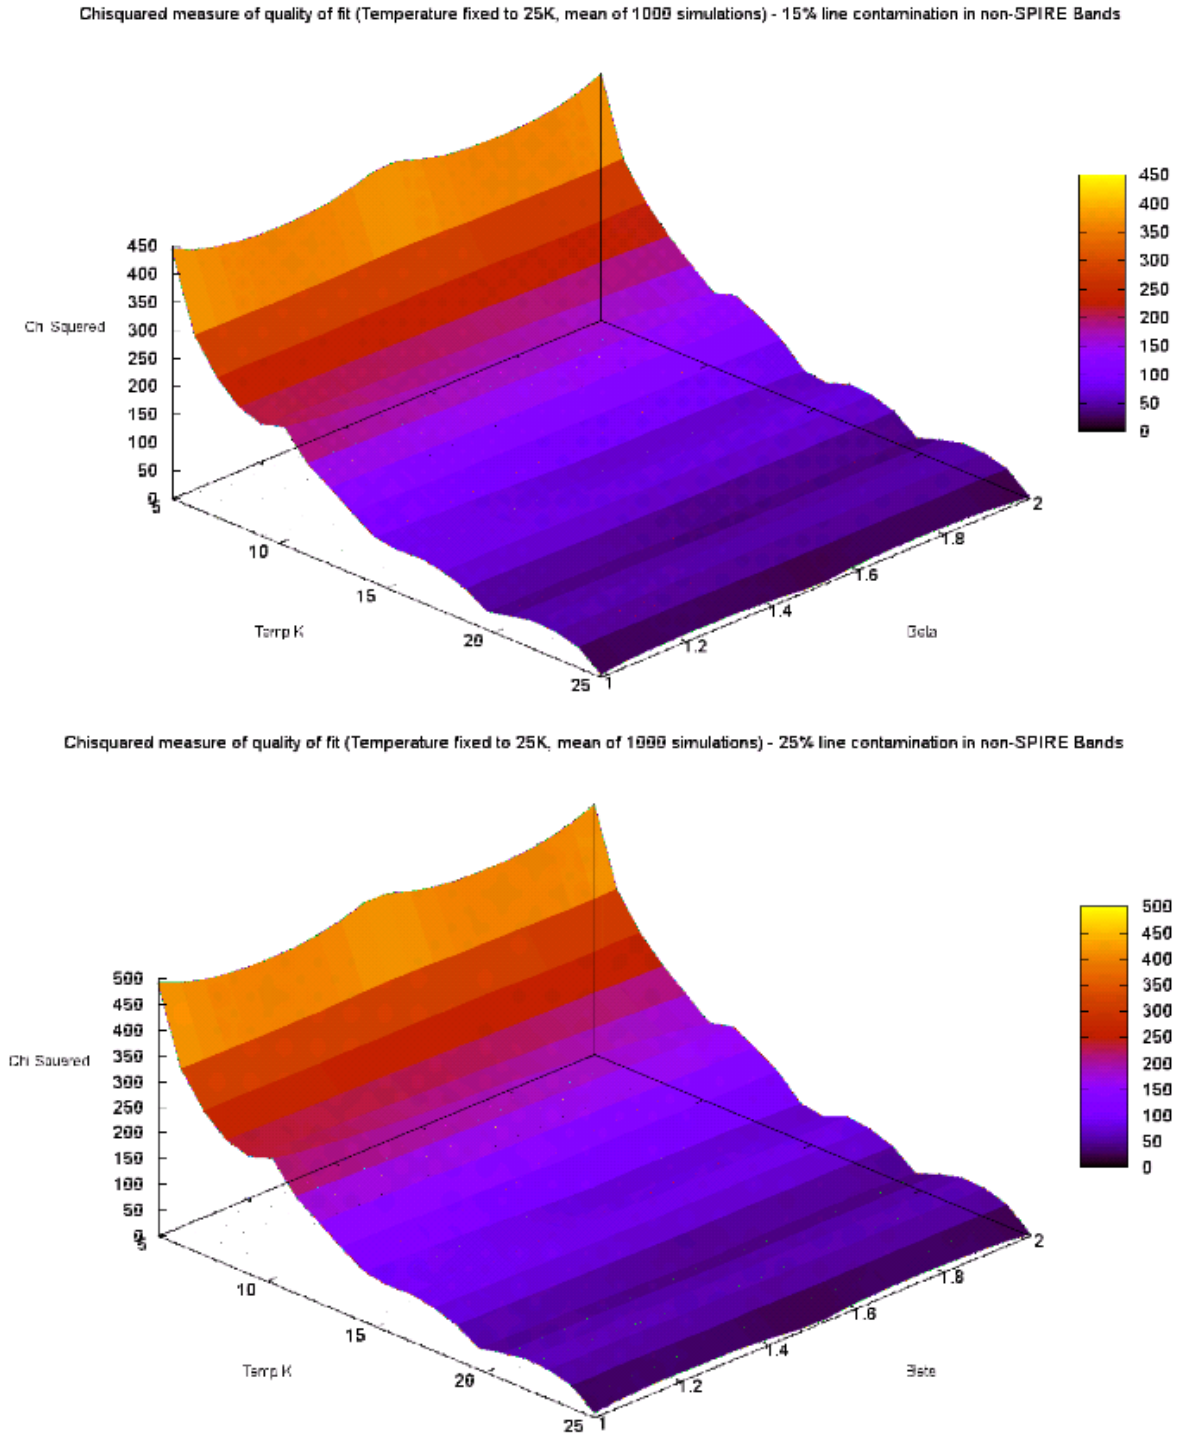

Figure 5.70: Line contamination with $T_{rec} = 25K$ fixed across non-SPIRE bands χ, 15% and 25% line contamination percentage.

5.6.4 Greybody fitting and expanding the simulations

We can examine the effects of global SPIRE line contamination in more detail to examine how the turnover feature of β_{rec} in its central value behaves. This involves expanding the line contamination simulations as follows:

Increase fit resolution: We change the gridding steps of β_{model} and T_{model}, to 0.1 and 2.5K respectively.

Freedom of fit for β_{rec}: We remove the constraint on β_{rec} from being limited to between 1 and 2

Higher temperature simulations: We examine a 40K greybody case, substituting the simulated IRAM and SCUBA wavelengths with simulated PACS observations at $70\mu m$ and $160\mu m$

5.6.4.1 Increase fit resolution

The earlier simulations revealed an unusual behaviour of T_{rec} in the central range. However the small amount of steps in β_{model} means it makes sense to expand the β_{model} and T_{model} ranges and examine this turnover behaviour in greater detail. In order to explore these changes we now plot ΔT and $\Delta\beta$, where the quantity ΔX is defined by $\Delta X \equiv X_{model} - X_{rec}$. This form enables us to explore the deviation between the fitted parameter and the initial value of the parameter in the model.

Figure 5.71 is a plot of ΔT against $\Delta\beta$ for the original gridding, i.e where β_{model} was gridded to $1, 1.5, 2$ and T_{model} was gridded to $5K \ldots 25K$. This plot form was chosen as it shows clearly the results of the simulations for both β_{rec} and

T_{rec} for the models. On this figure, we see a natural grouping of the models with line contamination percentage. We also see that ΔT has a naturally wider range than $\Delta\beta$. These groupings naturally fall into a grouping by β_{model} and increase laterally from the $\Delta T = 0$ axis in order of increasing greybody temperature. So as $|\Delta T|$ increases, so does T_{model}. A clear trend from the data is that the $\beta_{model} = 1$ simulations have T_{rec} less than T_{model}; $\beta_{model} = 2$ has T_{rec} greater than T_{model}; and for $\beta_{model} = 1.5$ they are roughly equal. As T_{model} increases, so does the degree of divergence. T_{rec} also becomes higher than T_{model} for the $\beta_{model} = 1.5$ simulations. As line contamination increases, the behaviour of the recovered temperature for the extremes of β_{model} is consistent, but the central behaviour tends to that of the higher value of β_{model}.

Above the $\Delta\beta = 0$ axis (and hence, where $\beta_{model} > \beta_{rec}$) lie the $\beta_{model} = 2.0$ simulations, where β_{rec} can only be fitted low due to the fit limits. Here we see that as line contamination percentage increases, $\Delta\beta$ decreases. For the 15% and 25% line contamination models, we see that the behavior for the $\beta_{model} = 1.0$ and 1.5 simulations show no turnover such as was seen earlier. For the 5% and 10% line contamination models, we see around $\Delta\beta = -0.1$ the turnover behaviour where the simulations cluster either side of the $\Delta T = 0$ axis, showing the changeover in behavior from T_{rec} being driven high to being driven low. This shows graphically the behaviour we noted earlier. Figure 5.72 shows the error bars on this plot, as a seperate figure for clarity. Even with the different scale of the two axes, it can be seen that on the whole, the error in temperature T_{rec} tends to be much larger than the error in β_{rec}.

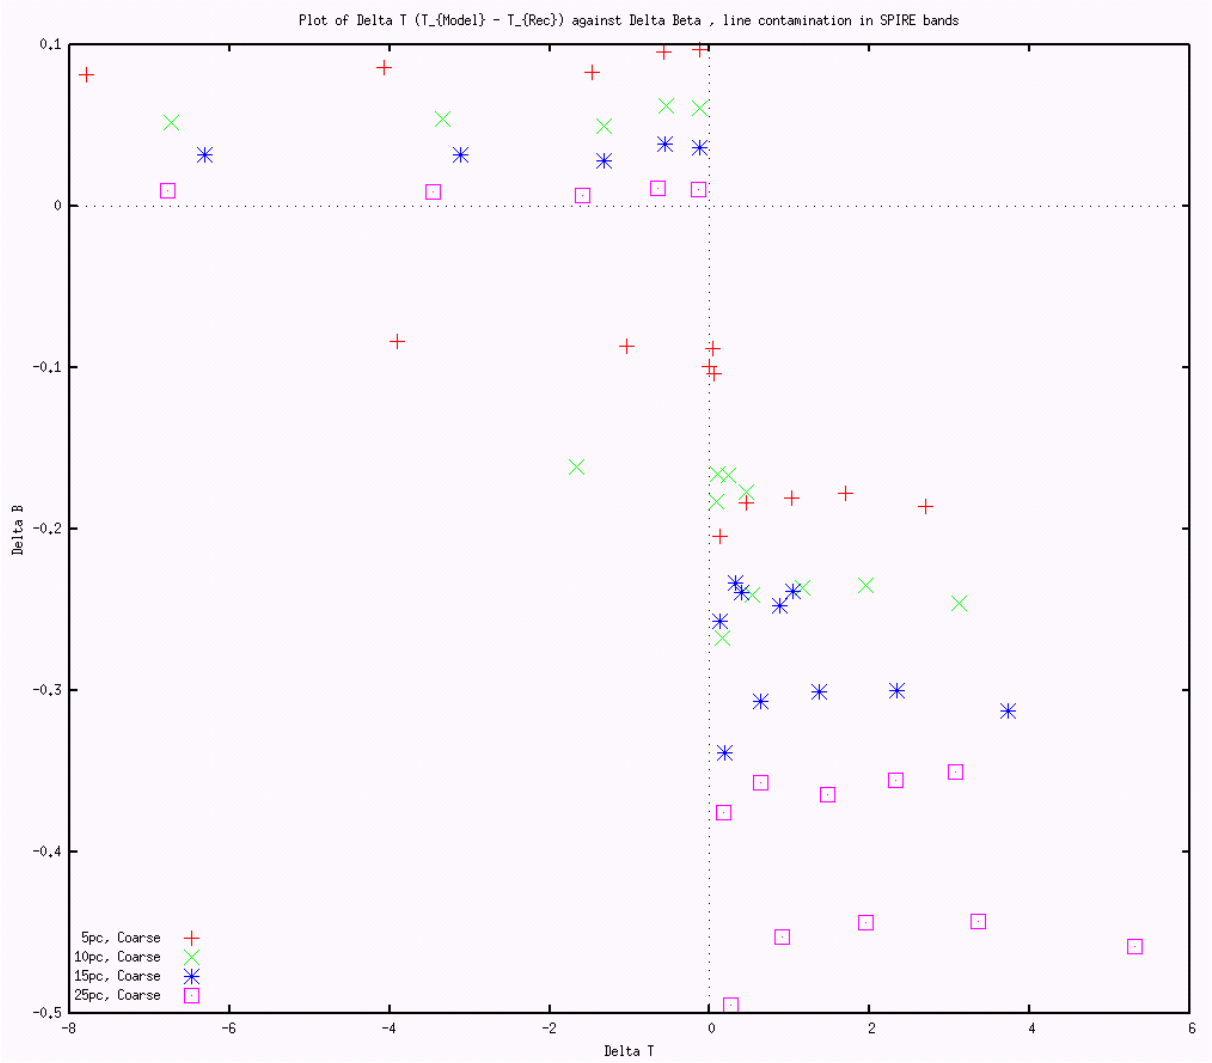

Figure 5.71: A diagnostic plot of ΔT against $\Delta\beta$ for the original resolution simulations.

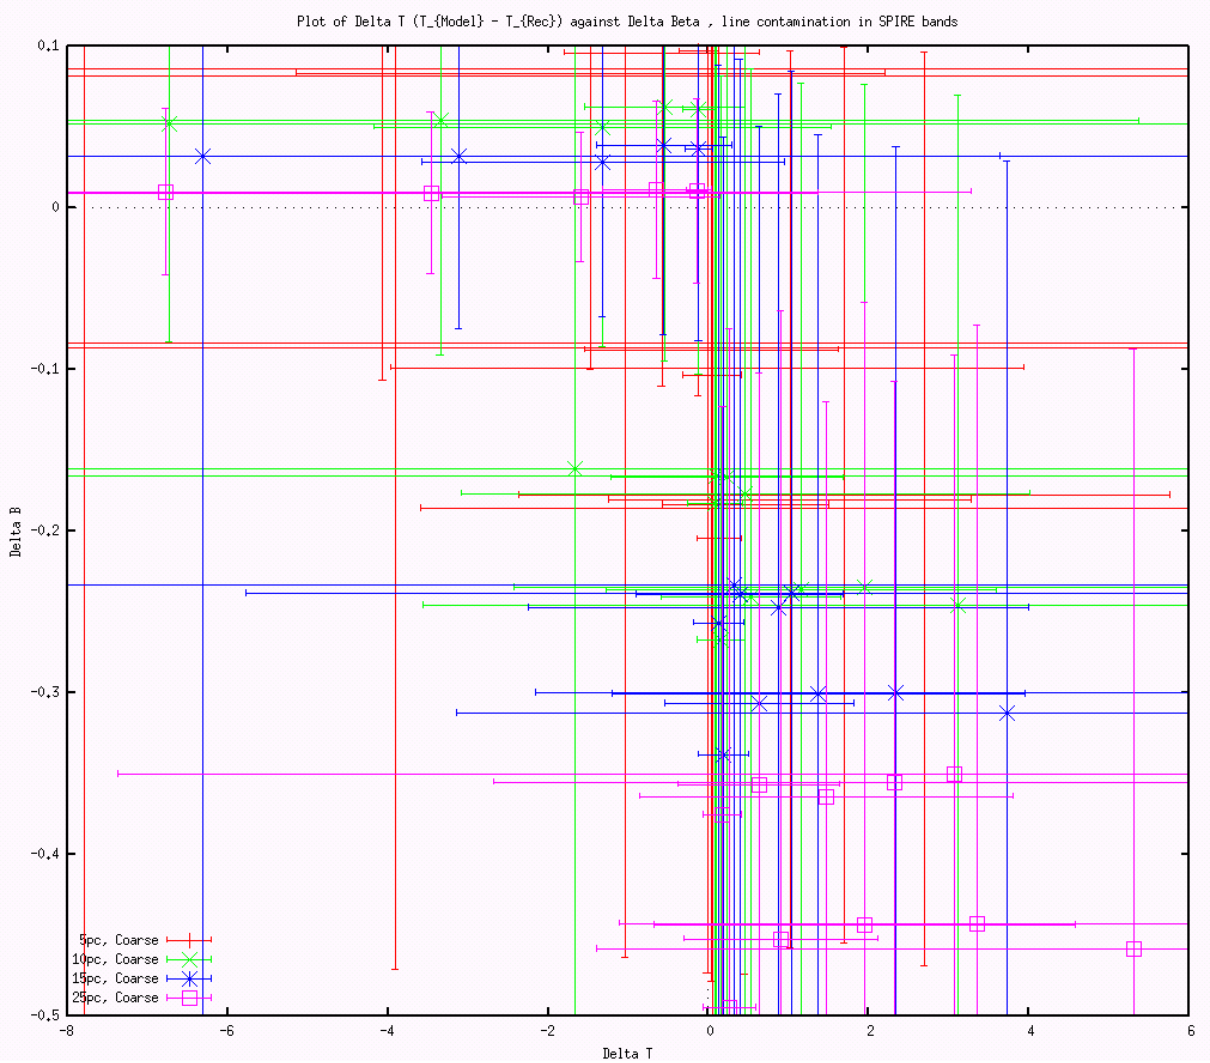

Figure 5.72: A diagnostic plot of ΔT against $\Delta\beta$ for the original resolution simulations with errorbars.

Figure 5.73 is a plot of ΔT against $\Delta\beta$ for the increased resolution simulations. The increased number of simulations reduce the legibility of this plot compared to Figure 5.71, but do show the evolution of the turnover behaviour. Once more ΔT has a naturally wider range than $\Delta\beta$. There is a degree of symmetry in the shape of the diagram, and we see that there is pronounced localisation of simulations on this diagram. For example, the 5% simulations group naturally at the higher end of

$\Delta\beta$, and the 25% simulations group naturally at the lower end. The turnover point where the sign of ΔT changes is seen where $\Delta\beta \approx -0.1 \ldots -0.2$. This expanded diagram shows that this turnover behaviour is seen throughout the simulations at all β_{model} values. The groupings by simulation percentage reveal that lower line contamination percentages tend to give $\Delta\beta = \pm 0.1$ and drive T_{rec} higher than T_{model}, where $\Delta T \leq 0$. However, as β_{model} approaches the upper fit limit of 2.0 we see the beginnings of the turnover behaviour and T_{rec} is driven lower than T_{model}.

For the higher line contamination percentages, $\Delta T \geq 0$ and $\Delta\beta \leq 0$. This behaviour is quite complex compared to the earlier simulations. For the 25% simulations this temperature turnover happens at $\beta_{model} \approx 1.8$ for $T_{rec} = 25K$, and $\beta_{model} \approx 1.9$ for $T_{rec} = 20K, 15K, 10K$. At 15% this happens at $\beta_{model} \approx 1.6$ for $T_{rec} = 25K$ and $\beta_{model} \approx 1.7$ for $T_{rec} = 20K, 15K$,.

For the lower line contaminations, at 10% the turnover is at $\beta_{model} \approx 1.3$ for $T_{rec} = 25K$, $\beta_{model} \approx 1.5$ for $T_{rec} = 20K$ and $\beta_{model} \approx 1.7$ for $T_{rec} = 15K$ and $10K$. At 5% this turnover is seen at $\beta_{model} \approx 1.3$ for $T_{rec} = 25K$ and $20K$, $\beta_{model} \approx 1.4$ for $T_{rec} = 15K$ and $\beta_{model} \approx 1.7$ for $T_{rec} = 10K$. For β_{rec}, in the simulations it is the $\beta_{rec} = 2.0$ values that sit above the $\Delta\beta$ axis. For lower line contamination percentages some of the $\beta_{rec} = 1.9$ simulations are also grouped here.

Figure 5.74 shows the error bars on this figure, again as a seperate figure for clarity. Once more we note that the error bars in T_{rec} are higher than those on β_{rec}.

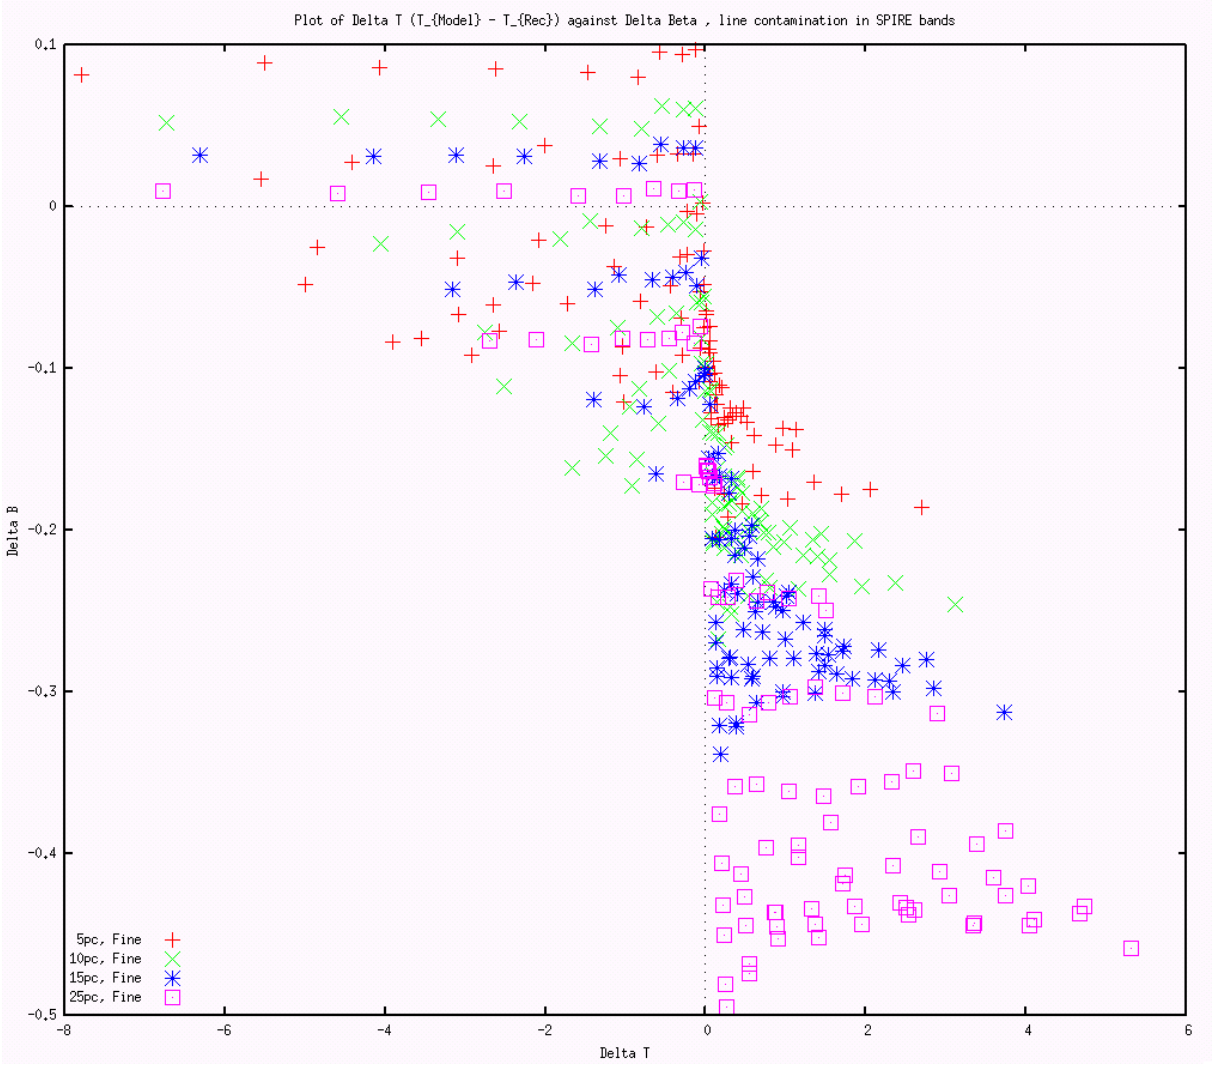

Figure 5.73: A diagnostic plot of ΔT against $\Delta\beta$ for the increased resolution simulations.

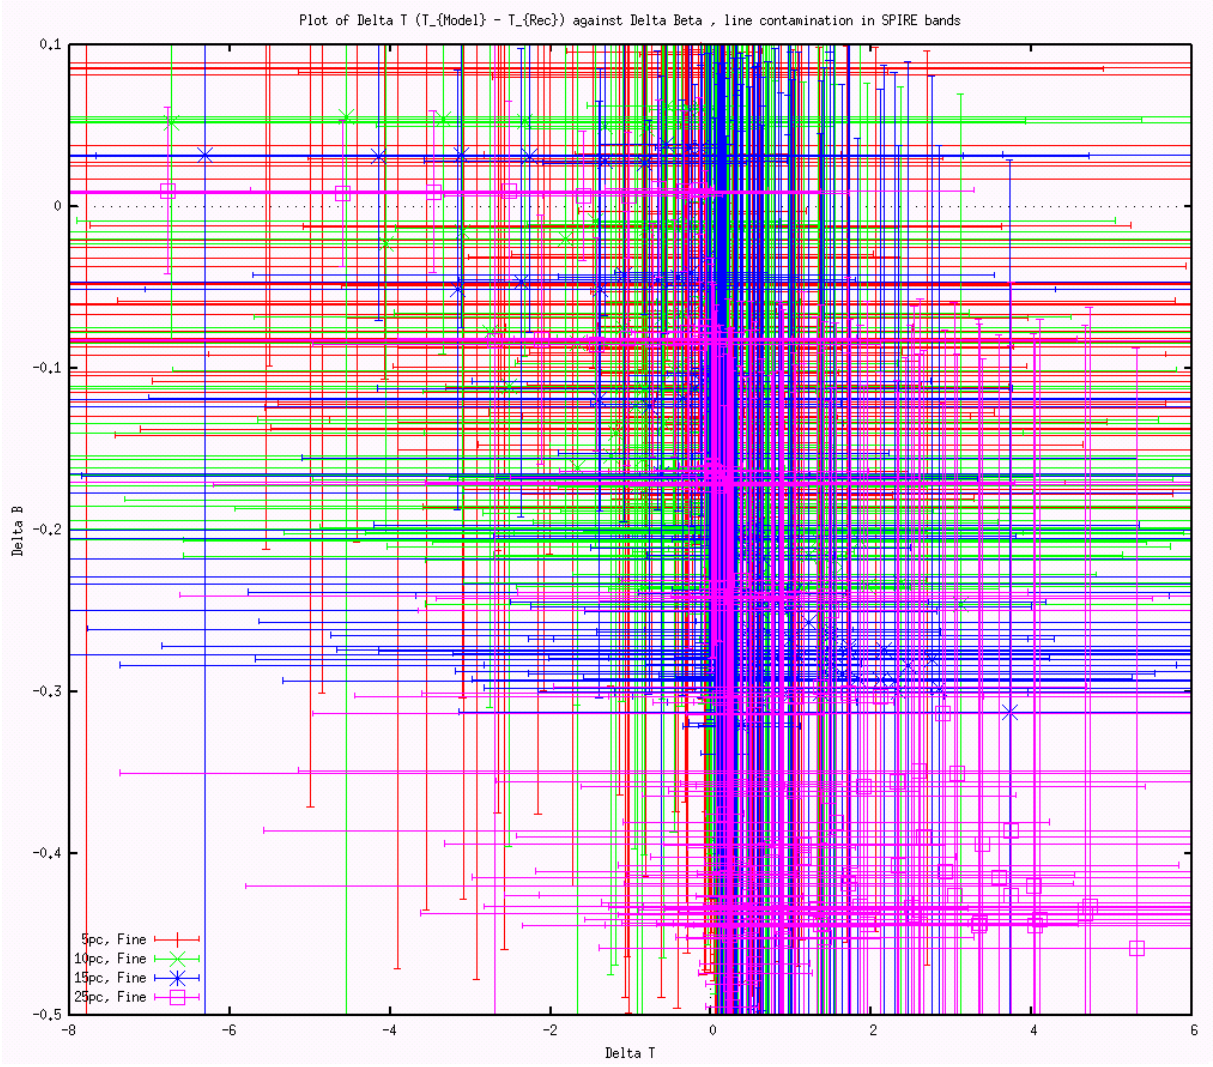

Figure 5.74: A diagnostic plot of ΔT against $\Delta\beta$ for the increased resolution simulations with errorbars.

5.6.4.2 Freedom of fit for β_{rec}

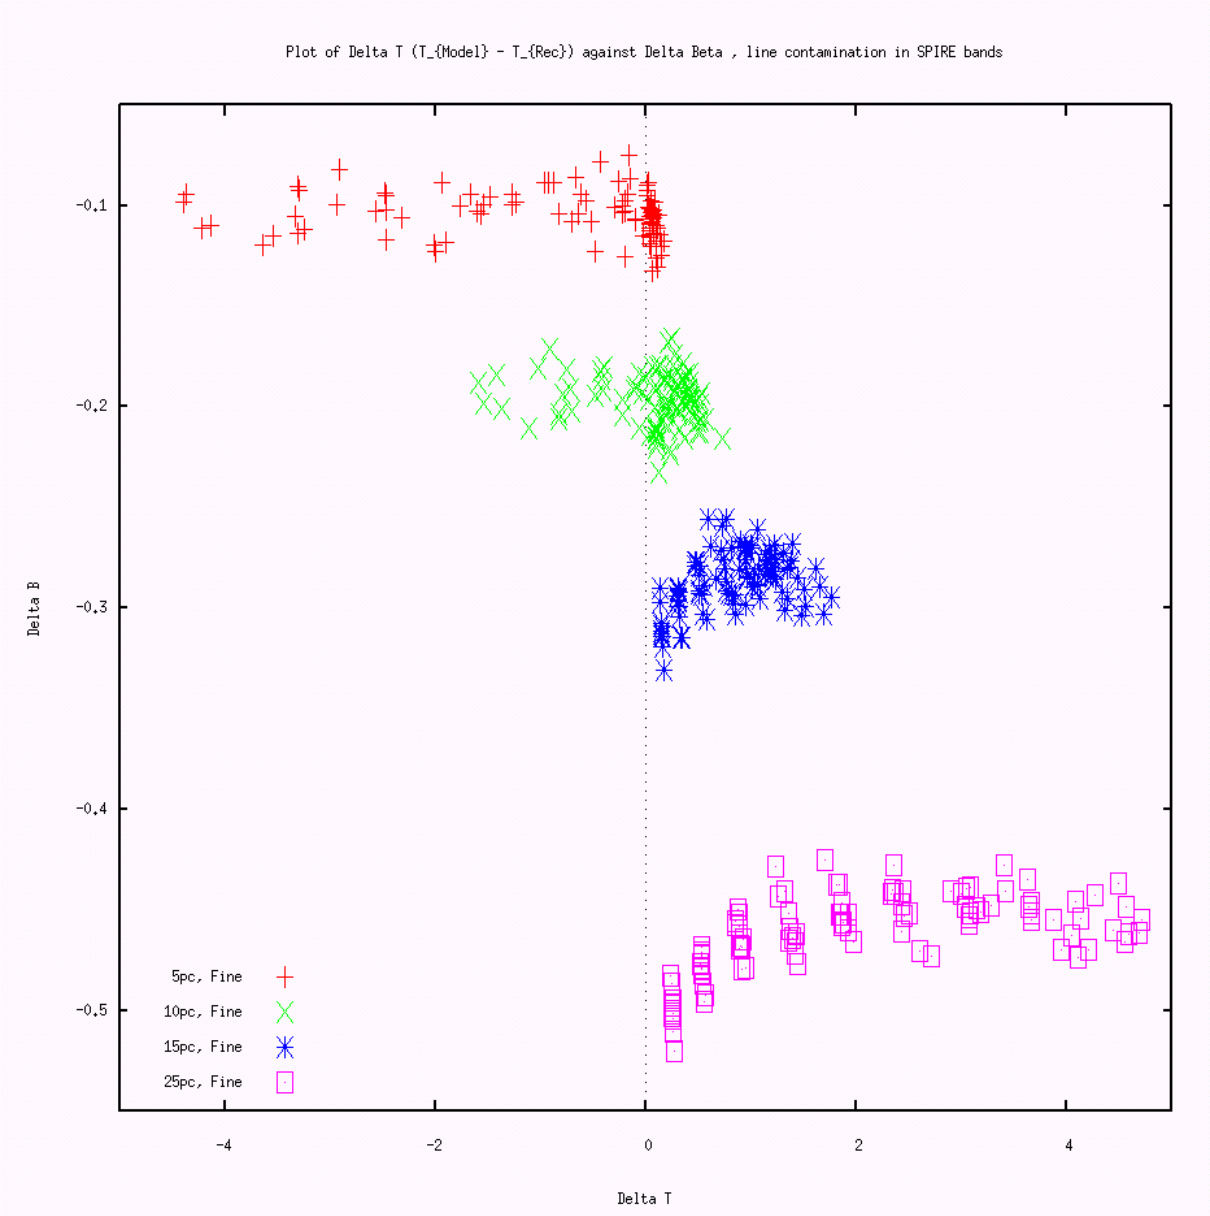

Figure 5.75: A diagnostic plot of ΔT against $\Delta\beta$ for the increased resolution simulations, removing the fit constraint on β_{rec}.

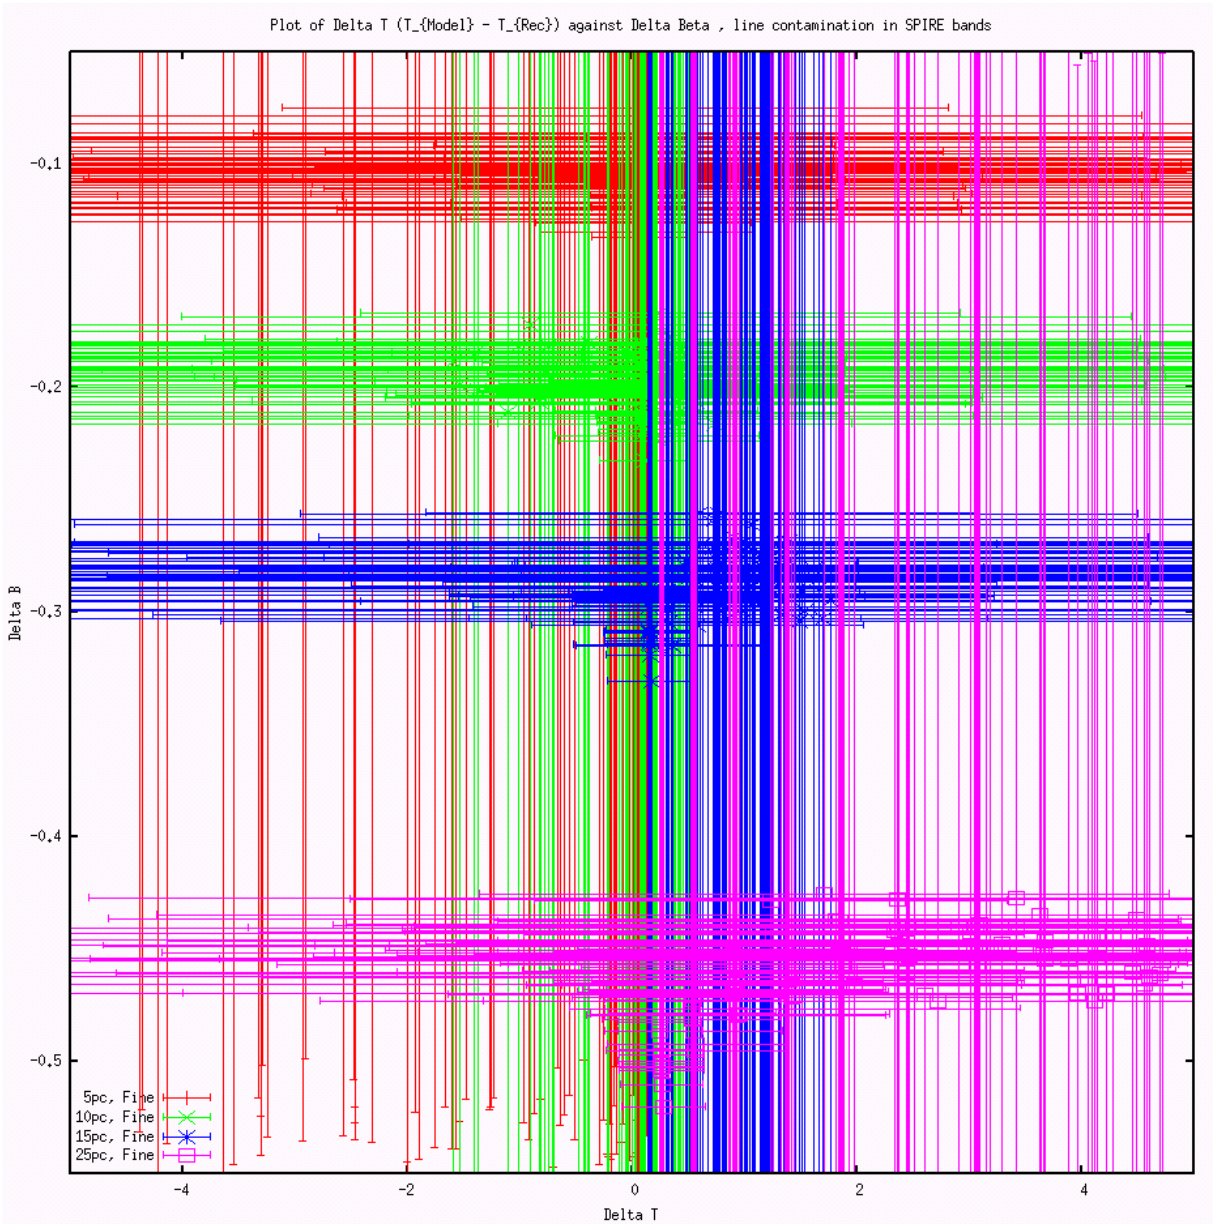

Figure 5.76: A diagnostic plot of ΔT against $\Delta\beta$ for the increased resolution simulations, removing the fit constraint on β_{rec} with errorbars.

Figure 5.75 is a plot of ΔT against $\Delta\beta$ for the increased resolution simulations, where the constraint of fit on β_{rec} has been removed. As a result these simulations are probably the most interesting, as there is little constraint on the fit. Here we can probe quite deeply the behaviour of line contamination on fitting. The

constraint on temperature of 2K upwards was used in the increased gridding fits, as there is a physical limit on T_{rec} of 0K. In these simulations, we see the behaviour of the fitting where the temperature T_{rec} and β_{rec} are largely unconstrained.

We see the data points form a distinct tendency to group. The 5% line contamination simulations group along the $\Delta\beta = -0.1$ (and hence β_{rec} is around 0.1 greater than β_{model}) with a range of $\Delta T \approx 0$ to $-5K$. In fact, over these simulations it seems that $\Delta\beta \approx -0.02 \times LC$ where LC is the line contamination percentage. For these 5% line contamination simulations, we see for $T = 7.5 - 12.5K$ that $T_{model} \leq T_{rec}$, i.e. $\Delta T \geq 0$. Above this temperature the turnover in behaviour in T occurs, and ΔT increases with increasing T_{model}.

We see from Figure 5.75 that the 10% simulations show similar behaviour, but the turnover temperature here is higher at around 20K. We see this graphically with the increased clumping of the data points around the $\Delta T = 0$ axis. Above this line contamination percentage, $\Delta T \geq 0$, and this behaviour is consistent across all values of T_{model} and β_{model}. For the 5% and 25% simulations the spread of ΔT is at its largest at around 5K. For the other line contamination percentages, where the results are grouped around the $\Delta T = 0$ axis, the spread is only around 2K.

Again, we note that Figure 5.76 urges caution, as error bars are large. However, we can now consider each line contamination percentage in isolation to see if we can identify more trends within the data.

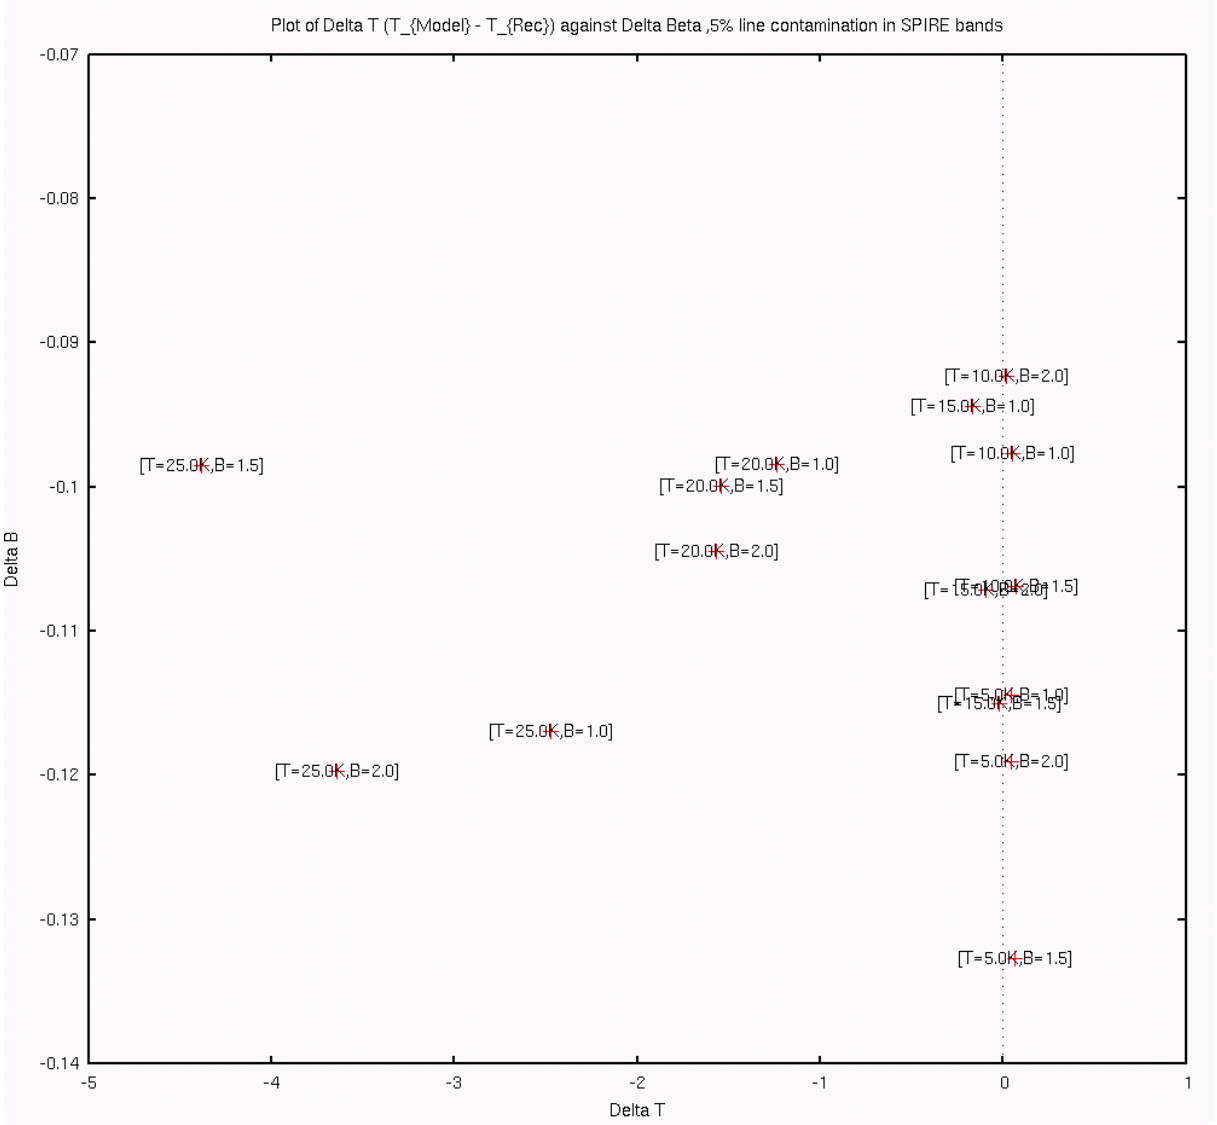

Figure 5.77: A diagnostic plot of ΔT against $\Delta\beta$ for 5% line contamination, coarsely gridded.

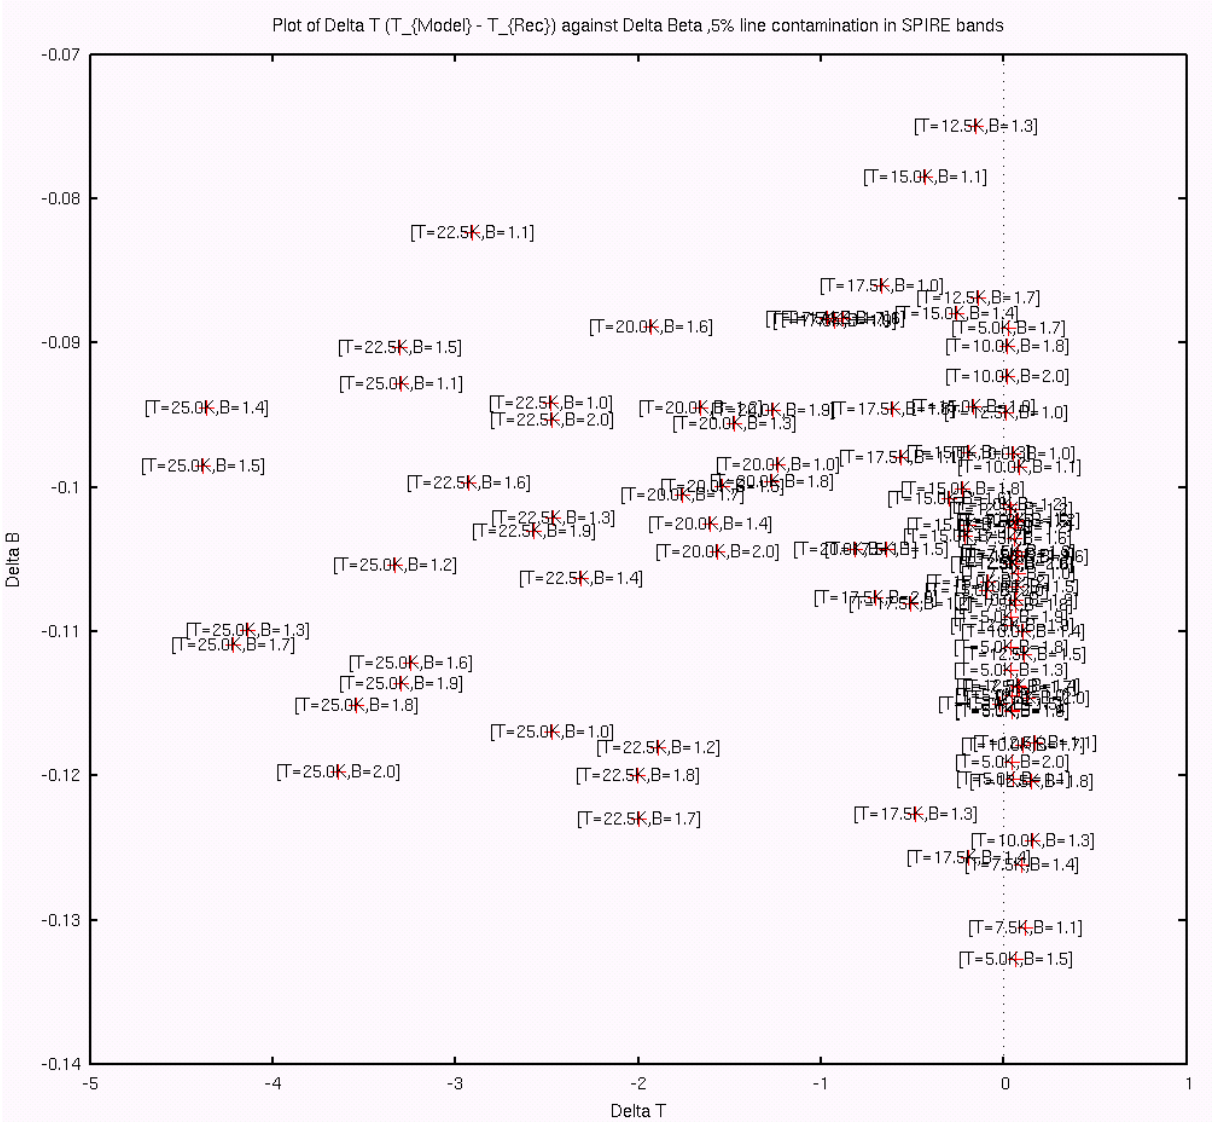

Figure 5.78: A diagnostic plot of ΔT against $\Delta\beta$ for 5% line contamination, finely gridded.

Figure 5.78 shows just the 5% line contamination simulations, with each point labelled with T_{rec} and β_{rec}. The finer gridded simulations provide a lot of data, and the plot is difficult to read. As a result, we return to the coarser gridded simulations to identify trends.

From Figure 5.77 for the 5% line contamination case, we note that ΔT is

increasingly negative for higher T_{model} simulations. For $T_{model} \leq 15K$, we see very little difference in recovered temperature. $\Delta\beta$ is small across the range of simulations, with these lower values of T_{model} showing the greatest dynamic range of β_{rec}. Here we see that line contamination has the greatest effect on T_{rec}, which we see from Figure 5.75 is consistent for each percentage. Figure 5.78 confirms this for the finer gridded simulations.

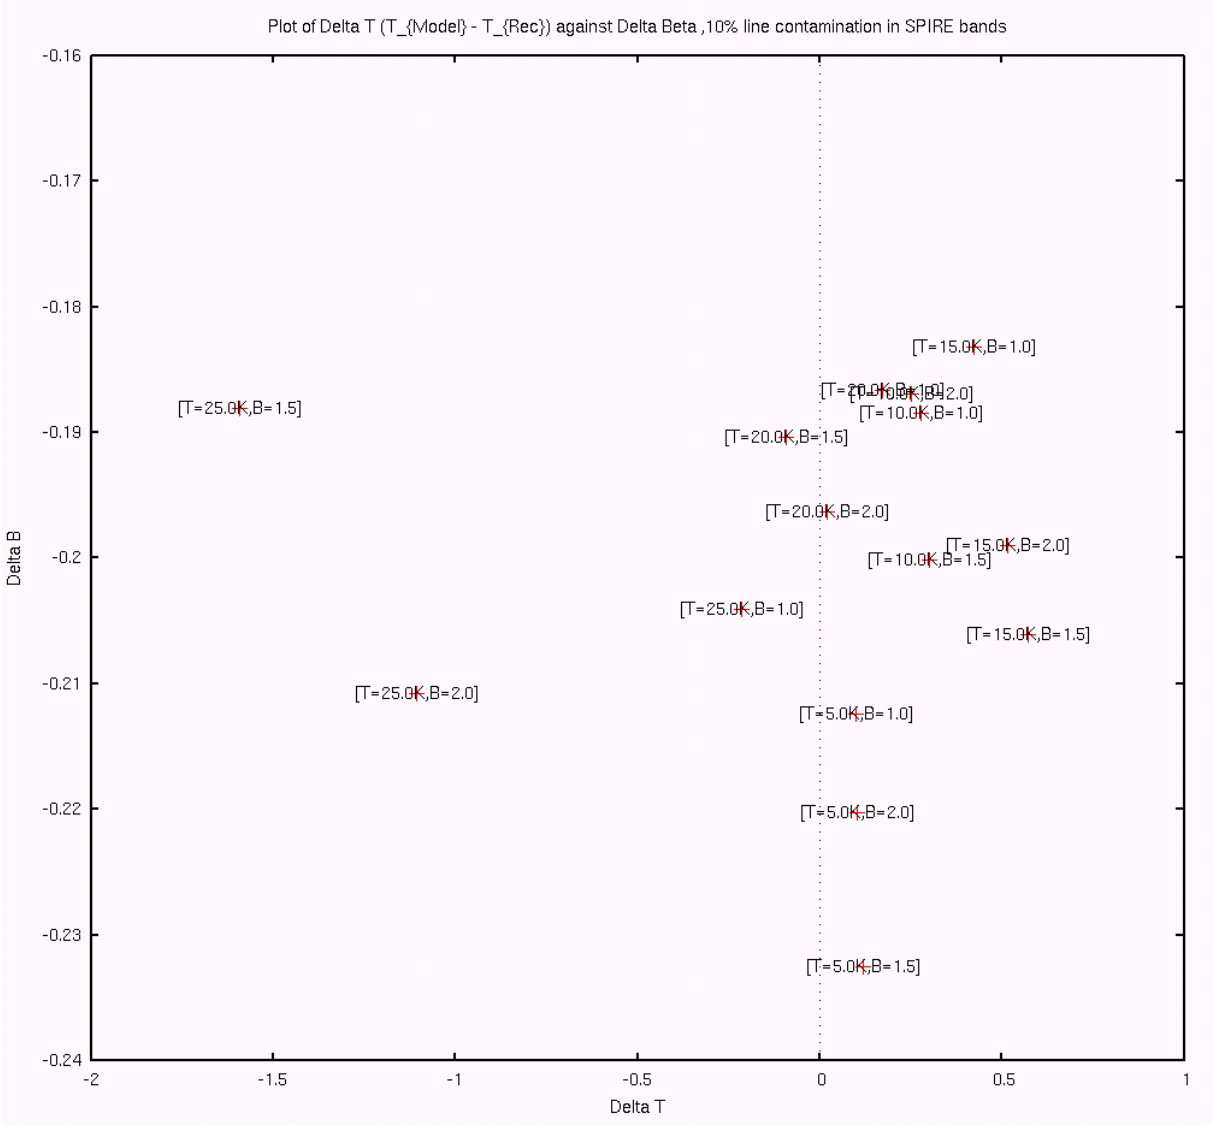

Figure 5.79: A diagnostic plot of ΔT against $\Delta\beta$ for 10% line contamination, coarsely gridded.

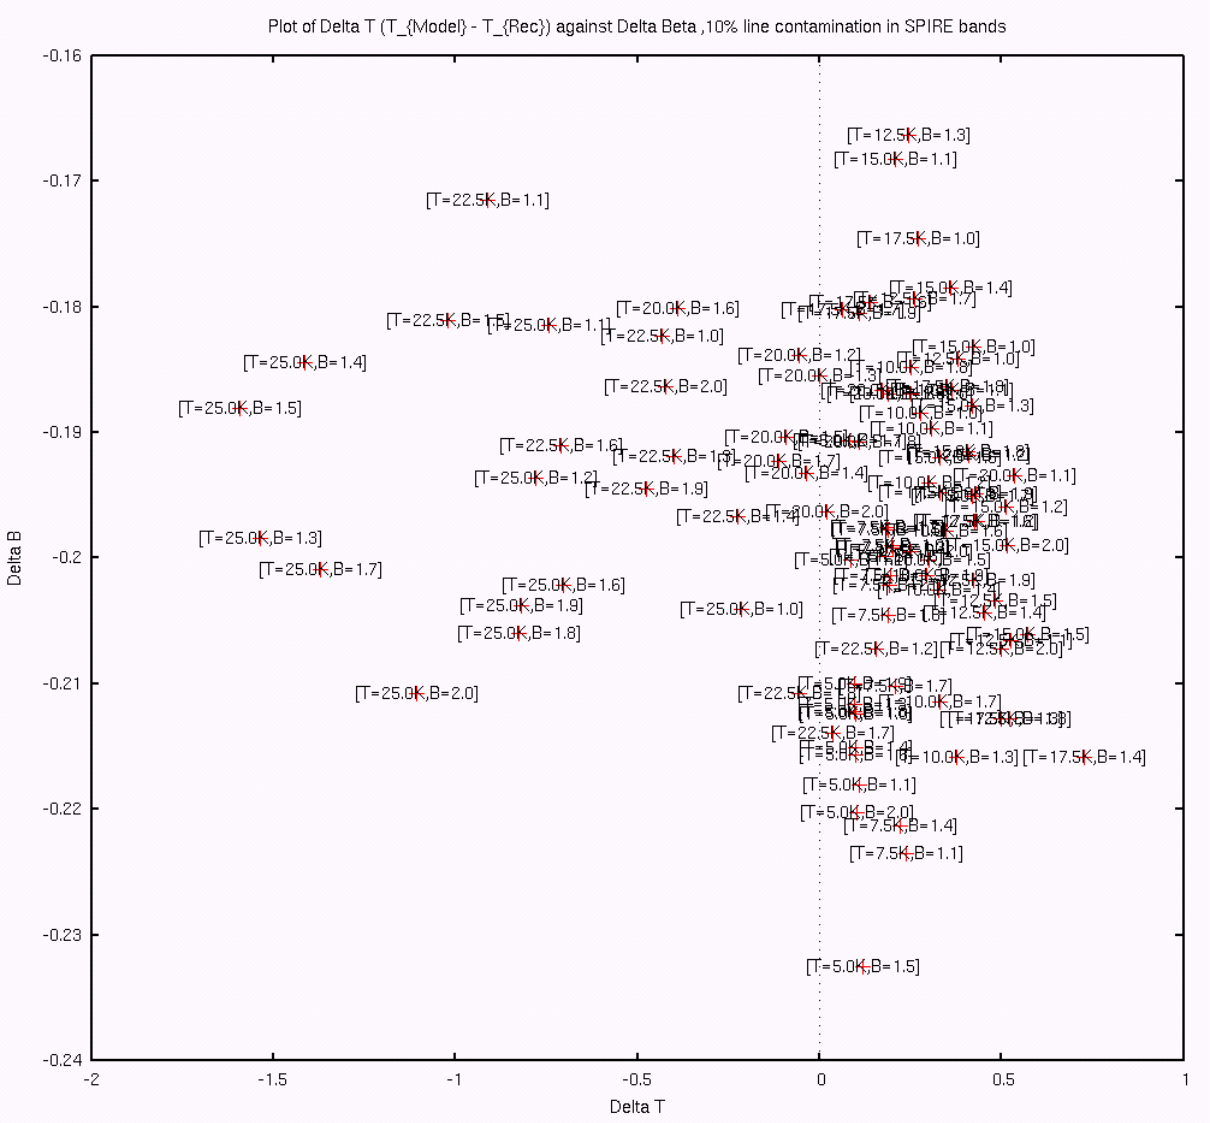

Figure 5.80: A diagnostic plot of ΔT against $\Delta\beta$ for 10% line contamination, finely gridded.

Figure 5.79 shows that the conclusions for the 10% line contamination case are broadly similar to the 5% case. The $T_{model} = 25K$ and $\beta_{model} \geq 1$ simulations show the largest negative value for ΔT, and hence T_{rec} is recovered higher than T_{model}. In all other cases $T_{rec} \approx T_{model}$. Figure 5.80 shows a more indepth view of this case, which shows that the $T_{model} \leq 20K$ simulations recover the most accurate

temperature in this case, and that $\Delta\beta$ has a broadly similar dynamic range. $\Delta\beta$ is approximately twice that of the 5% case, suggesting a roughly linear relationship between $\Delta\beta$ and line contamination percentage.

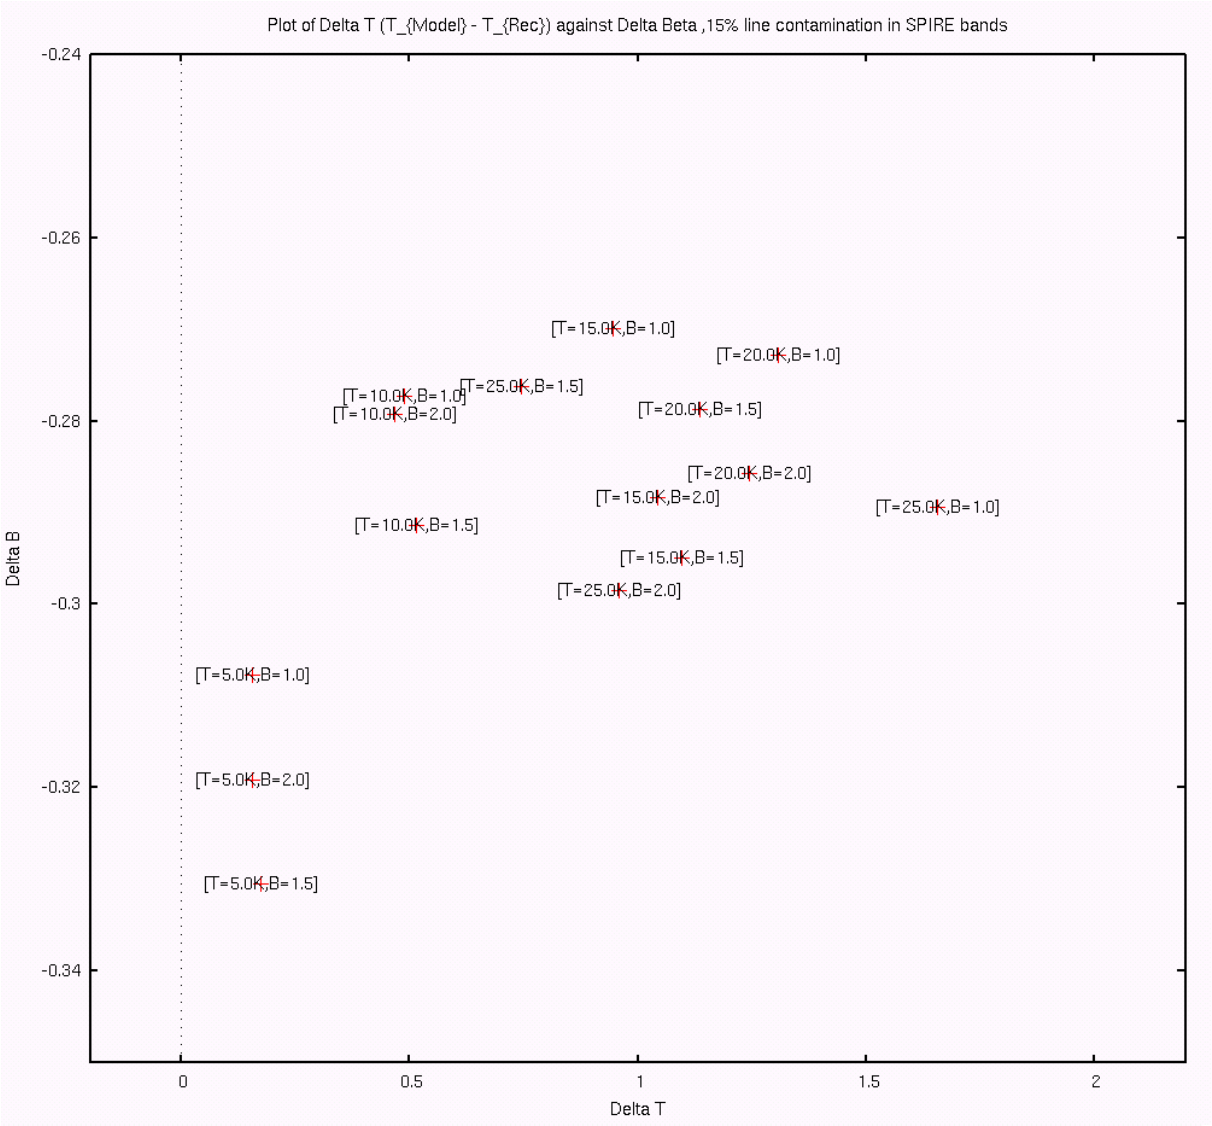

Figure 5.81: A diagnostic plot of ΔT against $\Delta\beta$ for 15% line contamination, coarsely gridded.

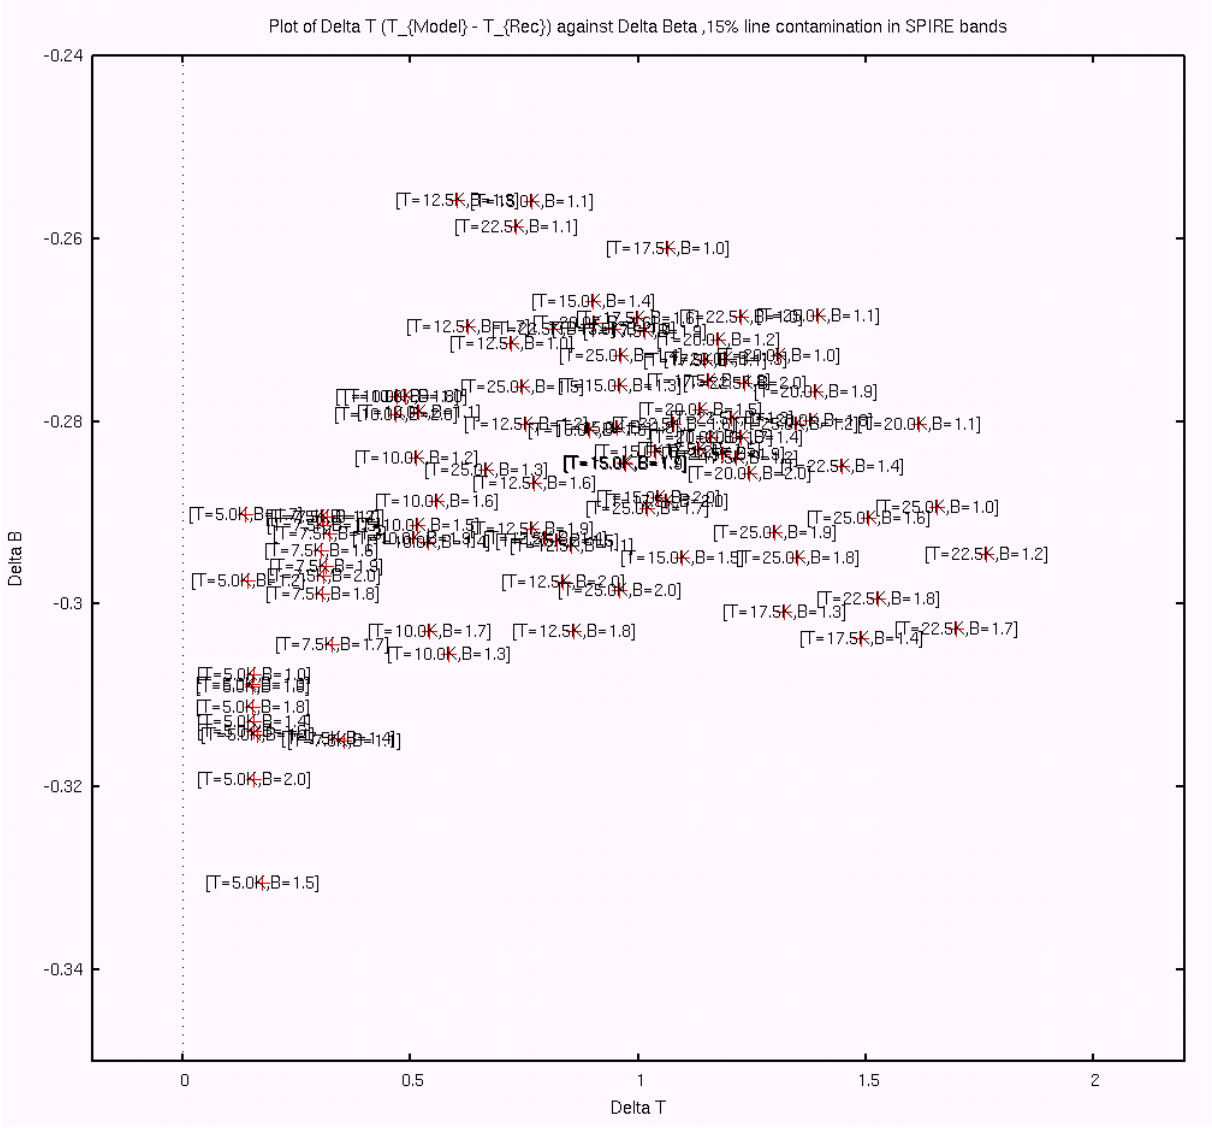

Figure 5.82: A diagnostic plot of ΔT against $\Delta\beta$ for 15% line contamination, finely gridded.

Figure 5.81 again suggests a roughly linear relationship between $\Delta\beta$ and line contamination percentage. Again, we see that the lower T_{model} simulations show a lower value of ΔT and hence show greater accuracy in recovering temperature. Figure 5.82 shows that there is a vertical banding of T_{model} against ΔT. There also seems to be a weak correlation, with higher β_{model} simulations showing increasingly

negative values for $\Delta\beta$. This suggests that β_{rec} is driven increasingly high as β_{model} increases.

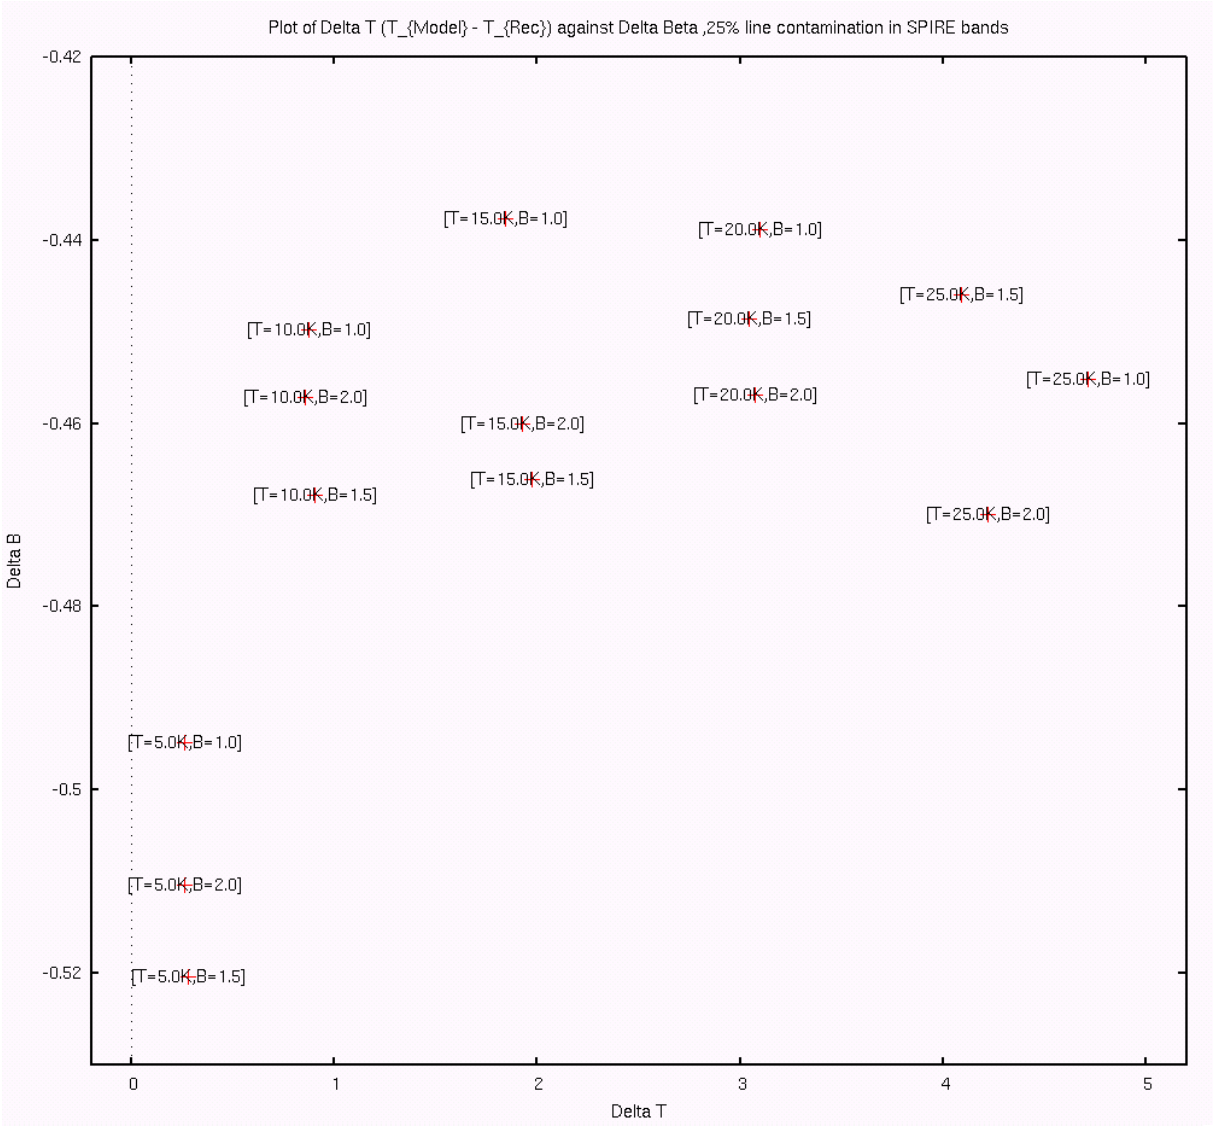

Figure 5.83: A diagnostic plot of ΔT against $\Delta\beta$ for 25% line contamination, coarsely gridded.

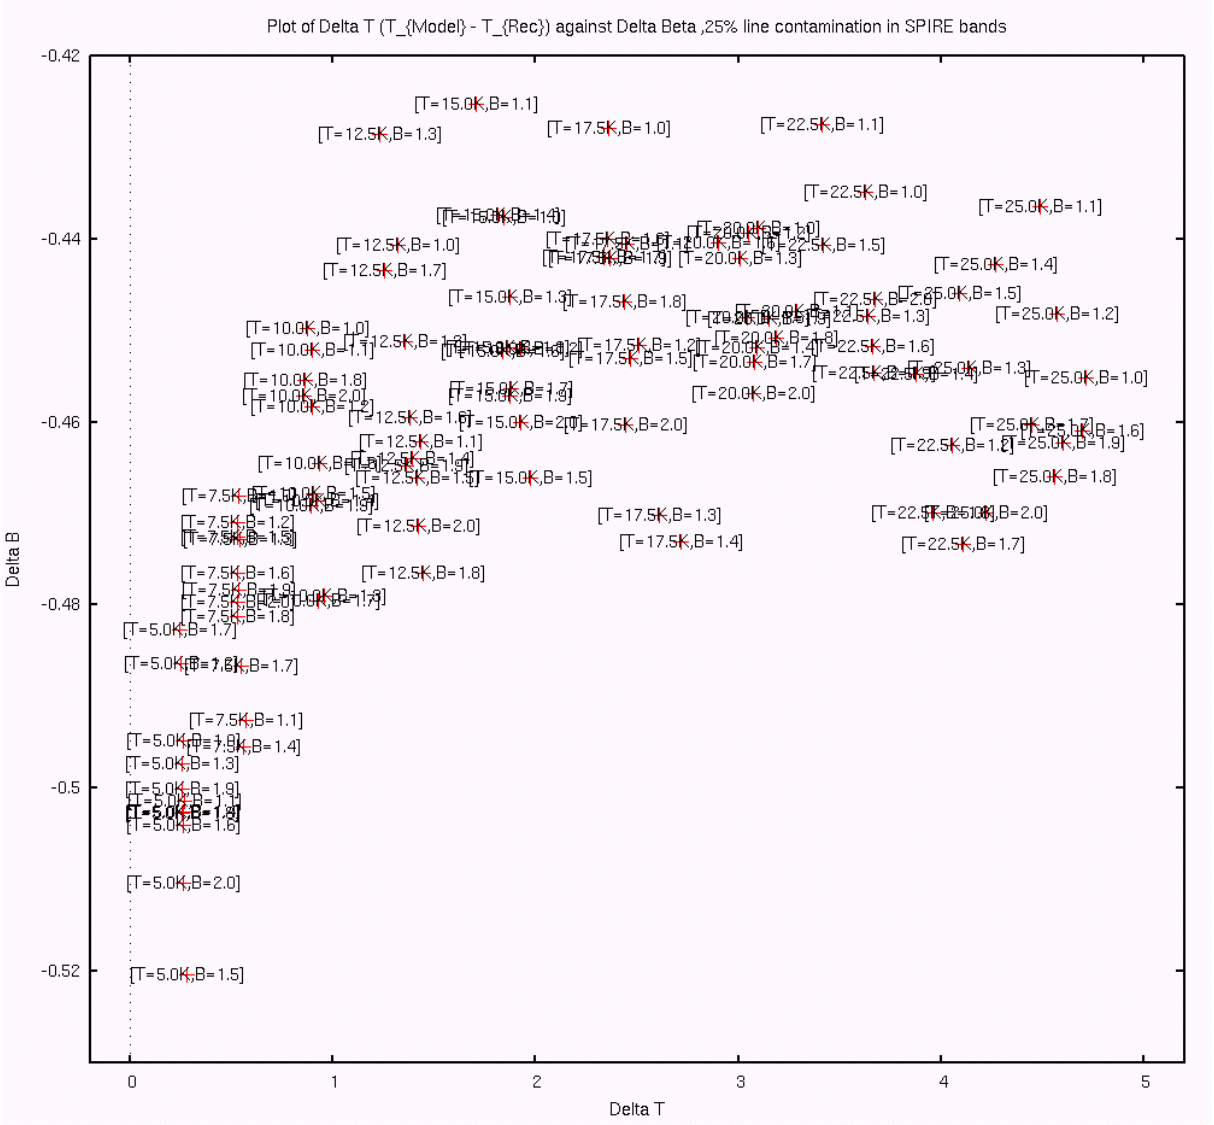

Figure 5.84: A diagnostic plot of ΔT against $\Delta\beta$ for 25% line contamination, finely gridded.

Figure 5.83 again seems to confirm a simple relationship between $\Delta\beta$ and line contamination percentage. Lower T_{model} simulations again show smaller values for ΔT. This seems to suggest that lower T_{model} simulations show greater accuracy in recovering T_{rec}. Figures 5.83 and 5.84 also seem to confirm that higher β_{model} simulations show increasingly negative values for $\Delta\beta$ – and hence β_{rec} is driven

increasingly high as β_{model} increases.

In all line contamination cases, the lower T_{model} simulations also tend to show the greatest negative value of $\Delta\beta$ – and hence for lower T_{model} it seems β_{rec} is increasingly recovered higher than β_{model}

In conclusion we see that, as line contamination increases for simulations where β_{rec} is free, the fit behaviour is dependent on line contamination and T_{model}. At 5% line contamination the temperature is driven high, and a turnover in behaviour is seen, at all temperatures above $T_{model} = 12.5K$. At 10% the turnover in this behaviour is at $T_{model} = 20K$. No turnover is seen at higher line contamination percentages. In all simulations, $\Delta\beta \leq 0$ and is approximately given by $\Delta\beta \approx -0.02 \times LC\%$. Either side of the 10-15% line contamination percentage, ΔT increases in magnitude. However, over all the line contamination percentages the spread of $\Delta\beta$ for each percentage is approximately constant.

Comparing each line contamination percentage individually, we note that lower values of T_{model} are affected least by line contamination. The range of $\Delta\beta$ is roughly constant, but the values of $\Delta\beta$ seem to show a simple linear relationship between line contamination percentage and $\Delta\beta$. There also appears to be a correlation, with higher β_{model} simulations showing increasingly negative values for $\Delta\beta$. This suggests that β_{rec} is driven increasingly high as β_{model} increases within each line contamination percentage band. Lower initial T_{model} simulations show that β_{rec} is increasingly recovered higher than β_{model}.

Figure 5.76 showed the error bars on the fit results, but the sheer number of datapoints meant that this plot only gives an overview. We can examine how the errors in the fitting proccess behave by going back to the coarsely gridded simulations.

Instead of plotting error bars, we shall label each point with $T_{rec}, \beta_{rec}, T_{error}, \beta_{error}$ repspectively. Even though these are the coarsely gridded datapoints, they do yield a valuable insight.

Figures 5.85 to 5.88 show each point labelled as above. Figure 5.85, which shows the error values for 5% line contamination, reveals that errors in T_{rec} scale approximately with T_{model} and errors in β_{rec} are largely constant. For the largest values of T_{model} and β_{model} errors increase to a very large amount indeed. Figure 5.86 shows that the trend in error continues, although there are no very large errors seen. However, we note the error in T can be quite significant. These trends continue in Figures 5.87 and 5.88.

Looking at the error trend across line contamination percentages, we note from these figures that the error in β_{rec} is largely constant across the range considered. However, the error in T_{rec} decreases as line contamination percentage increases. Hence it is the smaller line contamination that yields the greatest uncertainty in T_{rec} but the larger line contamination that yields the greatest value of ΔT. This means that line contamination, T_{model} and β_{model} have a very weak effect on the error in β_{rec}, but that line contamination has a stronger effect on T_{rec}. When errors are considered, we note that β_{rec} is recovered with far greater fidelity than T_{rec}. At higher temperatures and low line contamination levels, there is a lot of uncertainty in T_{rec}. This suggests that subtle levels of line contamination can provide a great deal of uncertainty in recovering T_{rec}.

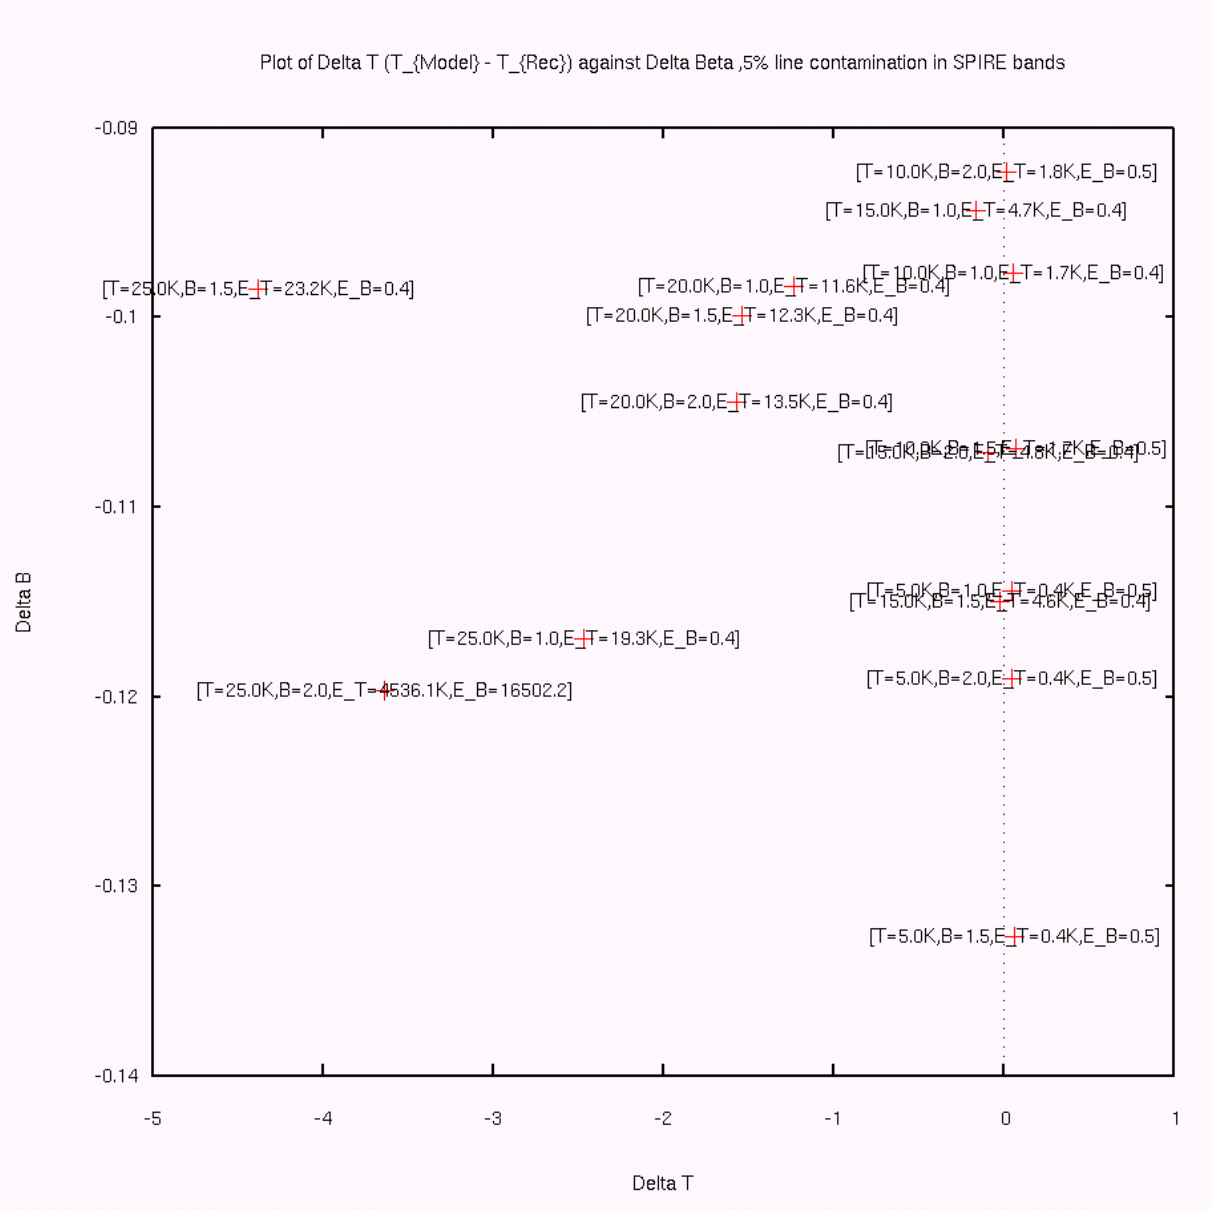

Figure 5.85: $T_{rec}, \beta_{rec}, T_{error}, \beta_{error}$ for 5% line contamination, coarsely gridded.

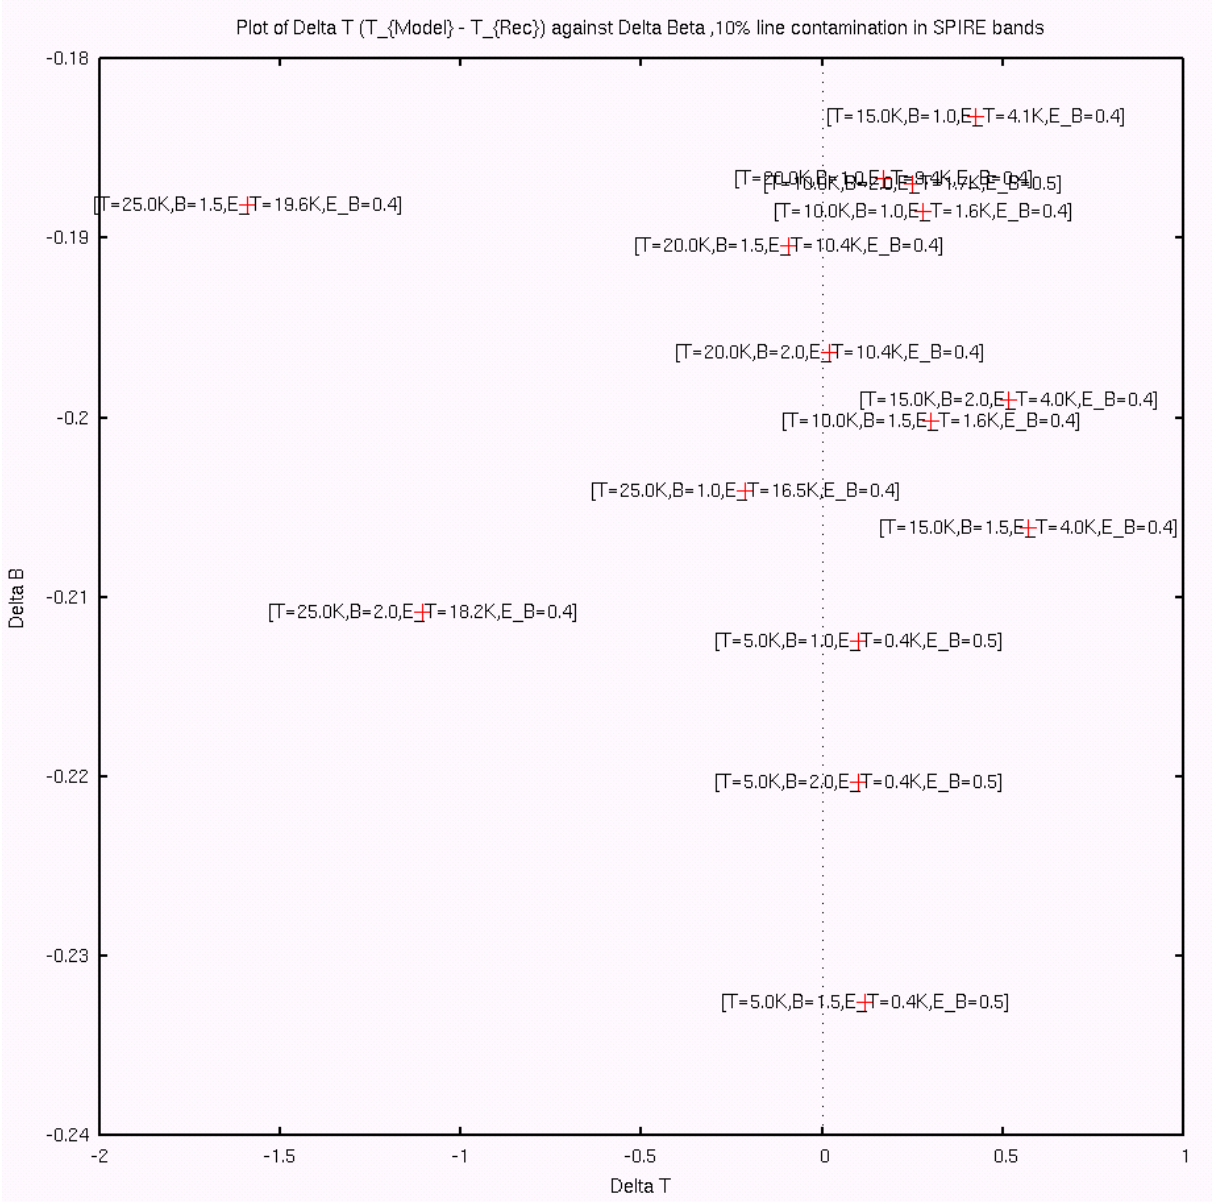

Figure 5.86: $T_{rec}, \beta_{rec}, T_{error}, \beta_{error}$ for 10% line contamination, coarsely gridded.

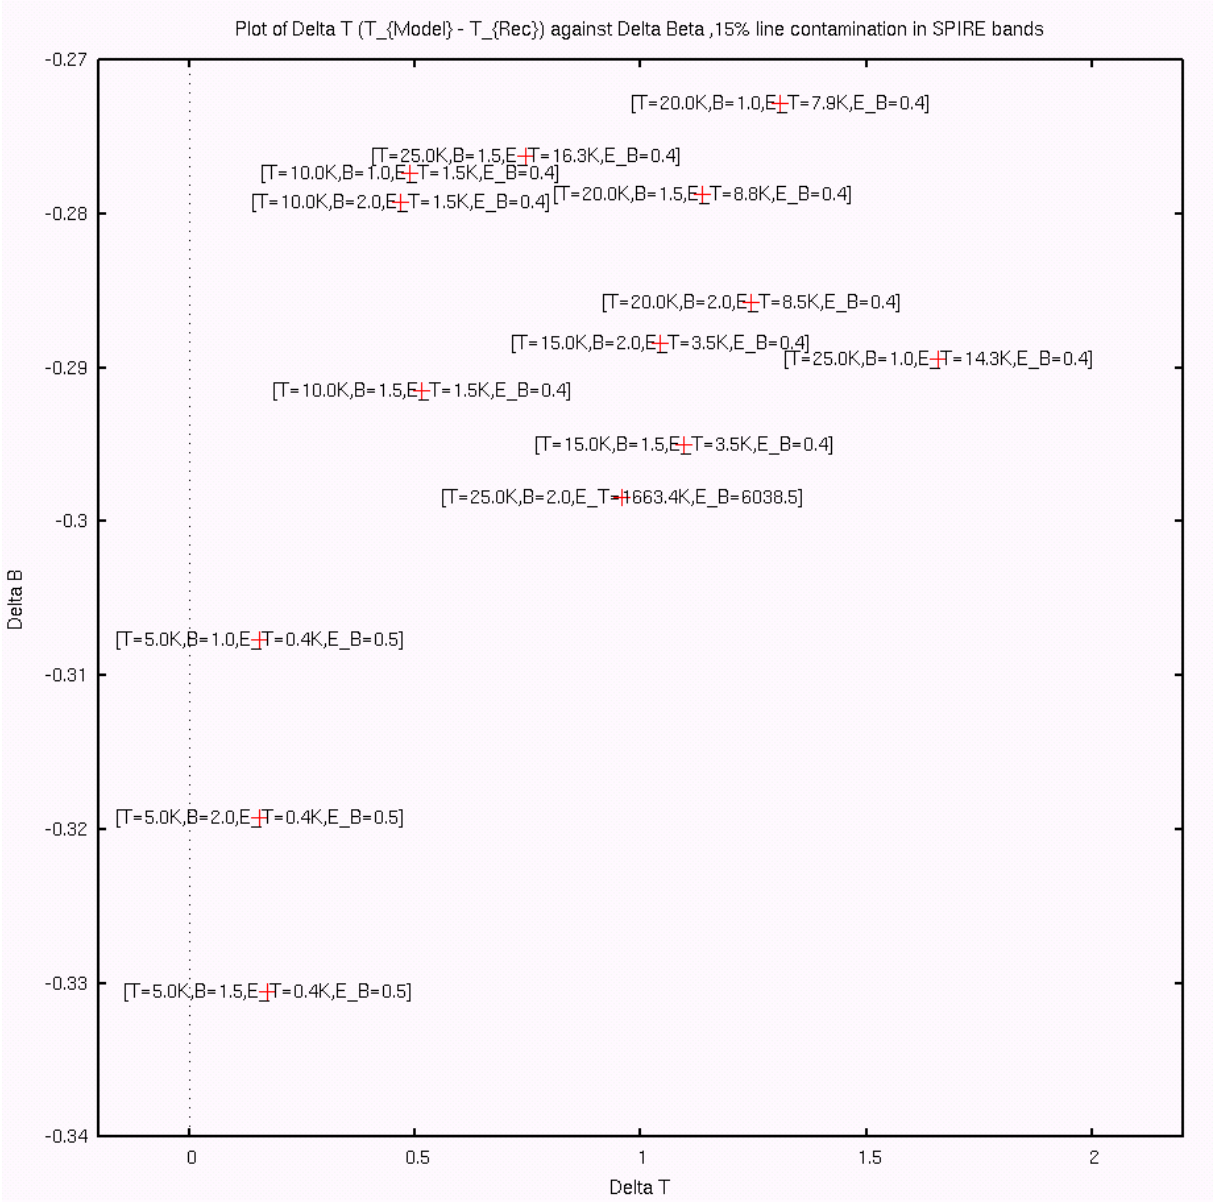

Figure 5.87: $T_{rec}, \beta_{rec}, T_{error}, \beta_{error}$ for 15% line contamination, coarsely gridded.

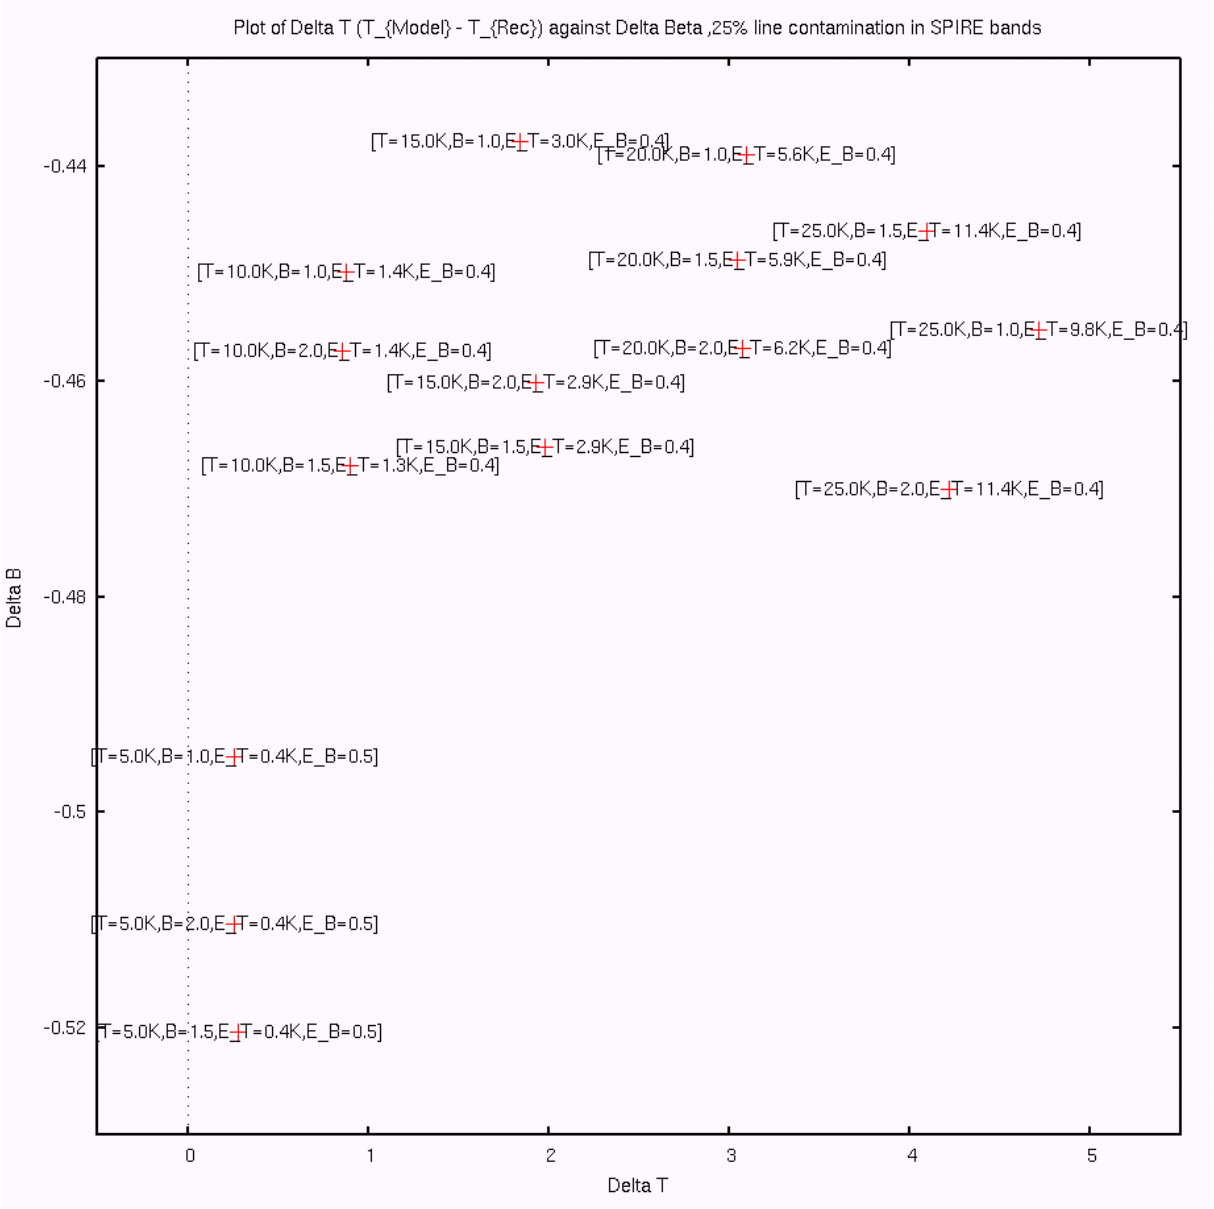

Figure 5.88: $T_{rec}, \beta_{rec}, T_{error}, \beta_{error}$ for 25% line contamination, coarsely gridded.

5.6.4.3 Higher temperature simulations

Figure 5.89 is a plot of ΔT against $\Delta\beta$ for the SPIRE and PACS simulations, where $T_{model} = 40K$ and $\beta_{model} = 1.0, 1.1, 1.2, \ldots 2.0$. Here, the interpretation is made clearer by the fact there is only one T_{model} in the simulations, and β_{model} and

line contamination percentage are the varying parameters.

As the greybody temperature is warmer than the previous simulations, the peak of the greybody emission moves out of the original wavelength range ($200 - 1200\mu m$). To constrain the peak, in these simulations we discard the simulated IRAM and SCUBA observations at 850 and $1200\mu m$ and add in simulated PACS observations at 70 and $160\mu m$. These observations ensure that the peak of the greybody emission (at approximately 73K) is just within our observational window.

In Figure 5.89 we clearly see a grouping of the four line contamination percentages. For each percentage, $\Delta\beta$ groups around 0.04 for 5%, 0.08 for 10%, 0.1 for 15% and 0.16 for 25%. All $\Delta\beta$'s are positive, so β_{rec} is being driven lower than β_{model}. The amount by which it is driven low is dependent on line contamination percentage - as line contamination increases, so does $\Delta\beta$. Also within each percentage grouping, $\Delta\beta$ is largely constant. Within each percentage grouping, we see a range of ΔT which shifts very slightly as contamination increases. The range of ΔT is between $-0.1K$ to $-0.7K$ and T_{rec} is higher than T_{model} for 5%. This becomes $-0.3K$ to $-0.9K$ for 25%. ΔT tends to become increasingly negative as β_{model} increases but this dependence is weak. This implies that ΔT is weakly dependent on β_{model} and has a stronger dependence on LC percentage. For $\beta_{model} = 1$, the range of T_{rec} is 40.38K (5%) to 40.37K (25%) but rises to 40.44K for 10% and 15%. However, the error in T at this value of β_{model} is around 4K and so within the error bars. At $\beta_{model} = 2$, this becomes for increasing contamination, 40.19K, 40.29K, 40.32K, 40.56K, but with an approximate error of 5K.

In these simulations we see that there is a reduced range in both $\Delta\beta$ and ΔT. ΔT varies between $-0.1K$ to $-1K$ and so the recovered temperature is not that dissimilar to the original. However, over the thousand simulations we find

the error bars in T are $\approx \pm 5K$ for the 5% simulations, increasing to $\approx \pm 6K$ for the 25% simulations. Here again, we calculate the final error in the recovered temperature as being the quadrature sum of the mean error for each combination of T_{model}, β_{model}, $LC\%$ over the thousand simulations, and the standard deviation of the thousand recovered parameters. For β_{rec}, the error is ≈ 0.25 overall. As shown in Figure 5.90, the error bars are large compared to the values of ΔT and $\Delta\beta$, and so mask the groupings we see in Figure 5.89. For the higher values of β_{model}, there is also a tendency for errors to rapidly increase, for all but the 25% line contamination case. The errors are very large for $\beta_{model} = 2$.

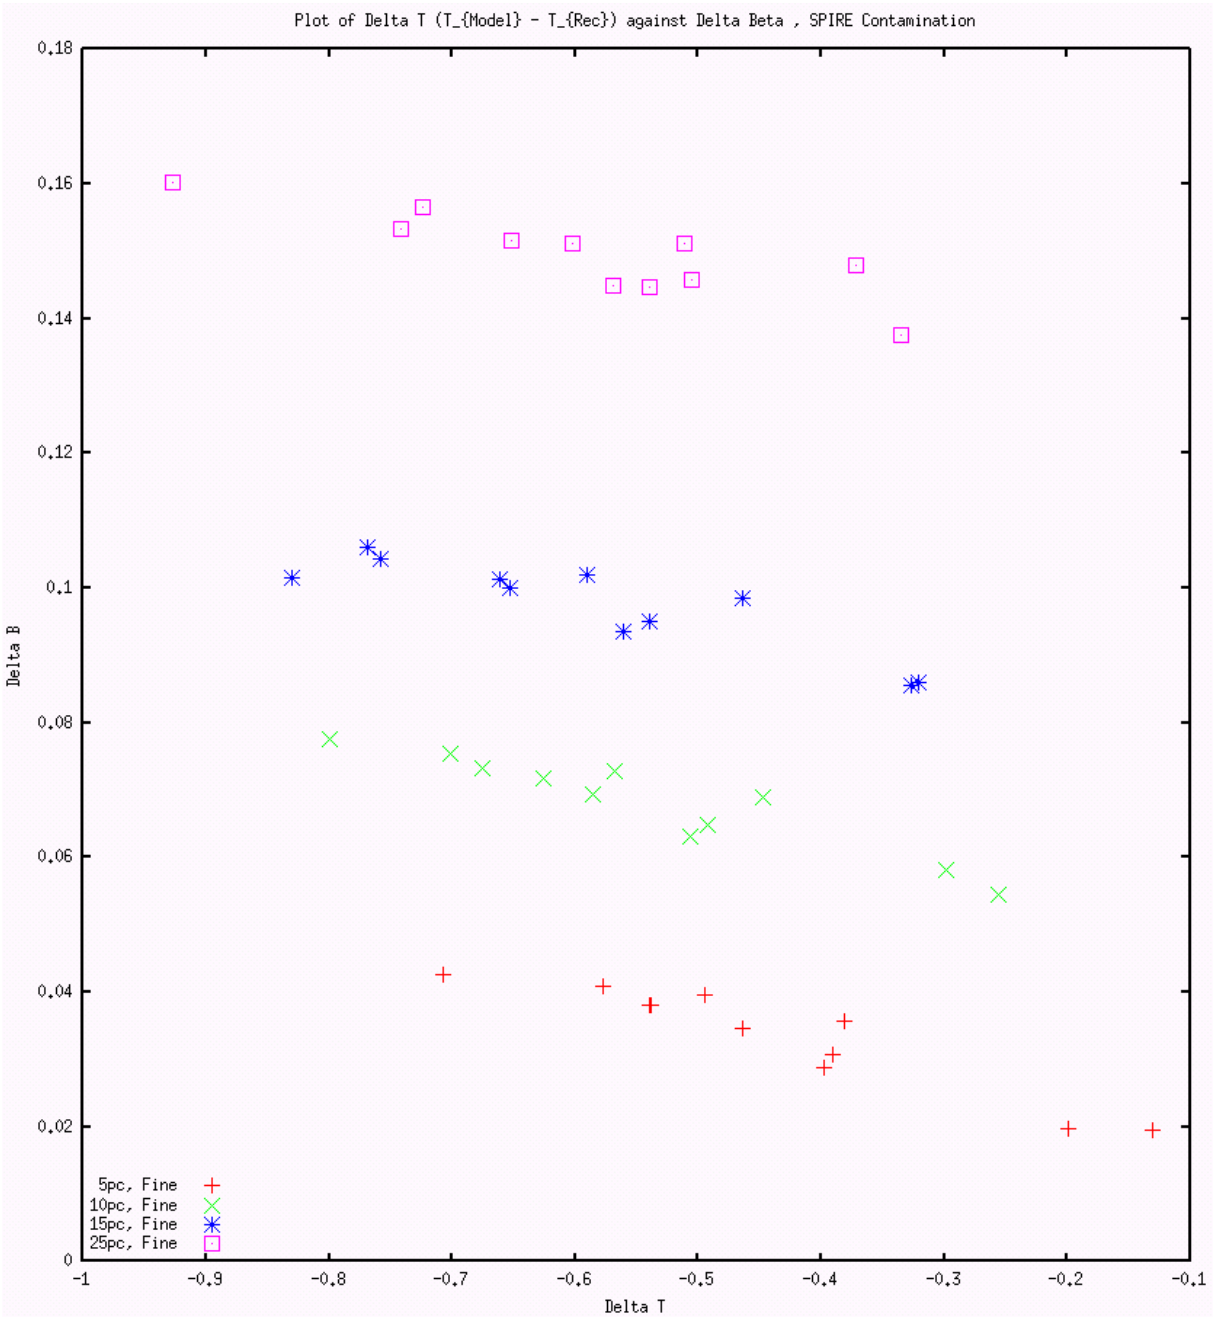

Figure 5.89: A diagnostic plot of ΔT against $\Delta\beta$ for $T_{model} = 40K$, simulated PACS/SPIRE observations.

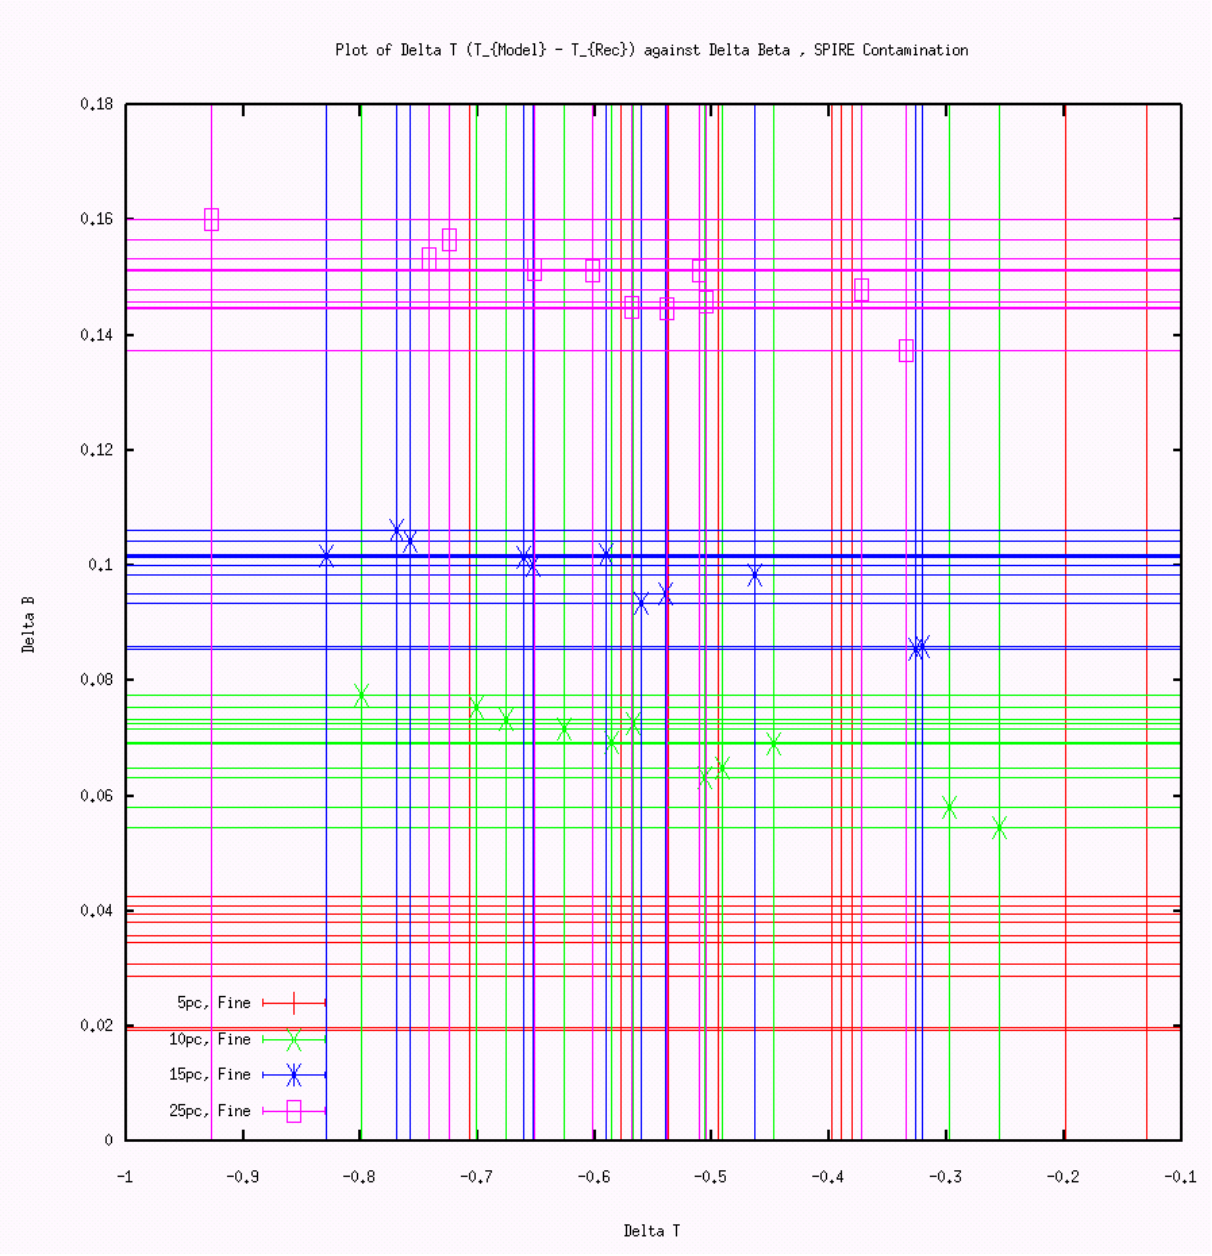

Figure 5.90: A diagnostic plot of ΔT against $\Delta\beta$ for $T_{model} = 40K$, simulated PACS/SPIRE observations with errorbars.

5.7 Conclusions from the greybody fitting

5.7.1 Conclusions from initial simulations

The initial simulations examined the effect of *LOCAL* and *GLOBAL* contamination. Locally, as the greybody temperature increased, we saw an increasing range of recovered temperatures as line contamination is applied. We also saw that line contamination in the $250\mu m$ band caused the recovered temperature to be driven higher than the model whereas line contamination in the $350\mu m$ and $500\mu m$ bands drove the fit temperature lower. The fit errors also increased and the range of recovered temperatures increased (to around 100K for a 25K model). These changes in recovered temperature were magnified as the greybody temperature increased. The variation in β_{rec} show no real diagnostic information, bar a slight tendency for the range of its variation to narrow with increasing temperature. Where line contamination is applied to the shorter wavelengths, β_{rec} is lower than β_{model}. This effect is reversed where line contamination is applied to longer wavelengths.

For global SPIRE line contamination, we note that ΔT is much lower across the simulations. In local simulations ΔT has a range of around 100K, and it is reduced to about 40K for global line contamination. Globally, we see that errors in T increase as T increases, and ΔT decreases and becomes negative. This shows the importance of having the peak within the waveband to constrain the fit. The magnitude of ΔT increases for larger β_{model}. In the global simulations, the effect of constraining β_{rec} to a range of 1 – 2 leads to differing behaviour at the limits. At the lower limit, β is driven high and T_{rec} is lower than T_{model}, therefore $\Delta T \geq 0$. ΔT increases as the line contamination increases in this case. At the upper limit, we see this temperature behaviour invert but the increase in ΔT with line contamination is

weaker. At the central value of β_{model} a turnover is seen. For lower temperature and line contamination simulations $\Delta\beta \leq 0, \Delta T > 0$. At higher temperatures the sign of ΔT changes - a turnover is seen at around 15K where the peak of the greybody moves out of the band. As line contamination increases above 10% this turnover is lost, and $|\Delta T|$ increases. For global non-SPIRE contamination, the behaviour is similar, although the magnitude of ΔT increases. As β_{model} increases, so does the recovered temperature in comparison to the global SPIRE simulations. β_{rec} shows inverted behaviour for global non-SPIRE simulations compared to SPIRE simulations. Lower β_{model} simulations show a smaller variation in β_{rec}.

Overall, for SPIRE line contamination we see a tendency for β_{rec} to increase and the fit temperature to decrease for low β_{model}. At higher β_{model} this situation is reversed due to the limit. Central β_{model} values show a differing behaviour dependent on line contamination. For non-SPIRE contamination we see that β_{rec} decreases and the fit temperature increases.

5.7.2 Conclusions from fixing β_{rec}

For simulations where $\beta_{model} = 2.0$, global SPIRE contamination yields a negative ΔT. It drives the recovered temperature higher than the model temperature. For the other simulations the sign of ΔT is positive, and this is accentuated by lower values of β_{model}. For global non-SPIRE contamination, ΔT is almost exclusively positive, although its magnitude decreases for increasing β_{model}.

5.7.3 Conclusions from fixing temperature

For the 10K simulations, for global SPIRE contamination where T_{model} is less than the fixed temperature, we see that $\Delta\beta$ is positive and the sign inverts for higher temperatures. Where $T_{model} = 10K$ we see that $\Delta\beta$ tends to be negative. The best recovery of β_{rec} is when $\beta_{model} = 2$. For non-SPIRE simulations at 10K, fixing the temperature means it is less likely to be returned as a fit limited value, and the 15K simulations show slightly better recovery. At the 10K limit, line contamination has little effect on the recovery. As line contamination increases, lower β_{model} simulations at this temperature show β_{rec} driven slightly higher until we are at the upper β_{model} level. For the 25K simulations, we see that β_{rec} is not returned as the limiting value until $T_{model} = 15K$ - where fit recovery improves. Here the best recovery of β_{rec} is where $\beta_{model} = 1$. For non-SPIRE simulations, at 25K we see better quality recovery at the lower β_{model} simulations, and overall the recovery is better than the SPIRE band simulations. The effect of line contamination on the simulations mirrors the SPIRE simulations.

5.7.4 Conclusions from the expanded simulations

Improving the resolution of the simulations shows the turnover for all percentages, where it is seen to be a function of line contamination percentage, β_{model} and T_{model}. Low line contamination percentages tend to a higher $\Delta\beta$ value than higher line contamination simulations. As line contamination increases, the tendency of ΔT to be positive increases. Lower values of T_{model} are affected least by line contamination. The range of $\Delta\beta$ is roughly constant, but the values of $\Delta\beta$ show a linear relationship between line contamination percentage and $\Delta\beta$. There also appears to be a correlation, with higher β_{model} simulations showing increasingly negative val-

ues for $\Delta\beta$. Lower initial T_{model} simulations show that β_{rec} is increasingly recovered higher than β_{model}.

5.7.5 Conclusions from free-β expanded simulations

All of these simulations showed a distinct grouping by line contamination percentage. This grouping can be defined as being approximately $\Delta\beta \approx -0.02 \times LC$. The lowest line contamination percentage simulations show $\Delta T \leq 0$. As line contamination increases, so does ΔT and above 15% it is exclusively positive. The range of ΔT also increases with line contamination.

5.7.6 Conclusions from higher temperature expanded simulations

These simulations show a grouping according to line contamination percentage. ΔT is negative, but very small. Each grouping shows a $\Delta\beta$ which is positive and increases with increasing line contamination percentage. Within each percentage grouping, we see a range of ΔT which shifts very slightly as contamination increases.

5.7.7 Implications for future observations

We can now consider the implications of line contamination on line fitting by considering the two extreme cases from the original study, IRC+10216 and IC342, which contained the most and least line contamination respectively. IRC+10216

was a point-source calibrated object with $\approx 25\%, 16\%$ and 11% line contamination across the PLW, PMW and PSW bands respectively in the synthetic observations. IC342 was an extended-source calibrated object with just 3% contamination, solely in the PLW band. We shall use these two observations to consider realistic astronomical situations where line contamination may be a significant effect.

5.7.7.1 IRC+10216 and line contamination

This source was point-source calibrated, and to obtain information about the PMW band contamination the SLW and SSW data were joined together. This was performed using the SSW data at the overlap to obtain a combined SED across the range 15 to 54 cm^{-1}. From this we used SLIDE to obtain line information (see Chapter 4). This method of obtaining a SED across the waveband assumes that both SLW and SSW are equally corrected for beam size effects. For a true point source, there should not be any issues arising, but if the source is not a true point source then this assumption may be problematic, and may be seen as a discontinuity between the SLW and SSW. In order to fully understand whether such effects are significant, we explore the combined SED as well as the SLW and SSW separately.

IRC+10216, also known as CW Leonis, was studied using SPIRE and PACS data by Ladjal et al. (2010). It is a carbon star, surrounded by an optically thick dust shell and envelope. Mauron & Huggins (1999) explored this source in the optical and observed an extended envelope out to approximately 200 arcsec. We can examine how line contamination will affect recovered greybody parameters for an example source of similar properties to our observation of IRC+10216. We shall consider the case of 25% PLW band line contamination on fitting the SLW SED

only. This will simulate a SLW observation only of this source, with the PMW line contamination weakened as it lies across the boundary join. The SLW SED is approximately 15 – 32 cm^{-1}, and from Figure 5.6 we can see that around half of the line contamination would be lost. We can therefore consider this a reasonable approximation of the SLW-only SED as our first IRC+10216-like observation.

We can then separately consider 25% line contamination on the SLW and SSW together. These simulations were performed with a constant line contamination percentage for each run, and so we cannot mix line percentages across the bands. We have already considered line contamination across three bands, and so we shall ignore the line contamination in PSW as the smallest value. Here we are smoothing the line contamination across the PLW and PMW bands, in essence transferring the PSW contamination to the PMW band. This will enable us to consider the effect of line contamination on a second IRC+10216-like observation – again one where the line contamination is significantly weaker in the shorter wavelengths.

For this comparison, we examine the effects of line contamination of 25% in the PLW band ($500\mu m$) on a temperature of $20K$. This was the fitted temperatures obtained by SLIDE. We examine the effects on T_{rec} and β_{rec} from our simulations, where we allowed β to be a fully floating value.

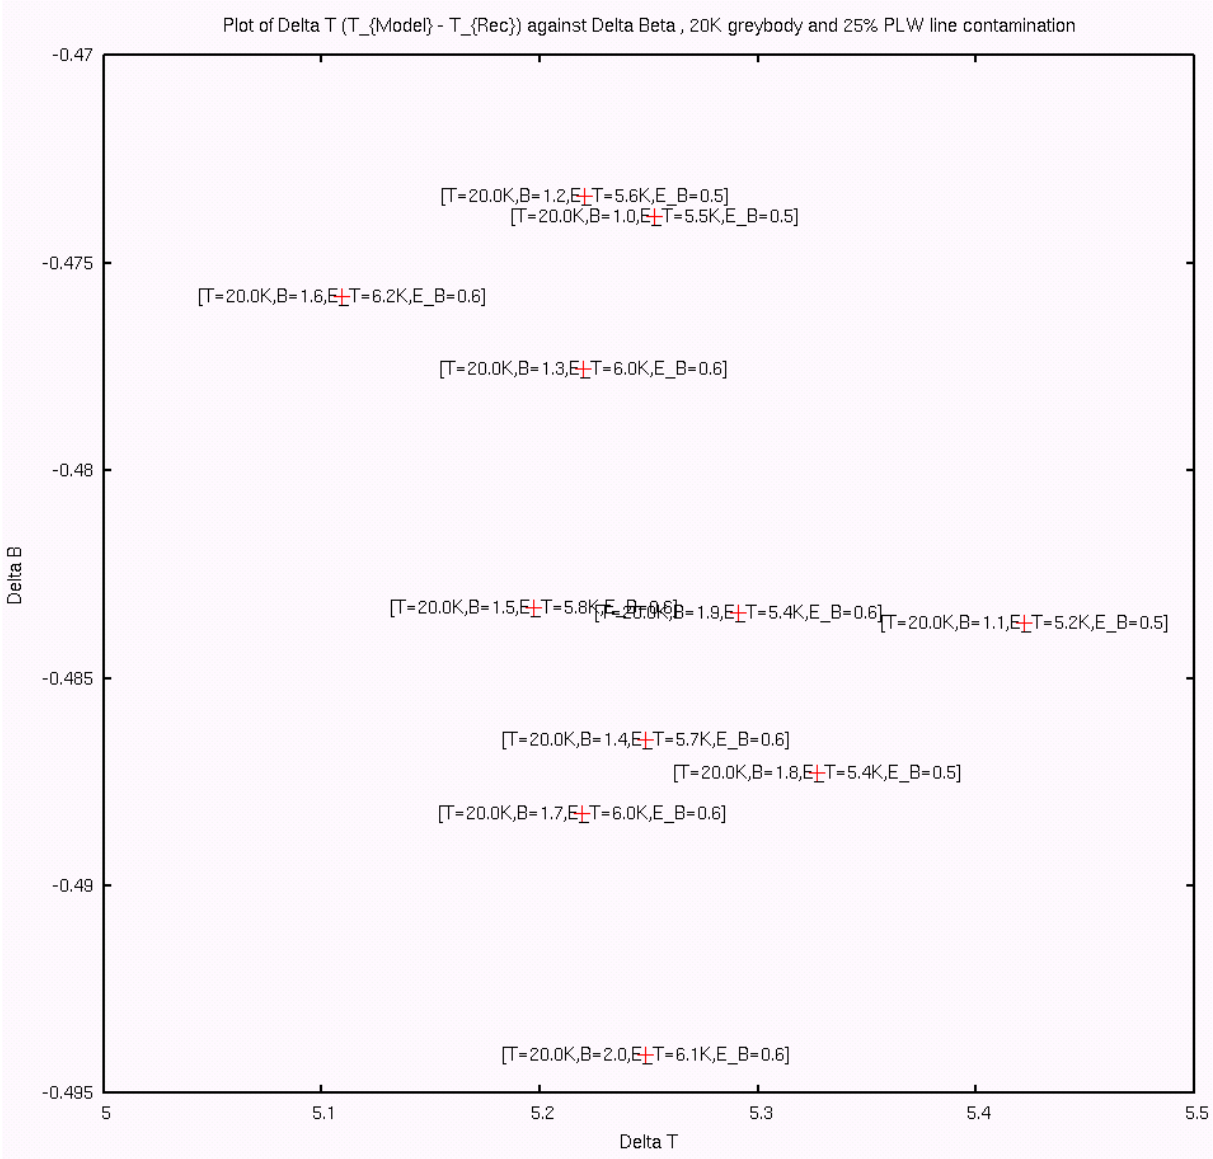

Figure 5.91: ΔT against $\Delta\beta$ for $T_{model} = 20.0K$, 25% PLW line contamination.

Figure 5.91 outlines the effect of 25% line contamination on the PLW band only, for the simulations at 20K. We saw earlier that the effects of line contamination on a global scale were such that the 1σ error bars in the fit masked $\Delta\beta$ and ΔT. Here, we see largely the same effect. However, in this case we see the error bars are comparable in magnitude to the amount of offset in both T and β. As a result, any offset in β and T is almost exactly at the limit of the 1σ error bars in the

parameter. Higher values of T_{model} yield larger values of ΔT but correspondingly larger error bars. Therefore, this suggests that the effect of line contamination in this case is impossible to detect.

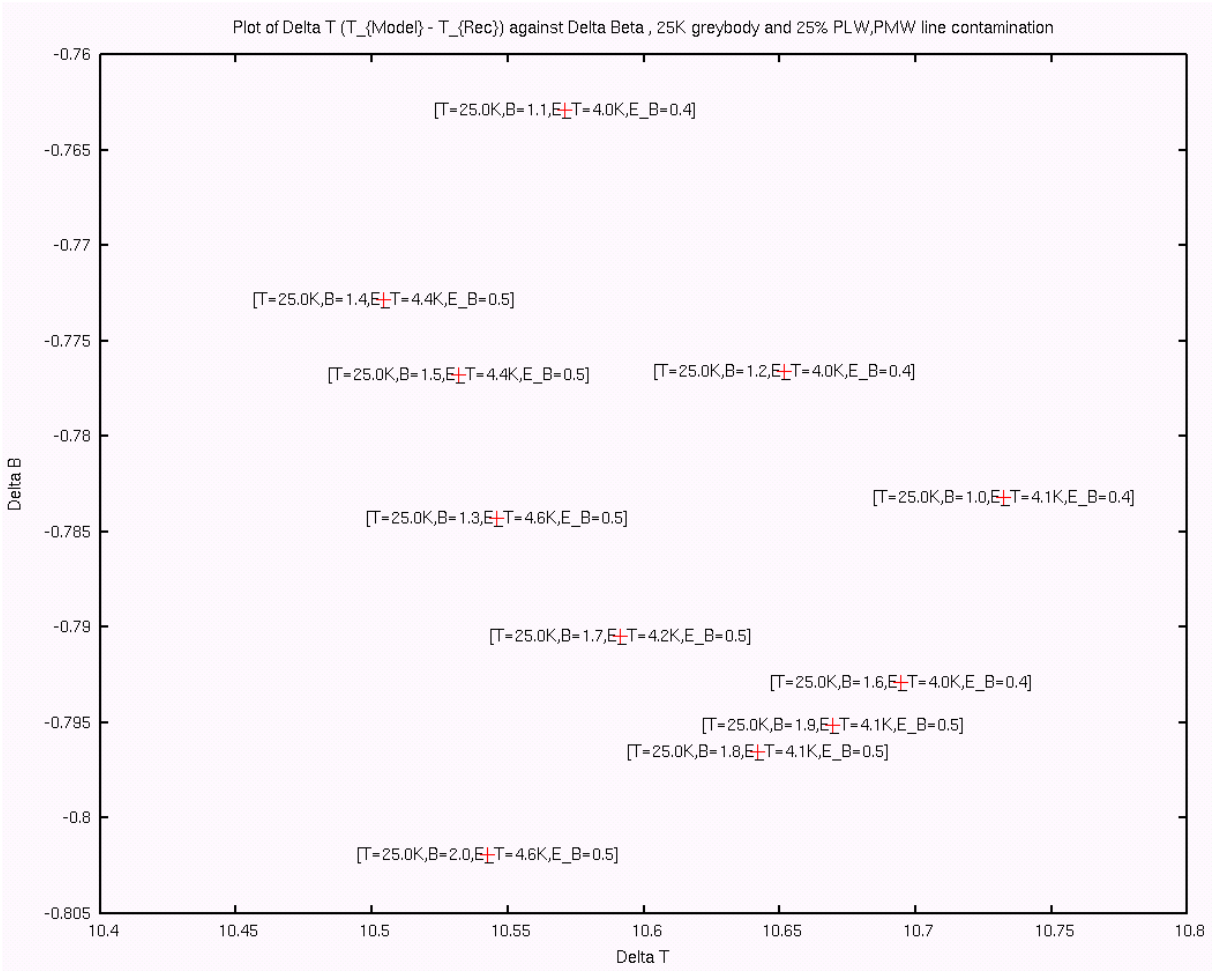

Figure 5.92: ΔT against $\Delta\beta$ for $T_{model} = 25K$, 25% PLW and PMW line contamination.

For a final comparison, we can take the mean value of the two greybody temperatures and compare the effects of 25% line contamination on both PLW and PMW. This will give us an approximate but effective comparison to the actual full SED of our example source with all line contamination. This will ignore the lesser line contamination percentage on PSW, as we have already examined the three band case. At these greybody temperatures, the peak emission is close to

the PSW value and so the line to continuum ratio would naturally be at is lowest value here for a fixed line value.

Figure 5.92 shows the result of this line contamination on a 25K greybody. Here we see the effects of line contamination on two bands simultaneously, and the values of ΔT and $\Delta\beta$ are outside the error bars of the fit. T_{rec} is lower than T_{model} by approximately 10K with error bars on the fitted temperature of about 5K. This means $\Delta T \approx 2.3\sigma$. Similarly for β_{rec} – it is higher than β_{model} by approximately 0.8 with error bars of around 0.5 meaning that $\Delta\beta \approx 1.6\sigma$.

With this we see the difficultly in predicting the effects of line contamination. Our simulations have included a range of line contamination percentage, greybody temperature and emissivity and there are few simple trends that seem to encompass the whole range. Our simulations did not include those with mixed line contamination percentages in different bands, as this would have increased the simulation time significantly. However we can identify some trends that have become apparent. For single bands, it appears that line contamination does affect T_{rec} and β_{rec} but the associated errors are large enough to mask it. Line contamination in a single band is largely masked as it appears the fit uses the remaining datapoints to constrain it and introduces larger parameter errors.

Line contamination applied to three bands appears to yield a similar effect. In this case, the addition of contamination to over half the number of datapoints in the SED does not unduly affect the recovery of T and β, but rather increases the 1σ errors. We suggested earlier that in this case the effect is to largely shift the SED in the y-axis, which would not affect these parameters.

However in this third case, we have applied line contamination to two central

bands of our five band SED. The line contamination in the PLW and PMW band has introduced a change in the values of β_{rec} and T_{rec}, but below the 3 σ level. When using only five wavebands to provide the range of the SED, in this case the distortion of the SED yields a change in the values of T_{rec} and β_{rec} to the order of $\sim 1.5 - 1.6\sigma$ for β_{rec} and $\sim 2.3 - 2.4\sigma$ for T_{rec}.

In this case, with the maximum line contamination percentage present, the change in T_{rec} and β_{rec} is less than 3σ. This means that the effect is still undetectable.

5.7.7.2 IC342 and line contamination

IC342 is a spiral galaxy located in Camelopardis, and is an extended source. This source is one of the largest spiral galaxies in diameter in the sky, and Jarrett et al. (2003) give a diameter for this source as 372 arcsec from the Two Micron All Sky Survey (2MASS, Skrutskie et al. 2006) data. The observation was calibrated using SPIRE's extended-source calibrated mode, and when the level of line contamination present was investigated both SLW and SSW bands were considered separately, without combining the two SEDs into a single SED. SLW and SSW have differing beam sizes and so to accurately combine both SEDs to obtain the PMW information would have required detailed knowledge of the source extent. This meant that for the extended-source calibrated observations only PLW and PSW were considered. This source showed the least degree of line contamination, with only 3% present in PLW. We shall therefore consider the 5% line contamination in PLW as the example.

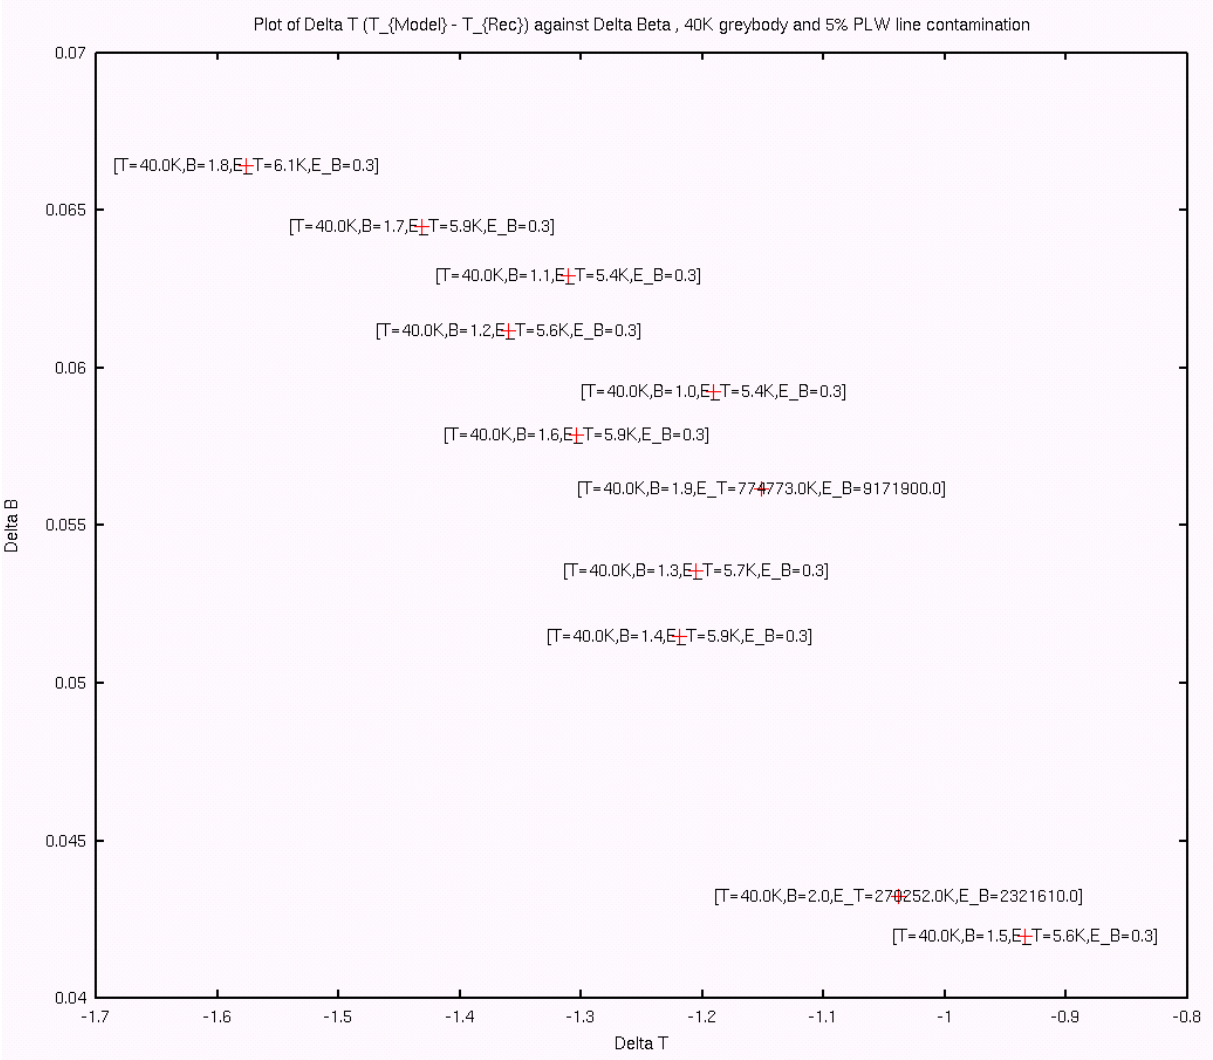

Figure 5.93: ΔT against $\Delta\beta$ for $T_{model} = 40K$, 5% PLW line contamination.

This small degree of line contamination would seem to suggest that any effect will not be detectable. The results are shown graphically in Figure 5.93. In this case, SLIDE fitted the SLW SED to a greybody of 55K, and so we shall examine the effect of 5% line contamination on the higher temperature simulations we performed.

Here, we see a range in ΔT of around $0.1K$ to $-0.5K$, and a range of $\Delta\beta$ of between 0.05 and 0.1. Error bars in T are an order of magnitude greater than

this value, and likewise with the errors in β. The error bars for T and β increase as β_{model} increases, and for $\beta_{model} > 1.9$ they become several orders of magnitude higher. From this we conclude that the effect of line contamination on a single band cannot be readily detected, as the changes are masked by the large error bars introduced.

5.7.8 Line contamination implications

For the two astronomical sources we have considered as extreme cases, the effect of line contamination has not been detectable. Adding line contamination to a single band, or a number of bands, has not resulted in a shift in T or β greater than 3 σ. This cannot be considered as a detectable change.

Chapter 6

SPIRE FTS observations of pre-stellar cores

6.1 Data

The data presented in this chapter are guaranteed time observations of pre-stellar cores taken by SPIRE. These were observations of L1521E and L1521F in Taurus, and L1689B in Ophiuchus. For L1689B, there were three observations taken: on-source, off-source and one in an intermediate position between the two. These observations were taken in SPIRE's low spectral resolution mode, with an observational resolution of around $0.89cm^{-1}$. The observations for L1521E and L1521F have only on-source positions.The details are listed in Table 6.1 below, which includes source name and Observation ID (OBSID) and the RA and Declination position of the central pixel of each observation. The OD (Observational Day) is also listed. Extended source flux calibration was used, as the sources all show evidence for extension.

Source Name	OBSID	RA	DEC	OD
L1521E	1342191211	67.309173	26.236110	288
L1521F	1342191208	67.165830	26.859419	288
L1689B (on source)	1342191221	248.701306	-24.634189	288
L1689B (off source)	1342191222	248.701335	-24.716494	288
L1689B (intermediate)	1342191223	248.658774	-24.634234	288

Table 6.1: SPIRE observations of pre-stellar cores.

Figures 6.1 to 6.5 shows the central pixel observations, with regular pipeline processing. We reconstituted the observations by taking the central SLW and SSW pixel observations, and rescaling the SLW observation to remove any discontinuity. A preliminary greybody fit was made and is overplotted.

Source Name	OBSID	T_{fit}	β_{fit}
L1521E	1342191211	11.8K	2.0
L1521F	1342191208	11.8K	2.0
L1689B (on source)	1342191221	13.5K	2.0
L1689B (off source)	1342191222	15.2K	2.0
L1689B (intermediate)	1342191223	14.7K	2.0

Table 6.2: Greybody fit parameters to prestellar-core data with regular pipeline processing.

Figure 6.1 shows the central pixel data for L1521E, with the greybody fit overplotted. Flux levels for this source are the weakest for the three on-source observations. The SLW data do not show a well defined continuum, and the SSW data are dominated by large error bars. There is also a ripple to the continuum. The greybody temperature recovered was approximately 12K. The fit is shown overplotted, and the poor quality of fit for this core is shown.

Figure 6.2 shows the central pixel data for L1521F, again showing an overplotted greybody fit. Flux levels for this core are higher than the previous core, and the SLW data show less distortion and more closely approximate the greybody slope expected. This source also shows higher error bars in the SSW region than the SLW, with pronounced underlying ripple in both regions.

Comparing Figures 6.1 and 6.2 to the data for the on-source observation of L1689B in Figure 6.3, we note that the L1689B on-source observation is of better quality. The SLW and SSW regions show a lesser degree of underlying ripple, and the error bars for both regions of the SED are of comparable size. A visual comparison of the greybody fit and the SED for these three sources shows that this source has a significantly better quality fit. For the remainder of this chapter I will concentrate on L1689B, which has the most data, and also the highest signal-to-noise ratio.

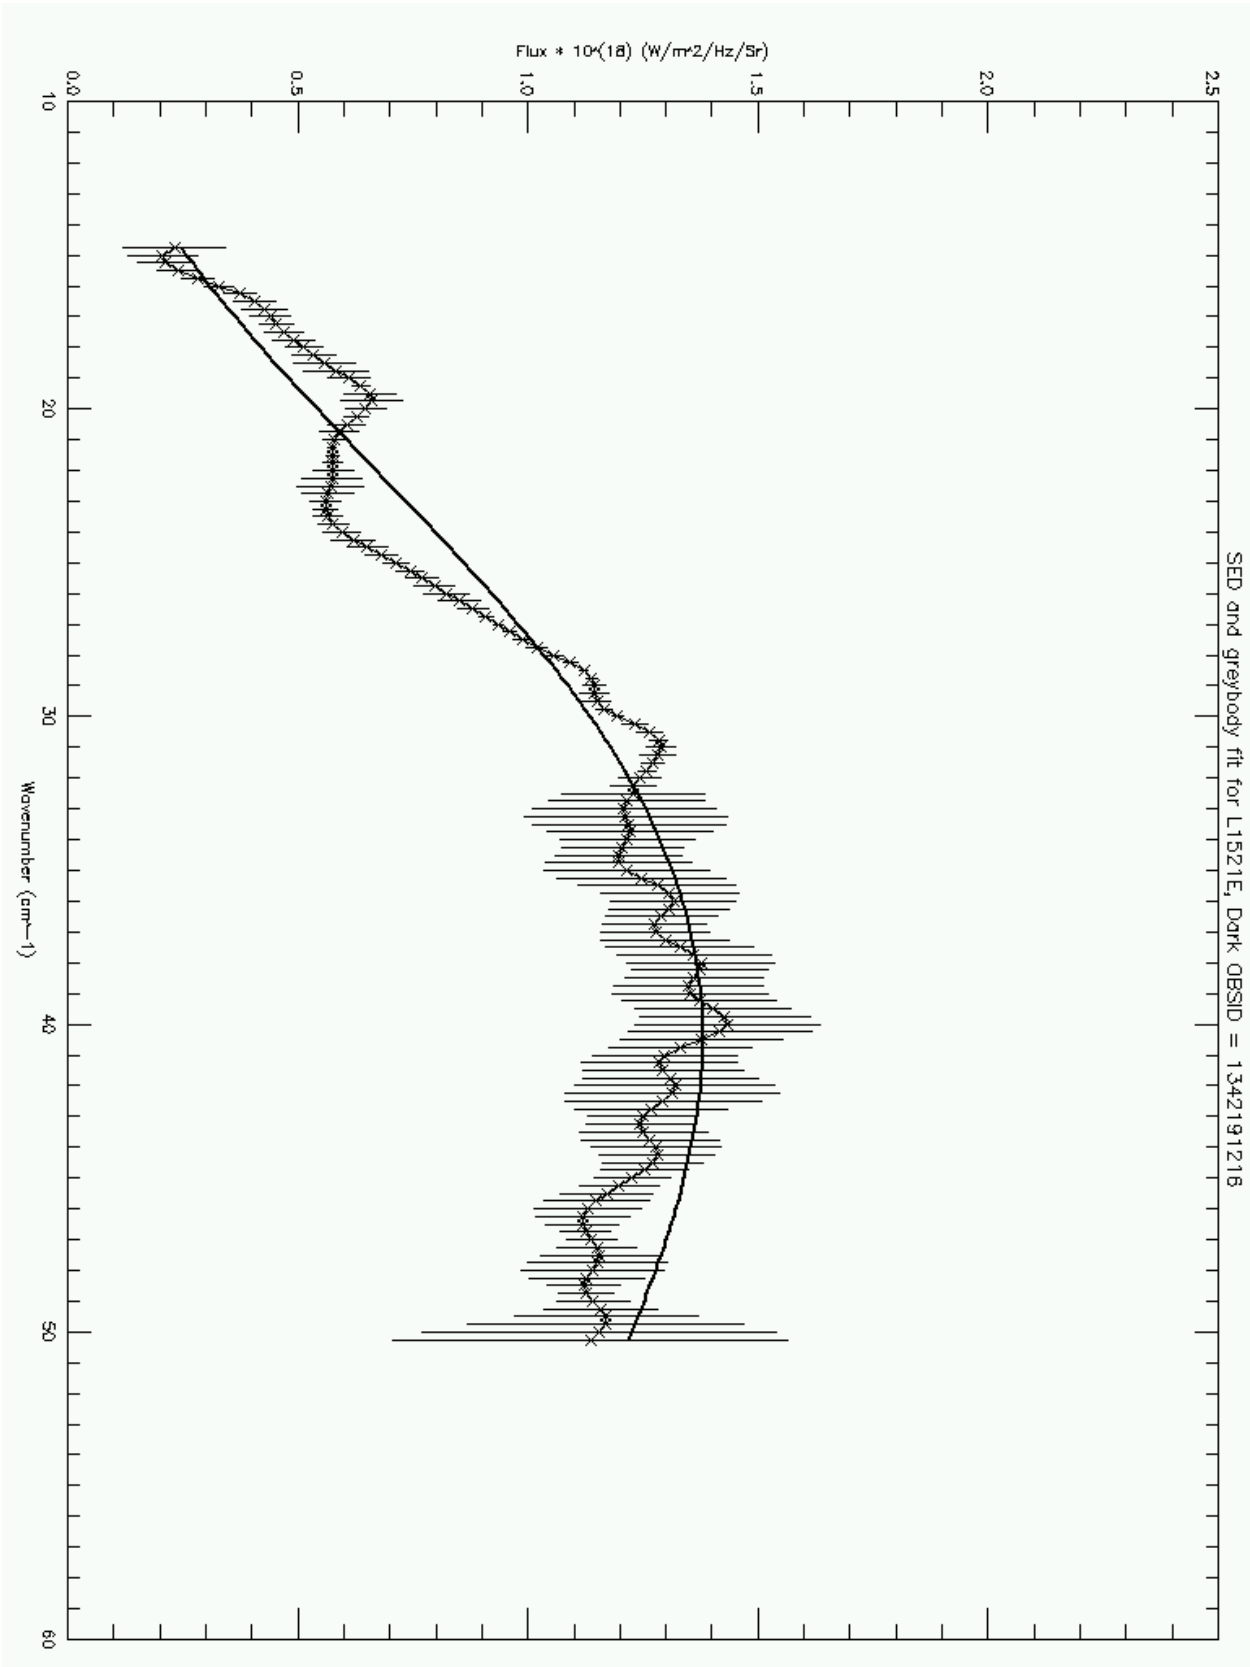

Figure 6.1: PSC observation of L1521E, central pixel only.

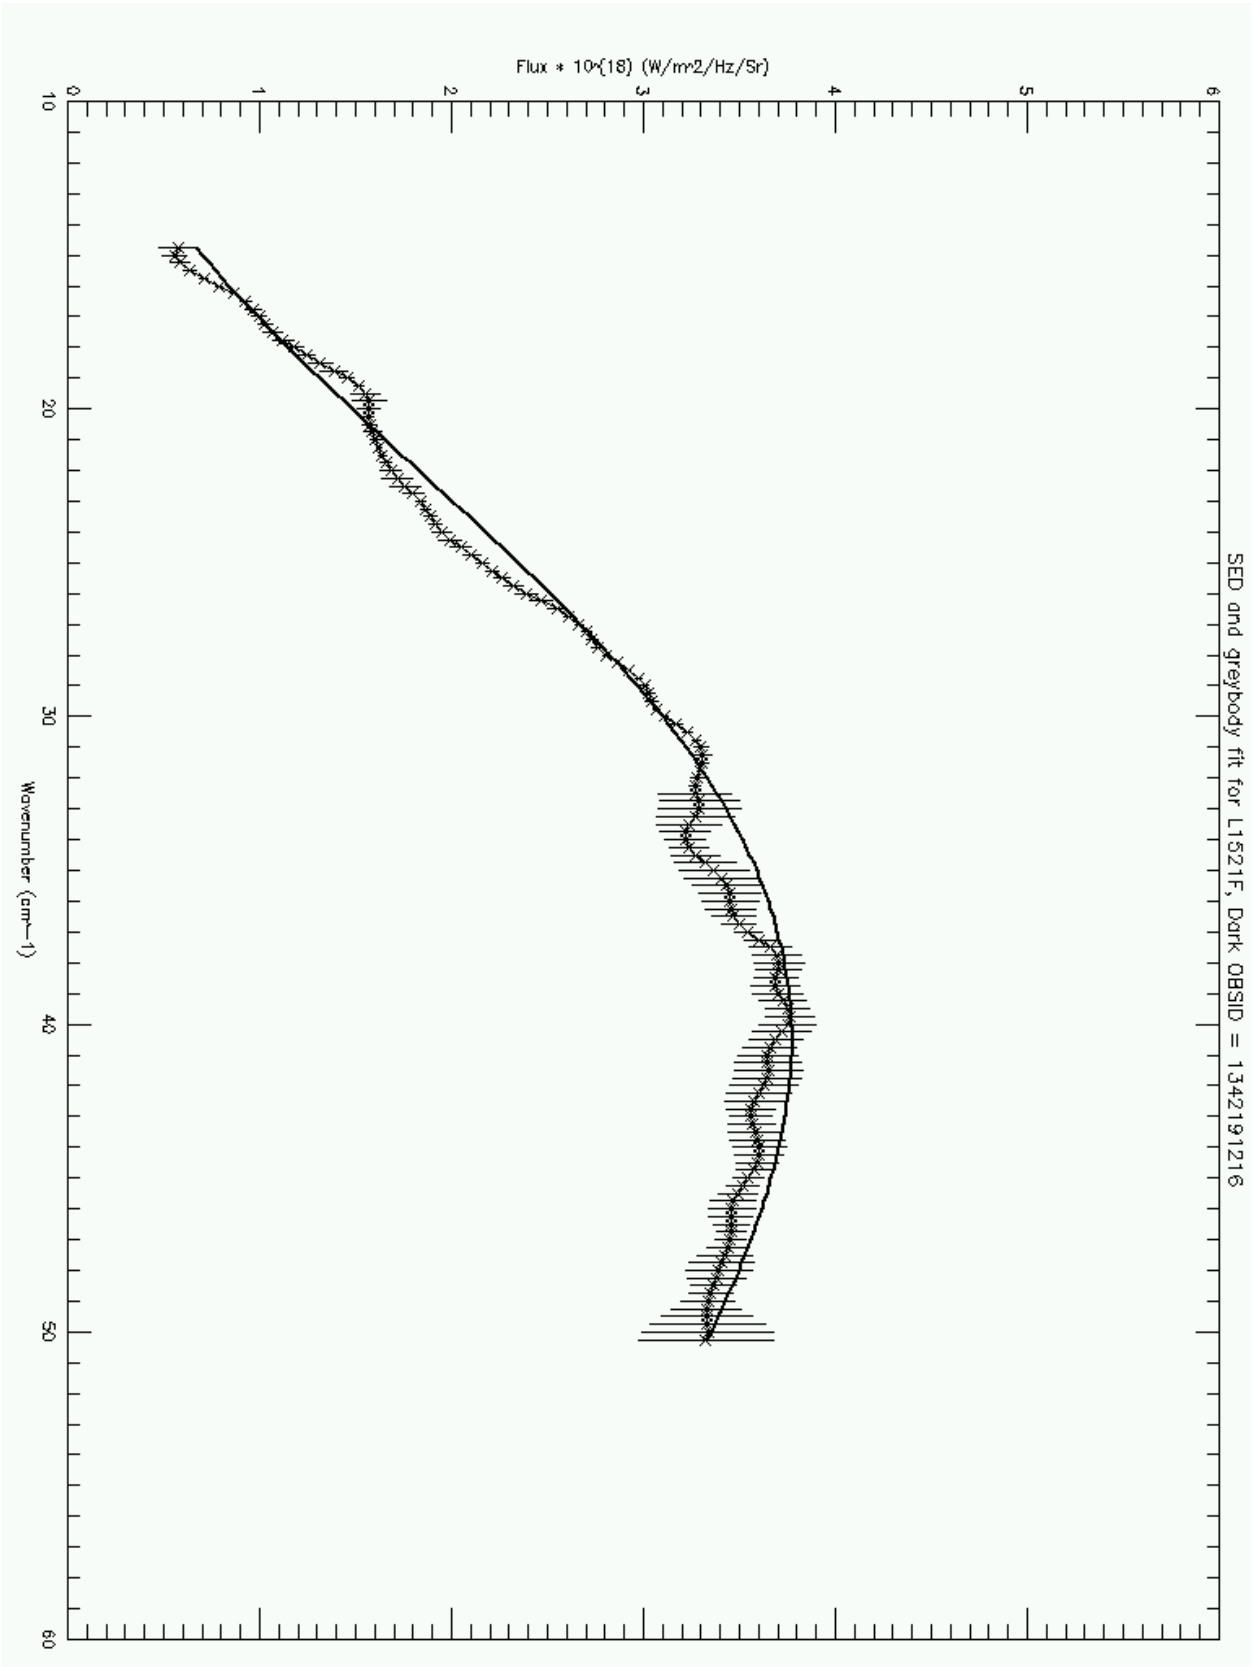

Figure 6.2: PSC observation of L1521F, central pixel only.

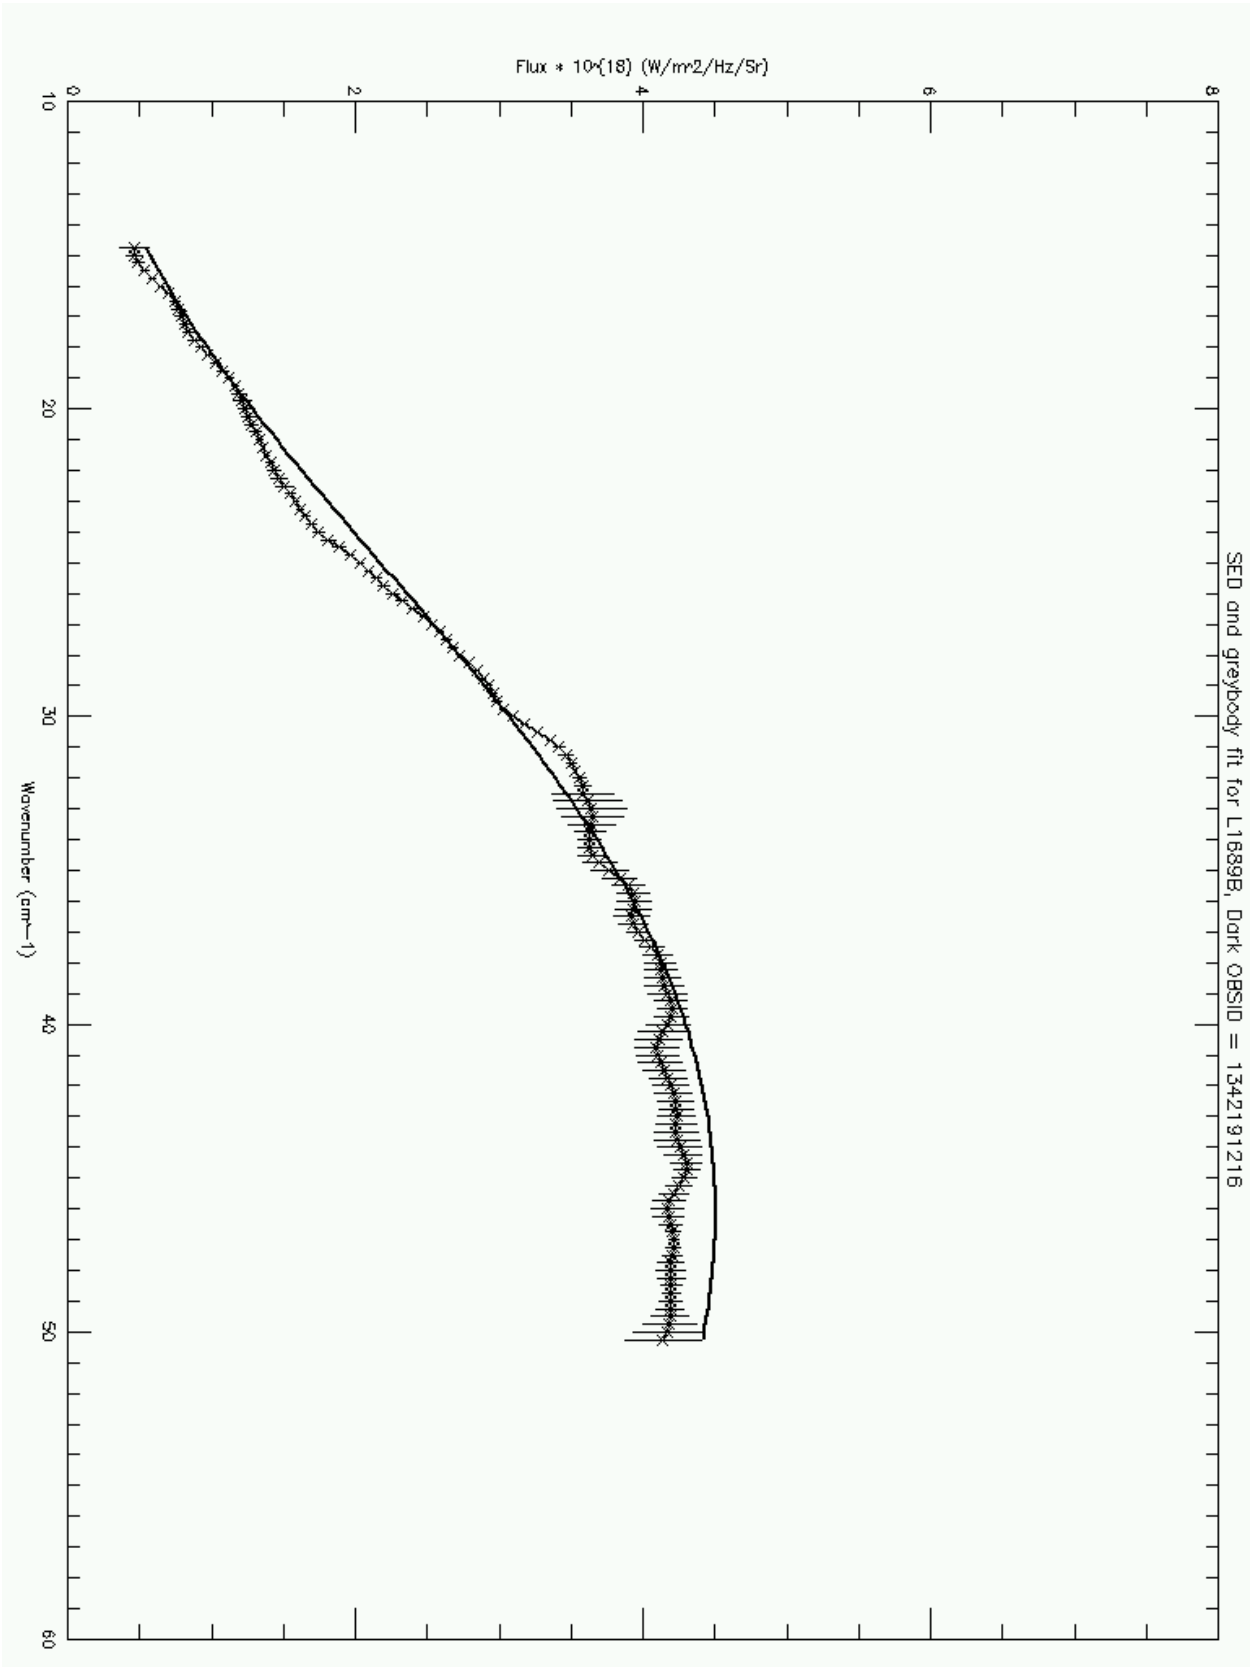

Figure 6.3: PSC observation of L1689B, central pixel only.

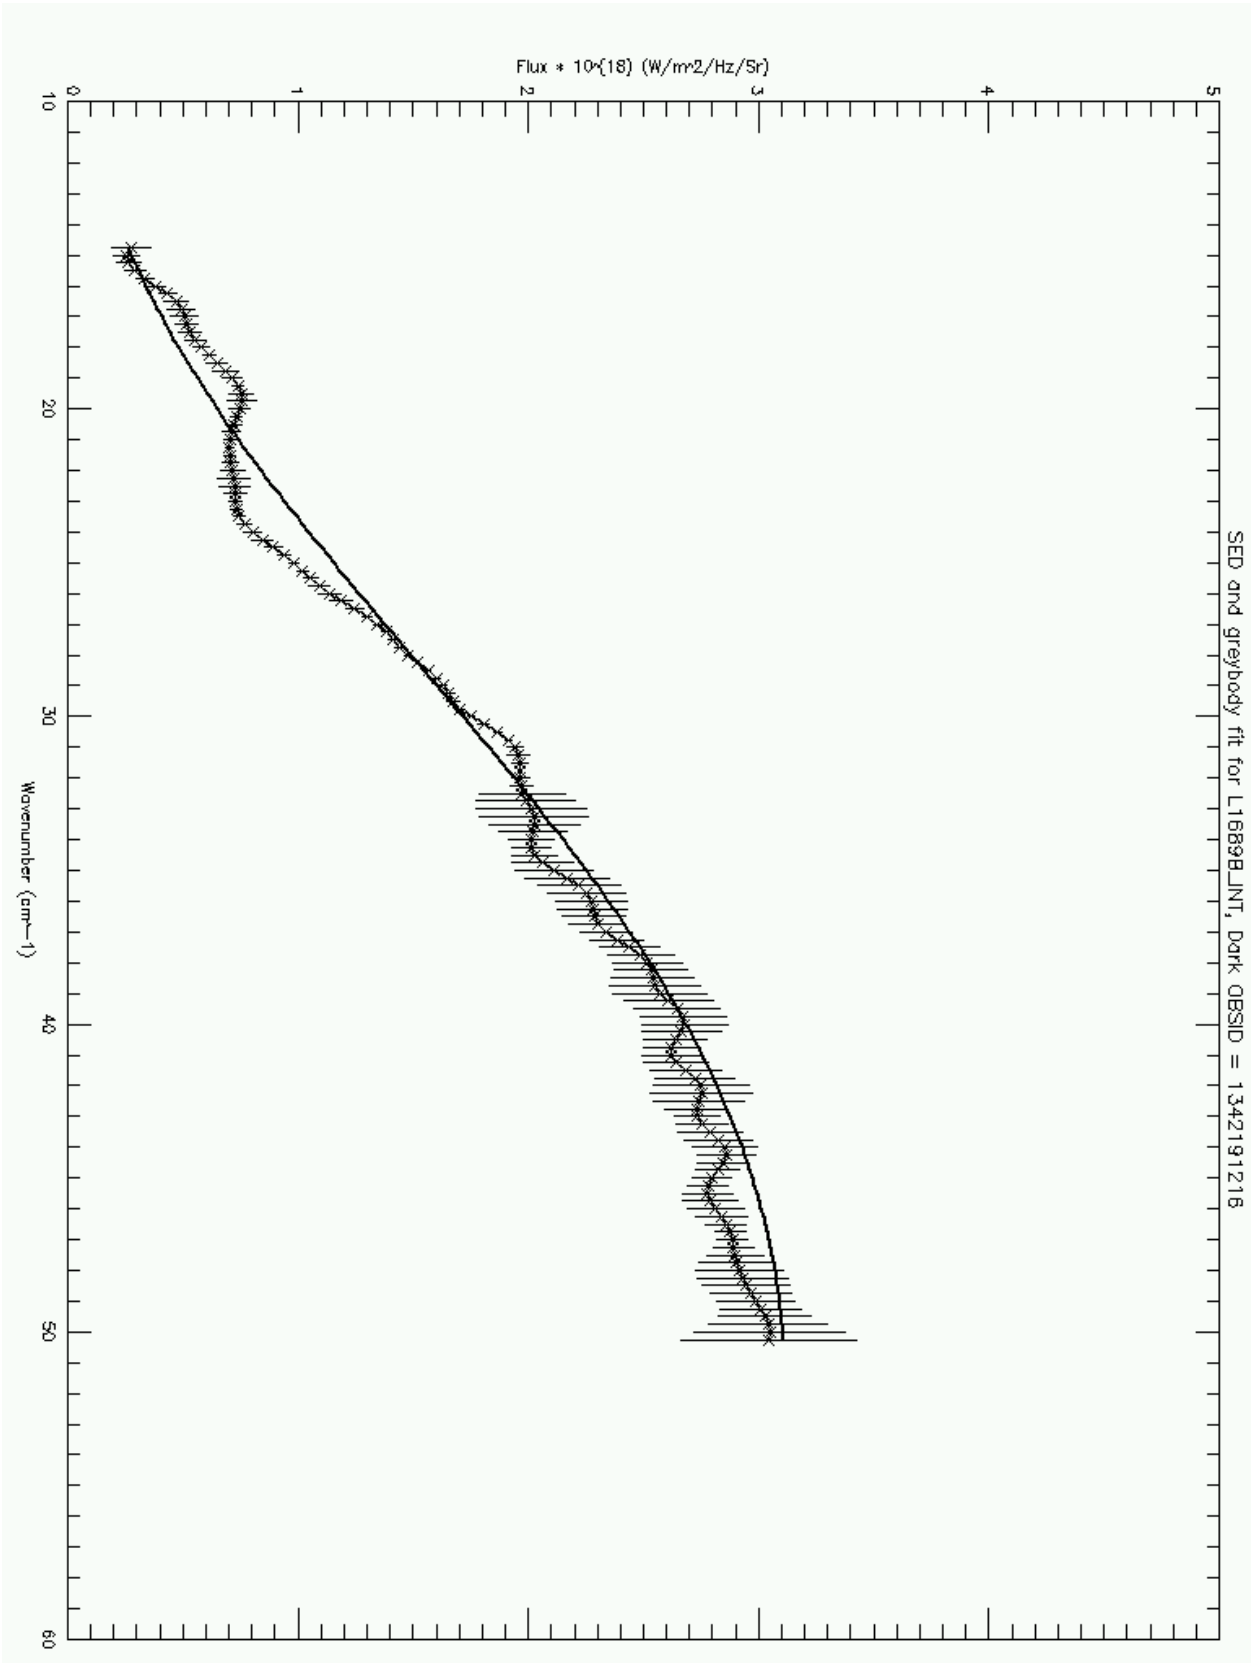

Figure 6.4: PSC observation of L1689B (middle position), central pixel only.

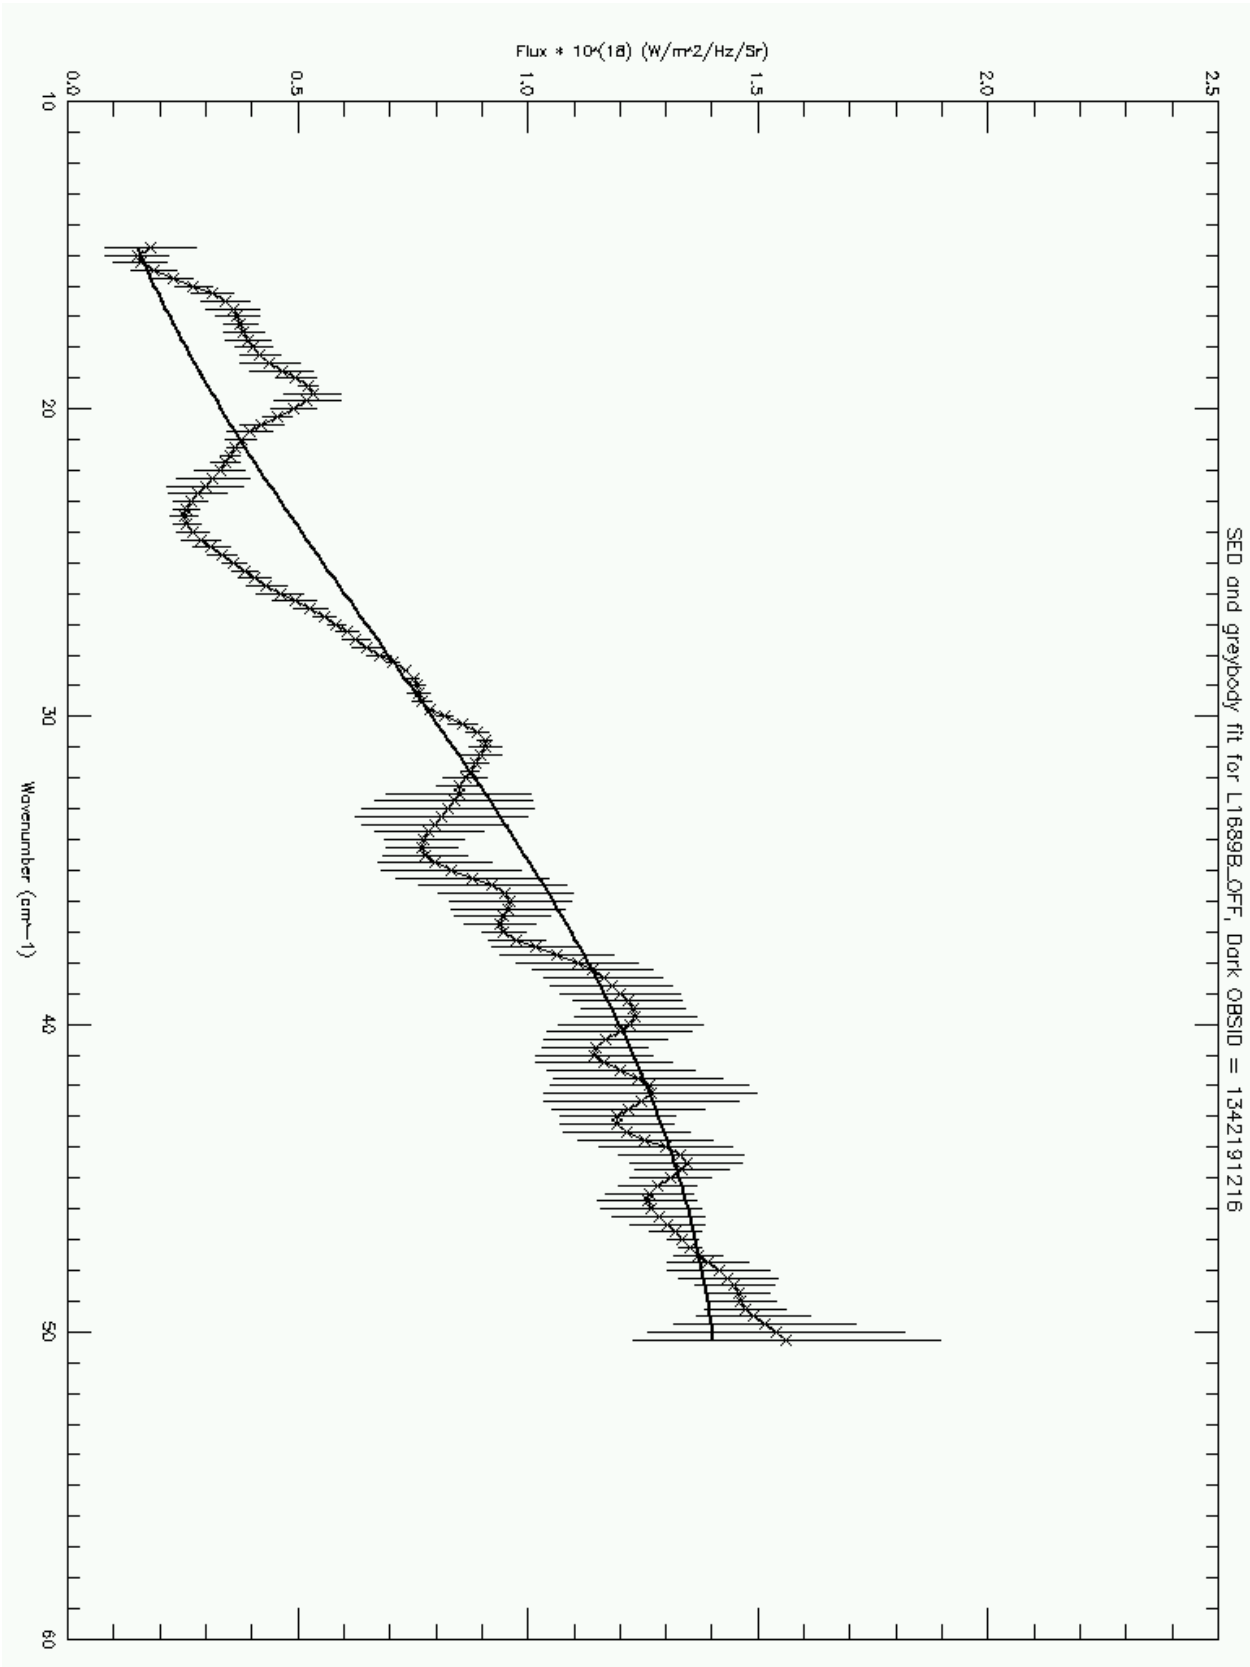

Figure 6.5: Off source PSC observation ofL1689B, central pixel only.

6.2 Previous work on prestellar core L1689B

Ward-Thompson et al. (1994) examined the radial flux density profile of L1689B, and found that they were unable to fit a single power-law profile over the full extent of the core. Figure 6.6 shows a log-log plot of the radial flux density profile. It was noted that for the outer regions (defined as approximately greater than twenty arc seconds) that the flux density follows $S_\nu(\theta) \propto \theta^{-1}$. Inside this region, $S_\nu(\theta) \propto \theta^{-0.25}$.

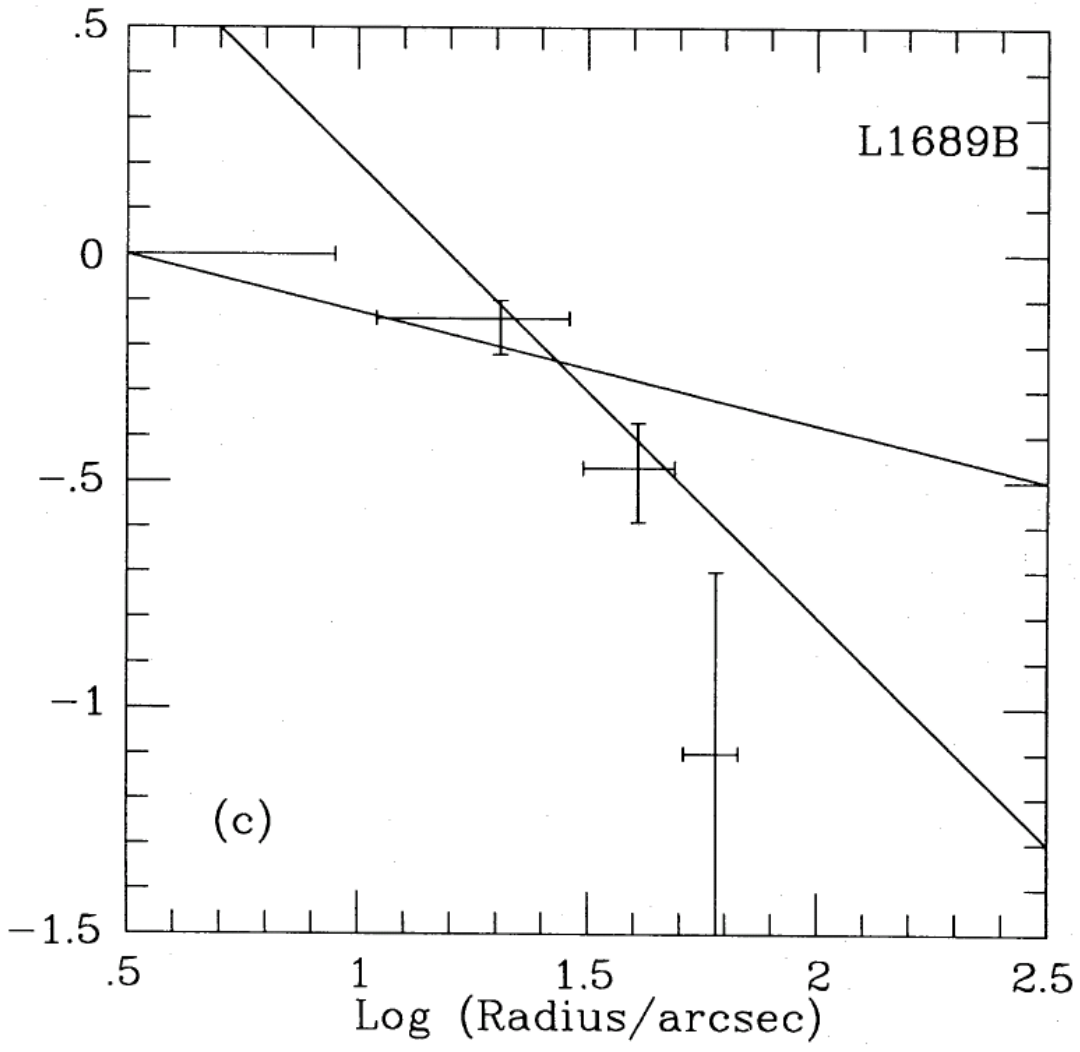

Figure 6.6: Two dimensional flux density of L1689B, from Ward-Thompson et al. (1994).

These profiles are two-dimensional and based on the observations, and so care must be taken to convert these into three dimensional radial density profiles. Ward-Thompson et al. (1994) assumed the cores were isothermal, and with spherical geometry. These arguments were made on the basis of the lack of a central source of heating within the core. From this, it was concluded that these two flux density profiles equated to radial density profiles of $\rho(r_{inner}) \propto r^{-1.25}$ and $\rho(r_{outer}) \propto r^{-2.0}$. This means that although the outer regions showed a radial density profile as predicted by the model of Shu (1977), the inner regions had a flatter profile. Also, the profile seen for starless cores was not in conflict with profiles observed in more evolved cores, and that earlier studies that seemed to agree with the Shu (1977) model may have been victims of poor observational resolution hiding the central flatter region.

From this two-stage radial density profile, it was surmised that either the cores were not yet evolved enough to have reached an isothermal profile or that this profile is affected by other processes not yet considered – such as magnetic support. The L1689B observation included a greybody fit to the SED and obtained a value for the greybody temperature of $T_{gb} = 15K^{+3}_{-7}$ using a value for $\beta = 2$. L1689B proved to be one of the most luminous prestellar cores, with L $\sim$ 1-2 $L_{\odot}$. The mass of L1689B was deduced using

$$M_{gas} = \frac{S_{\nu} D^2}{B(\nu, T)\kappa_{\nu}} \tag{6.1}$$

where S_{ν} is the measured flux density (using the flux density within the FWHM at 800μ m), D was taken as 160 parsecs, B is the blackbody function and κ_{ν} is the dust mass opacity. This was taken to be $\sim 0.01 cm^2 g^{-1}$ using $\beta = 2$ and a ratio of gas to dust mass of 100. T for the blackbody fit was T_{gb}. From this, a mass of

$1.2M_{\odot}$ was obtained.

Andre et al. (1996) detailed the results of a study of L1689B using the IRAM 30m telescope, mapping the continuum emission at $1.3mm$ from this core. The core was identified as being elongated in the east-west direction, and has an aspect ratio of around 0.7. This study also confirmed the dual layer structure of radial column density, and concluded that L1689B is around $0.6M_{\odot}$ and of size 0.045×0.067 pc. This agreed well with the figures from Ward-Thompson et al. (1994). For the central region showing the flatter density profile, they claimed a mass of $0.3M_{\odot}$ and an approximate diameter of $8000AU$.

This mass, and the approximate 30:1 ratio of volume density between L1689B and the surrounding cloud led Andre et al. (1996) to claim that it is in fact a pre-stellar core, that is self gravitating and with a permanent structure. Using the conclusions from Ward-Thompson et al. (1994) they surmised that this core is either in the early phases of, or about to start, collapsing into a protostar.

Analysing the core of L1689B was performed by mapping its radial density profile, using circular and elliptical annuli. Circular mapping is the simplest, and elliptical mapping allows for the elongated nature of the core by using ellipses of the same aspect ratio as the core. These maps were then used to compare the results of simulating various models of the pre-stellar core's density profile by generating artificial, realistic observations of the models. Simulating single- and double-power law structures for the density profile showed that a single power-law structure for the core was not adequate to reproduce the results. Although a single Gaussian model was used for numerical simulations of protostellar collapse at the time, Andre et al. (1996) deduced that a two-Gaussian model was needed, comprised of an unresolved point source and a broader elliptical Gaussian. Since the central

point source has very weak emission, the possibility of a centrally formed protostar was discarded and explained as possibly the begininng of an instability prior to collapse.

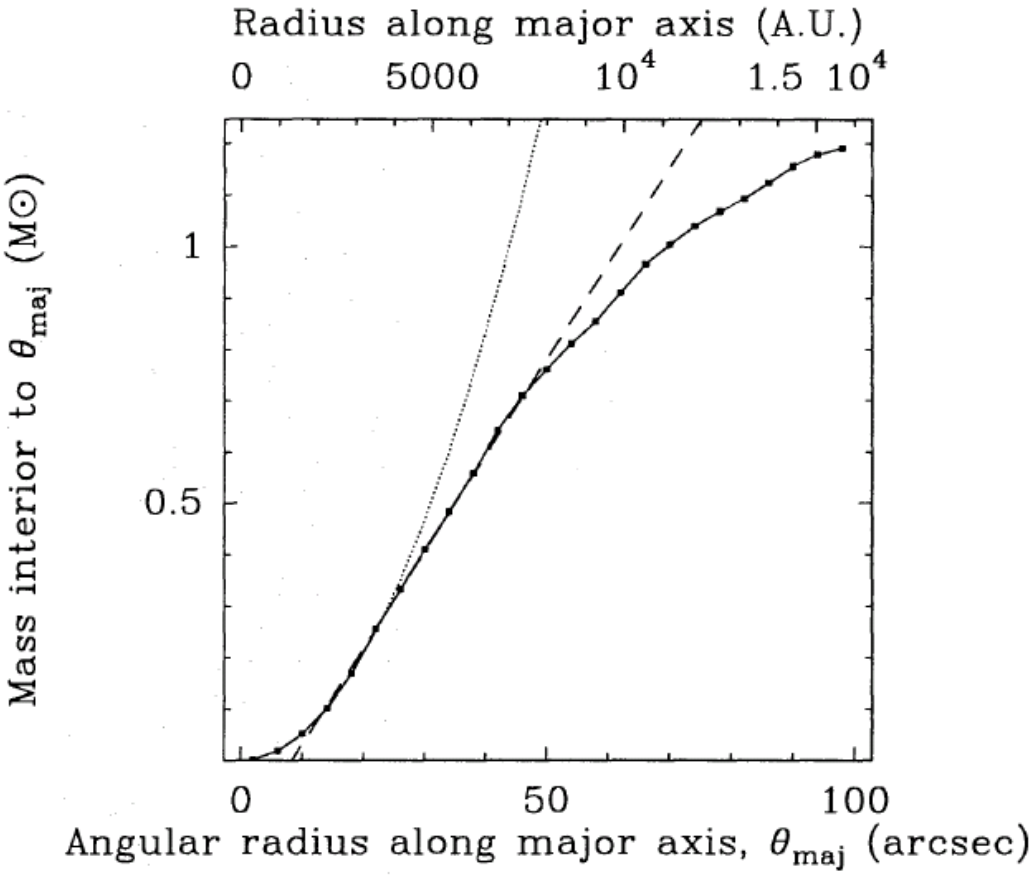

Figure 6.7: Mass distribution curve for ellipsoidal areas for L1689B (Andre et al., 1996).

Figure 6.8 shows the results of their modelling. This figure shows the radial flux density profile for L1689B. The flux density is normlised, and is an average of the observed emission at $1.3mm$. This figure has four sub-figures, labelled (a) – (d). The first two are for assuming a cicular symmetry and averaged over circular annuli, and the last two are for eliptical annuli. Within each pair, the first is plotted in linear space, the second in log space.

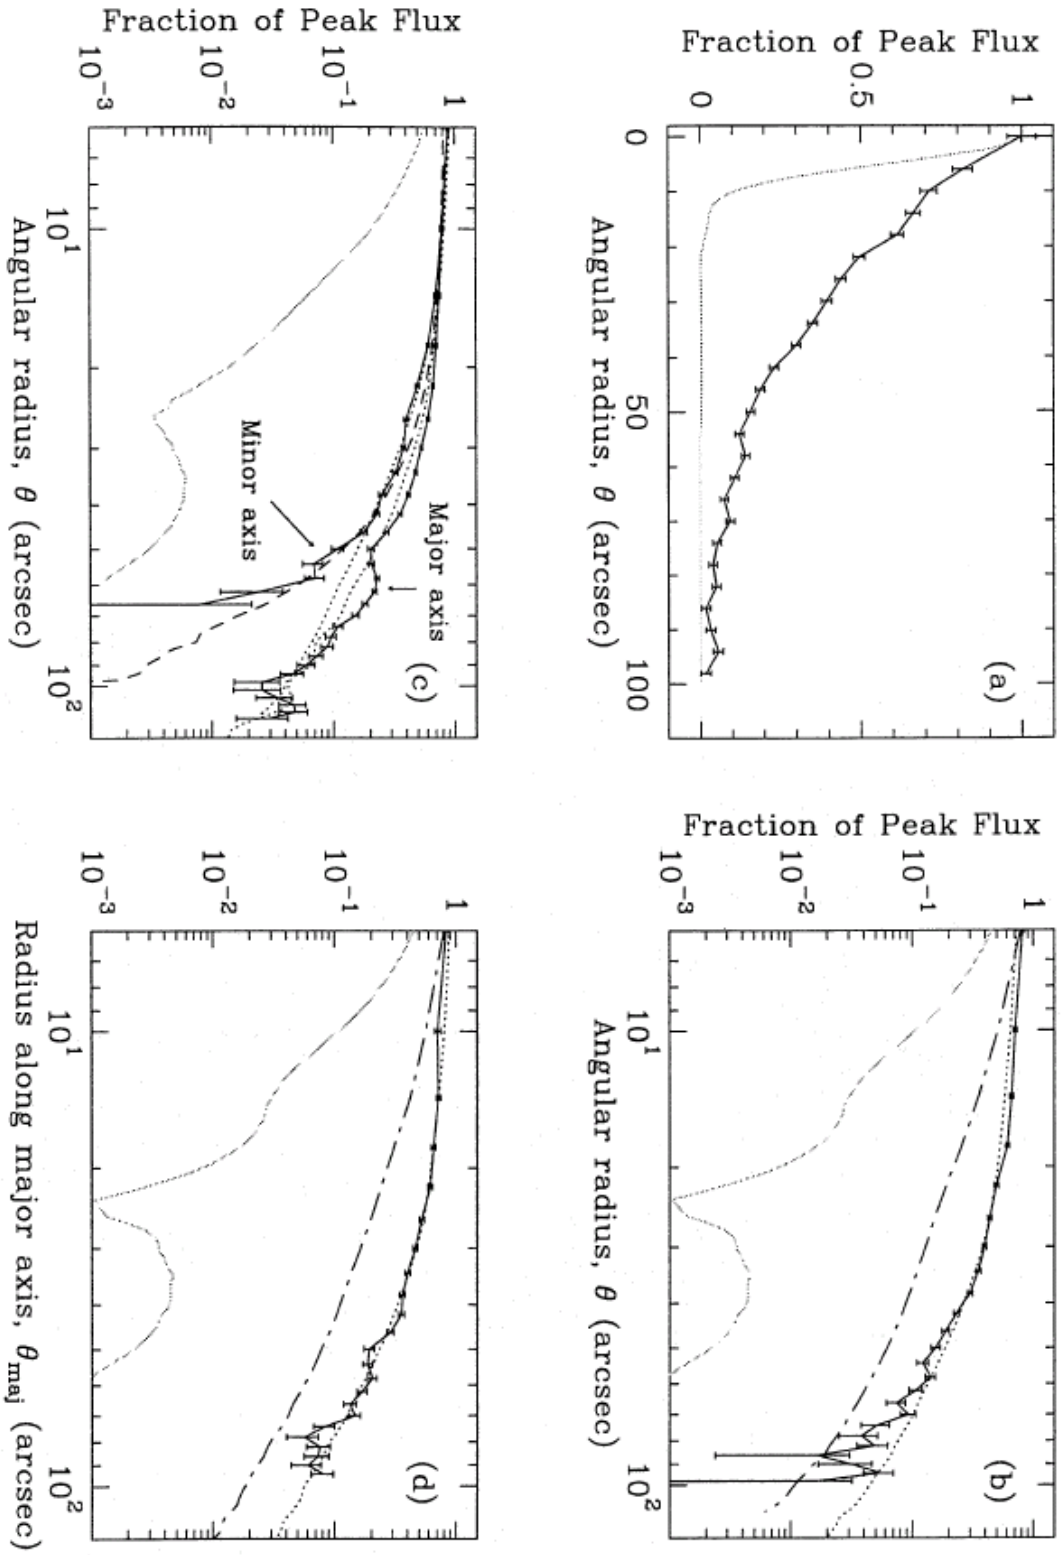

Figure 6.8: Results from radial density modelling from Andre et al. (1996) for L1689B.

On each plot the dotted solid line is the radial profile. The dotted line represents a two-tier structure, with a flatter inner profile to the $\rho(r) \propto r^{-2}$ outer structure. The dashed line is a single power-law strucure, and the thin dashed line is the simulated model profile. It can be seen from this figure that the two-tier structure is the better fit.

The mass distribution was deduced by using the submillimetre continuum emission, as it is optically thin. Andre et al. (1996) calculated the mass M within a radius R using,

$$M(R) = \frac{S_{1.3mm}(\theta)D^2}{\kappa_{1.3}B_{1.3}(T_{Dust})}, \tag{6.2}$$

although we note this is of a similar form, and equivalent to that used by Ward-Thompson et al. (1994). For T_{Dust} a mean of three earlier observations of other authors was used, of $18K$. From this, Figure 6.7 shows the derived mass distribution curve for elliptical annuli. This curve plots the angular radius θ and the mass within this radius. We note here the solid line is the derived mass from the observations, the dotted line is the fit for the inner flat region where the mass is proportional to θ^2, and the dashed line is the fit for the outer region where it is proportional to θ. The deviation of the dashed and solid lines where θ is greater than around 50 arcsec was attributed to flux loss due to the map being of finite size. From this, the estimate of the mass in the inner, flat region was calculated to be $\sim 0.3M_{\odot}$. An upper limit for the mass, for a region of two arcminutes, was given to be just over $2M_{\odot}$. The inner flat region of the core was believed to be around $4000AU$

Two possible explanations were given for this strucure. The first involves a

Bonnor Ebert sphere structure, a thermally supported self gravitating gas and dust structure. This is not considered isothermal, as the model considers the effect of the local ambient Interstellar Radiaiton Field on an isothermal Bonnor Ebert sphere. Conversely, an isothermal model which involves ambipolar diffusion as a magnetically supported core was considered. The conclusions drawn from the modelling were that the structure had an inner radial density $\rho(r) \propto r^p$ where $p = 1.0 \ldots 1.4$ and an outer radial density $\rho(r) \propto r^{-2}$. The magnetically supported model was less favoured as the magnetic field support needed ($\approx 80\mu G$) was higher than seen in the literature for the region.

Observational evidence for the freezing out of CO molecules within the central core of L1689B was outlined by Redman et al. (2002). This was performed using JCMT observations of the $C^{17}O$ $J = 2 \rightarrow 1$ line, using the structure of the hyperfine components of the line. This particular CO line was chosen as it is more optically thin and has a more complex line structure than the $C^{18}O$ molecule. With a strong line detection, the structure enables the calculation of optical depth effects, leading to greater confidence in the abundance obtanined. This molecular line shape consists of nine closely grouped hyperfine structure lines. Any effect of optical depth will be seen in the ratios of the weaker to the stronger components, and so the line profile contains information about the underlying optical depth.

Redman et al. (2002) calculated a conversion factor for $C^{17}O/H_2$, and used the $C^{18}O$ data from Jessop & Ward-Thompson (2001) as a cross check. Both values agree well for the column density. This was cross checked by using observational data from Evans et al. (2001) to consider the dust column densities across the core, which correlate with those obtained by Ward-Thompson et al. (2002).

Draine (2006) studied the effect of grain size on the spectral index β. From

a power-law relationship for grain sizes in protostellar discs, the resultant spectral index is related to the power-law index p of grain size a for a size distribution $dn/da \propto a^{-p}$. This yields a spectral index such that $\beta_\lambda = (p-3)\beta_s$ in the sub-millimetre region. β_s is the spectral index of the small particle limit, and was expected to be the same value of $\beta \sim 1.8$ for the interstellar medium.

In summary, the properties of L1689B as a prestellar core have been observed, modelled and debated since the 1980's and its structure has been explored thoroughly. It is believed to have a central, flat density profile inner core of size around 5000 AU (Andre et al., 1996; Ward-Thompson et al., 1999; Redman et al., 2002) with significant depletion of CO within this core (Jessop & Ward-Thompson, 2001; Redman et al., 2002). The depletion is believed to be due to the freezing out of CO molecules onto dust grains (Redman et al., 2002) inside the core whose temperature is around 10K (Redman et al., 2002; Ward-Thompson et al., 2002).

The mass has been variously estimated to be between 0.6 and 2.1 $M_\odot$ in the central region (Ward-Thompson et al., 1994; Andre et al., 1996), which is believed to be rotating compared to a quiescent outer envelope (Redman et al., 2004). Dynamically, the core is believed to be close to starting, if not already started, the process of contraction towards a protostar (Andre et al., 1996). We now consider the SPIRE observations and what further information we can obtain from them, in the context of the earlier observations of this prestellar core.

6.3 L1689B Data Processing

We use the fact that there are a number of pixels in the SPIRE FTS to analyse the region surrounding the core of L1689B (see Chapter 4). This involves processing the

observation in HIPE for all pixels, and then identifying the pixels whose positions are coincident. In this way we can create a wideband observation of the source at the seven coincident pixel positions as previously shown in Section 4.1.

The SPIRE observation for L1689B used in this chapter was a single pointing, extended source calibrated observation, taken on Observing Day 288 (OD288), 2010 February 26th, with an observational ID (OBSID) of 1342191221. The observation is low spectral resolution and unapodised, with a nominal resolution of $\sim 1cm^{-1}$, and an actual observational resolution of $0.89cm^{-1}$, supersampled at $0.25cm^{-1}$. The observation was processed in HIPE v6, and calibrated using a dark-sky observation (known as a 'dark') to correctly allow for the temperature of the telescope.

The observation was calibrated in the spectral domain by processing both observation and dark as if both were actual observations, with no associated dark sky (Polehampton, 2010). The dark used was also observed on OD288, with an OBSID of 1342191216. Each pixel was reduced separately. As the dark calibration takes place in the spectral domain, a separate HIPE program was produced to calibrate each pixel by subtraction of the dark from the observation. The resultant observation and associated error for each pixel was output to a text file for each pixel. The error for each pixel in this case was taken as the quadrature sum of the errors for the observation and the dark.

In this way, we consider the original observation to be

$$O_{original} = O_{source} + O_{telescope} \tag{6.3}$$

where O_{source} is the observation of the source with no other emission present. If we consider the dark observation to be

$$O_{dark} = O_{telescope}, \tag{6.4}$$

and solely the telescope's own emission, then we can reconstruct the source only observation by

$$O_{source} = O_{original} - O_{dark}, \tag{6.5}$$

which is the reduction strategy we have chosen in the spectral domain (Polehampton, 2010). The subtraction means that we can then calcualate the error in the source-only observation from the error in the original observation and the error in the dark observation by quadrature, yielding

$$E_{source} = \sqrt{\left(E^2_{original} + E^2_{dark}\right)}. \tag{6.6}$$

Section 4.1 shows a representation of the position of each pixel in the pixel array, and Figure 6.9 shows that for SLW and SSW there are 7 pixels that are coincident that are labelled C1-C7. Pixel C1 shares a common centre for SLW and SSW central pixels, and pixels C2-C7 are such that the SSW pixel is completely contained in the SLW pixel area. We then overplotted these pixel positions on a SPIRE Photometer map observation of L1689B to show the position of each pixel on the source, and this is shown in Figure 6.10.

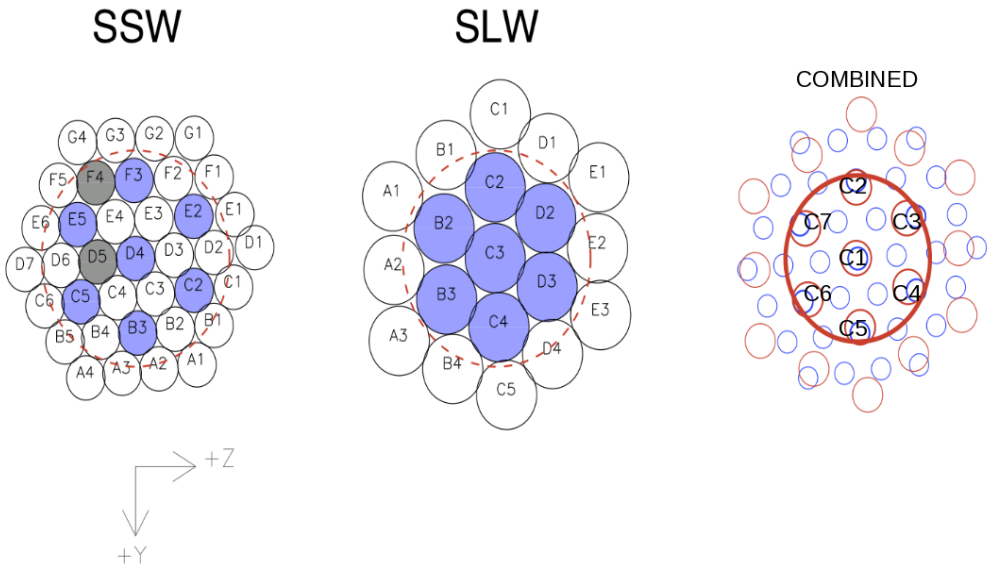

Figure 6.9: The SPIRE pixel array showing coincident pixels.

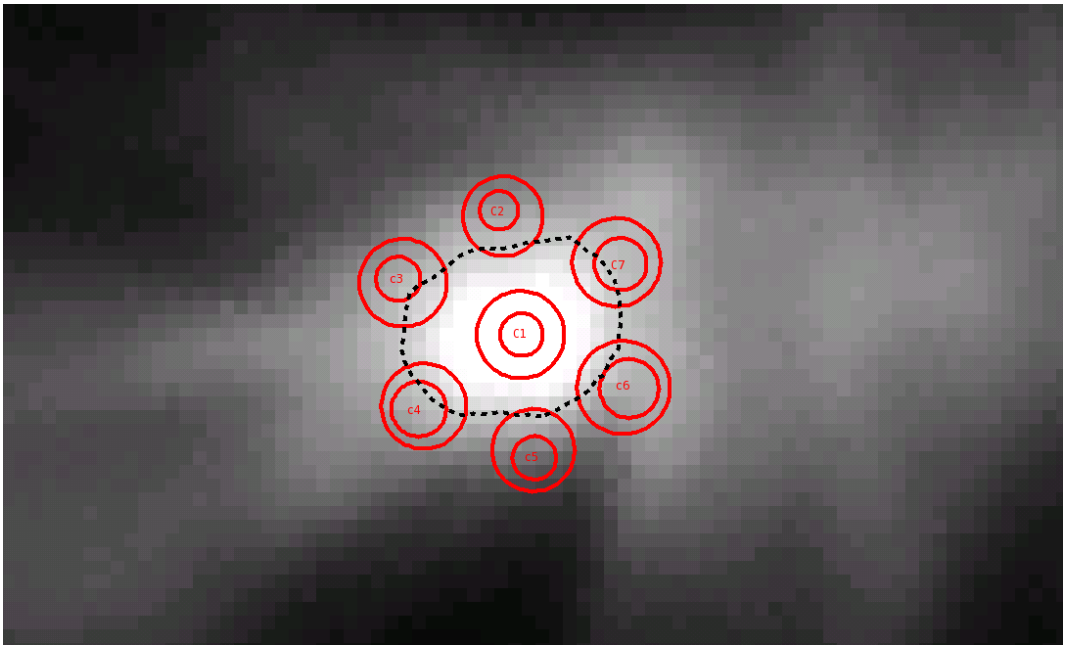

Figure 6.10: The coincident pixels overlaid on the photometric mapping observation of L1689B, the dashed line is an isophotal contour to outline the core morphology.

Each combined pixel was created from the relevant SLW and SSW pixel. With extended mode calibrated data, there is a likelihood of seeing a discontinuity between the SLW and SSW data as each pixel has a different beam size, and the beam size for each pixel shows a degree of wavenumber dependence. Discussion as to the best nature to correct for this is ongoing, and research continues on ways to adequately correct for this amongst the SPIRE FTS team. For these observations, we attempted to find our own way of correcting baseline artefacts on the observations by using the central pixel C1 for reference.

In this case, we created a combined observation for each pixel by splicing the SLW data and SSW data together. The centre of the overlap between each pixel is found, and the SLW data is scaled to match the equivalent position for SSW data. In this way, the seven combined pixel observations are created from the fourteen individual pixels involved. In each observation, there is some degree of ripple to the baseline as shown in Figure 6.11. This ripple has been discussed in earlier sections, and in this case it was deemed important to remove such ripple. The fringing function created for categorising SLIDE's effectiveness with real observations was used (see Chapter 4). We found the best result by using the central pixel data to characterise the ripple and apply this correction to all other combined pixels.

Using the greybody equation (Equation 5.2), we peformed a fit using a logarithmic form of the greybody equation (see Chapter 5). The free parameters were the temperature T and spectral index β, while Ω, the solid angle and f, the filling factor of the source within the observing aperture, were combined into a single fitting factor that we called the 'scale factor' of the greybody fit. We obtained a measure of the error of the fit from the returned 1σ errors.

We subtracted the residual of the fit from the central pixel in log space. This

residual was then used as a ripple subtraction function, subtracted from each pixel in turn in log space. The residual was rescaled as appropriate. This residual from the central pixel fit is shown in Figure 6.12. From this, we obtained values for β and T for each pixel by fitting, and this is shown in Figures 6.13 to 6.19 and Table 6.5. The error bars quoted are for the 1σ errors in the fit coefficients as returned from the fit routine. For these plots, we note that we removed the first five datapoints from each SLW observation, as at this end of the observation there is a pronounced tail that is a processing artefact.

6.4 Discussion

The fits shown in Figures 6.13 to 6.19 are for the SPIRE data only, although early on in the processing SPIRE Photometer data at 250, 350 and 500 μm, ISO data at 90, 170 and 200 μm, MIPS data at 160 μm, and SCUBA data at 450 and 850 μm were used to check flux calibration. These data points are shown on the indivdidual plots for each combined pixel, but they are not used in the fit. Pixel C1 is on source, and the remaining pixels are located anticlockwise from the 12 o'clock position. We note that from Figure 6.10 that C2 and C5 are further outside the elongated extended emission surrounding the source, and C3,C4,C6 and C7 are more within this extended position. Tables 6.3 and 6.4 show the fit result for temperature and β respectively. The central position shows the fit result for the central pixel, and the remaining values show the difference for each pixel fit position (Figure 6.10). We note that with a 48 arcsec spacing between C1 and C2–C7, and an assumed distance of 160 parsecs for L1689B that the radius of the circle bound by C2–C7 is approximately 8000 AU. Figure 6.20 shows the greybody SEDs in their respective pixel positions.

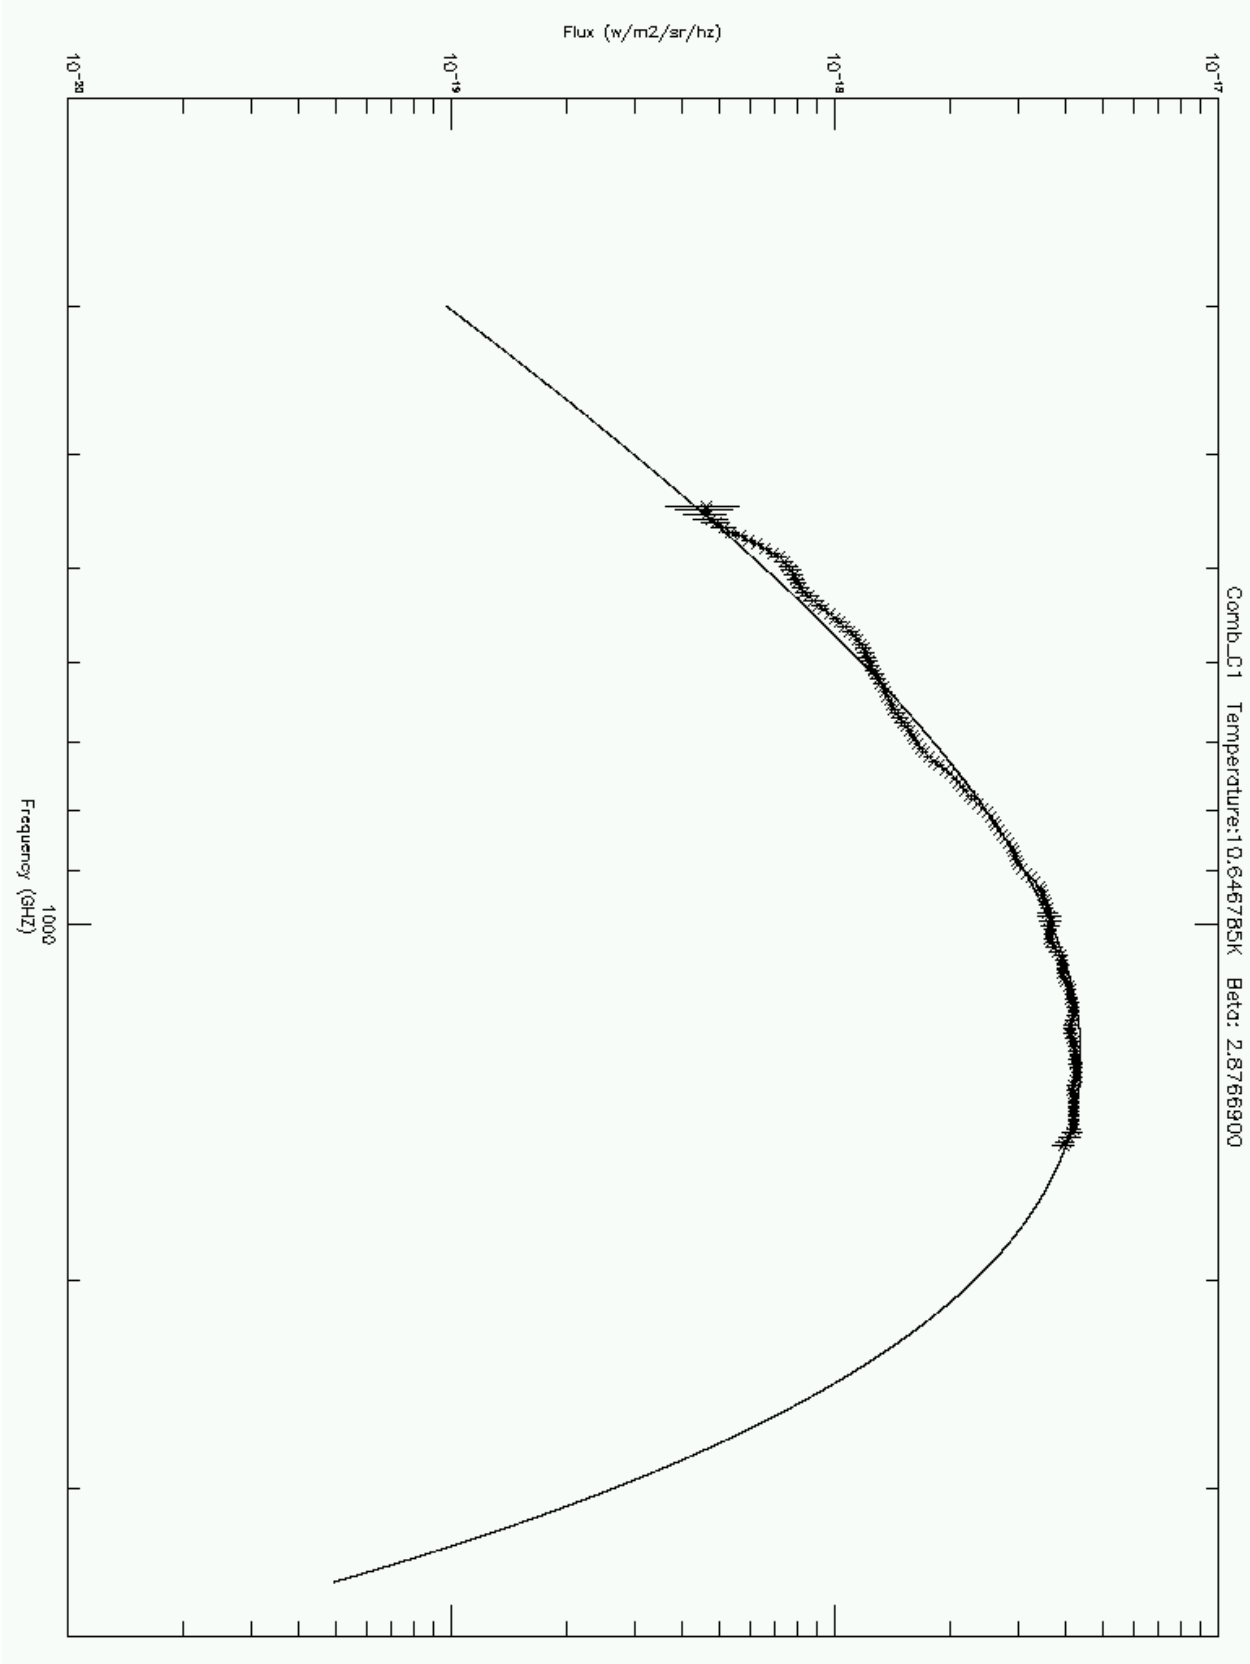

Figure 6.11: L1689B central pixel with overlaid fit.

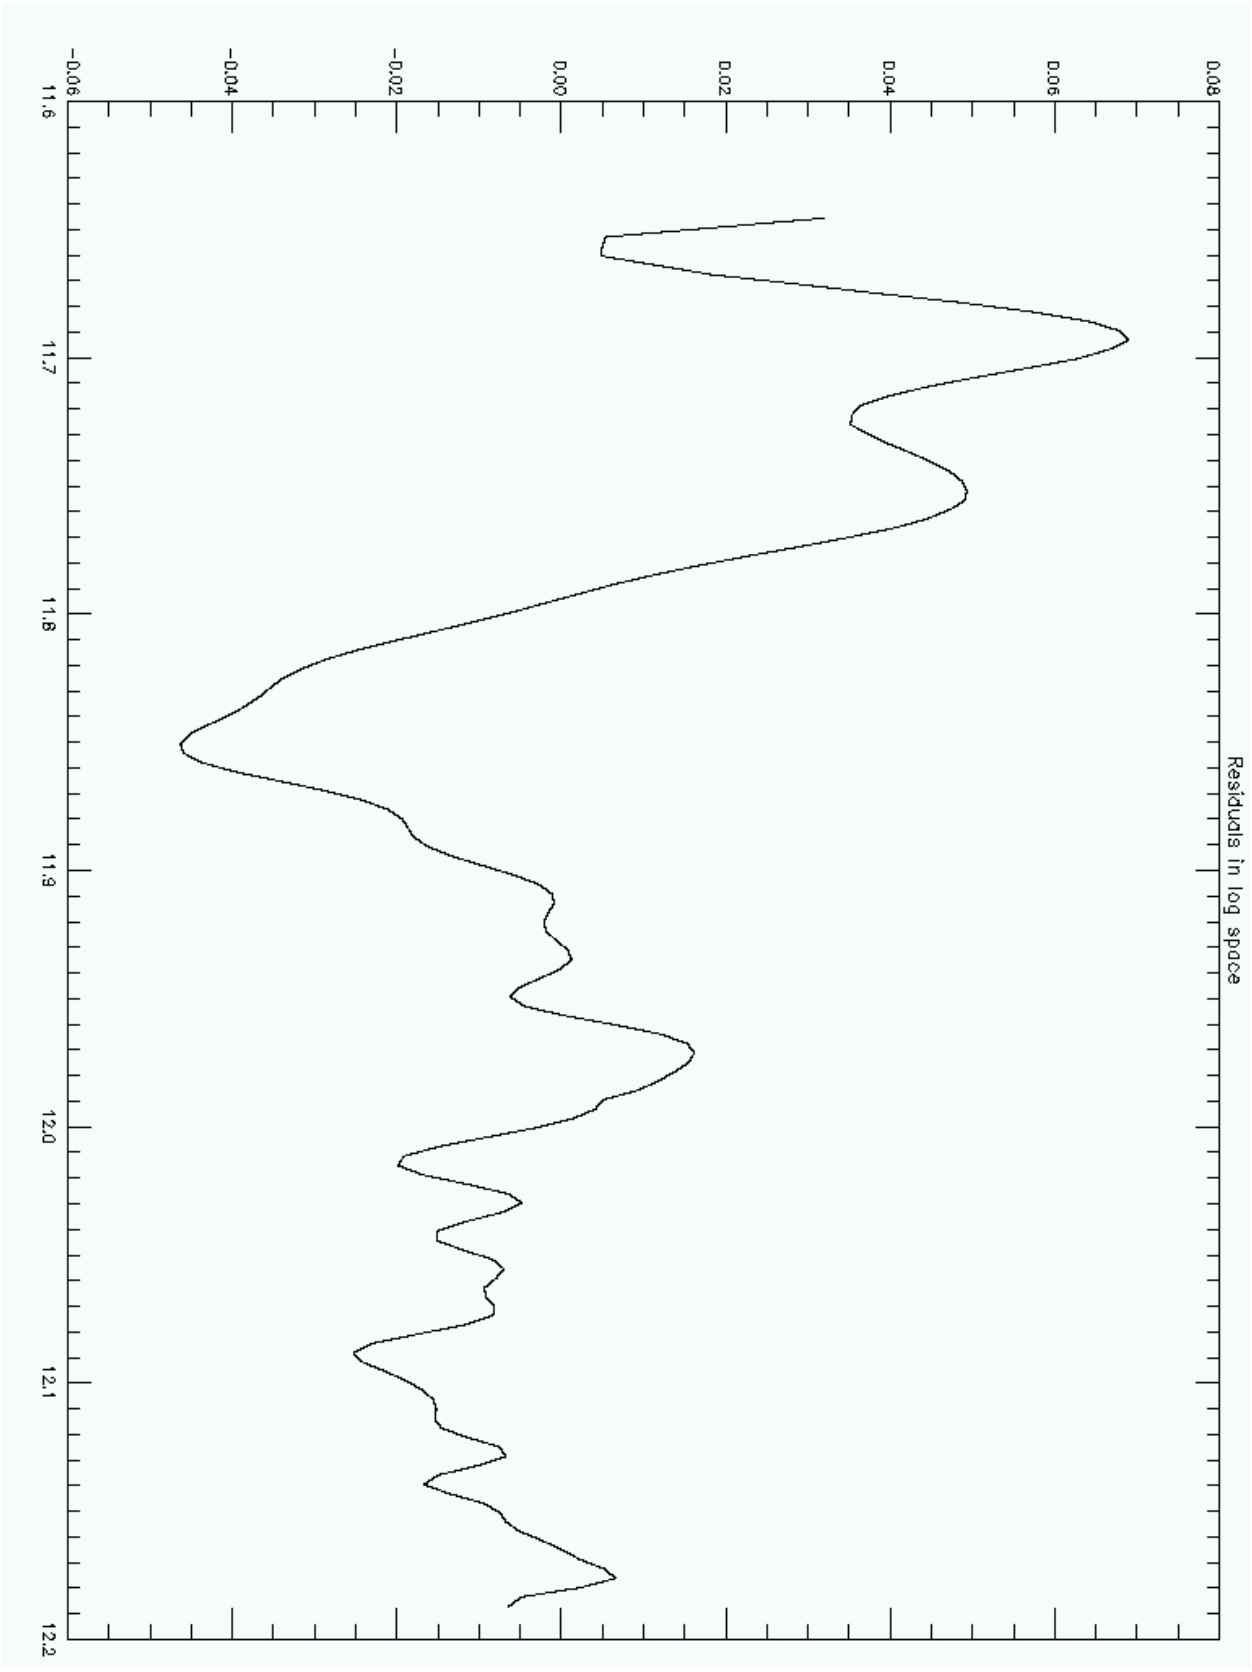

Figure 6.12: L1689B central pixel residuals from fit.

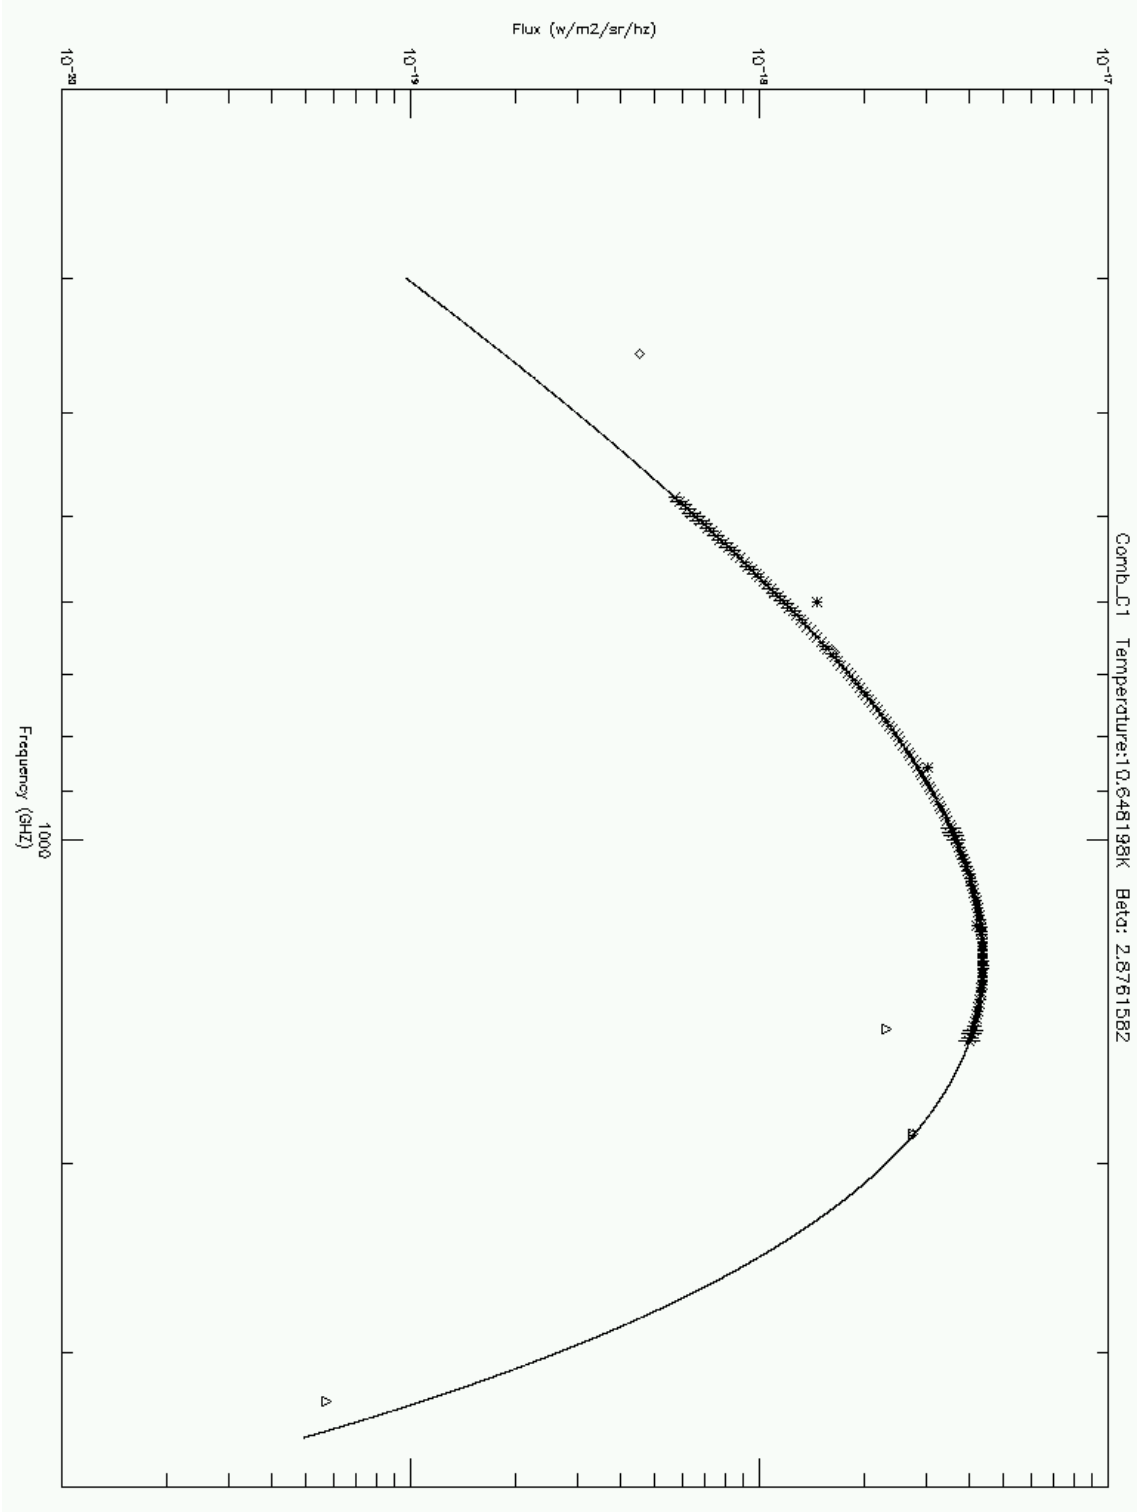

Figure 6.13: L1689B Greybody Fits - pixel C1.

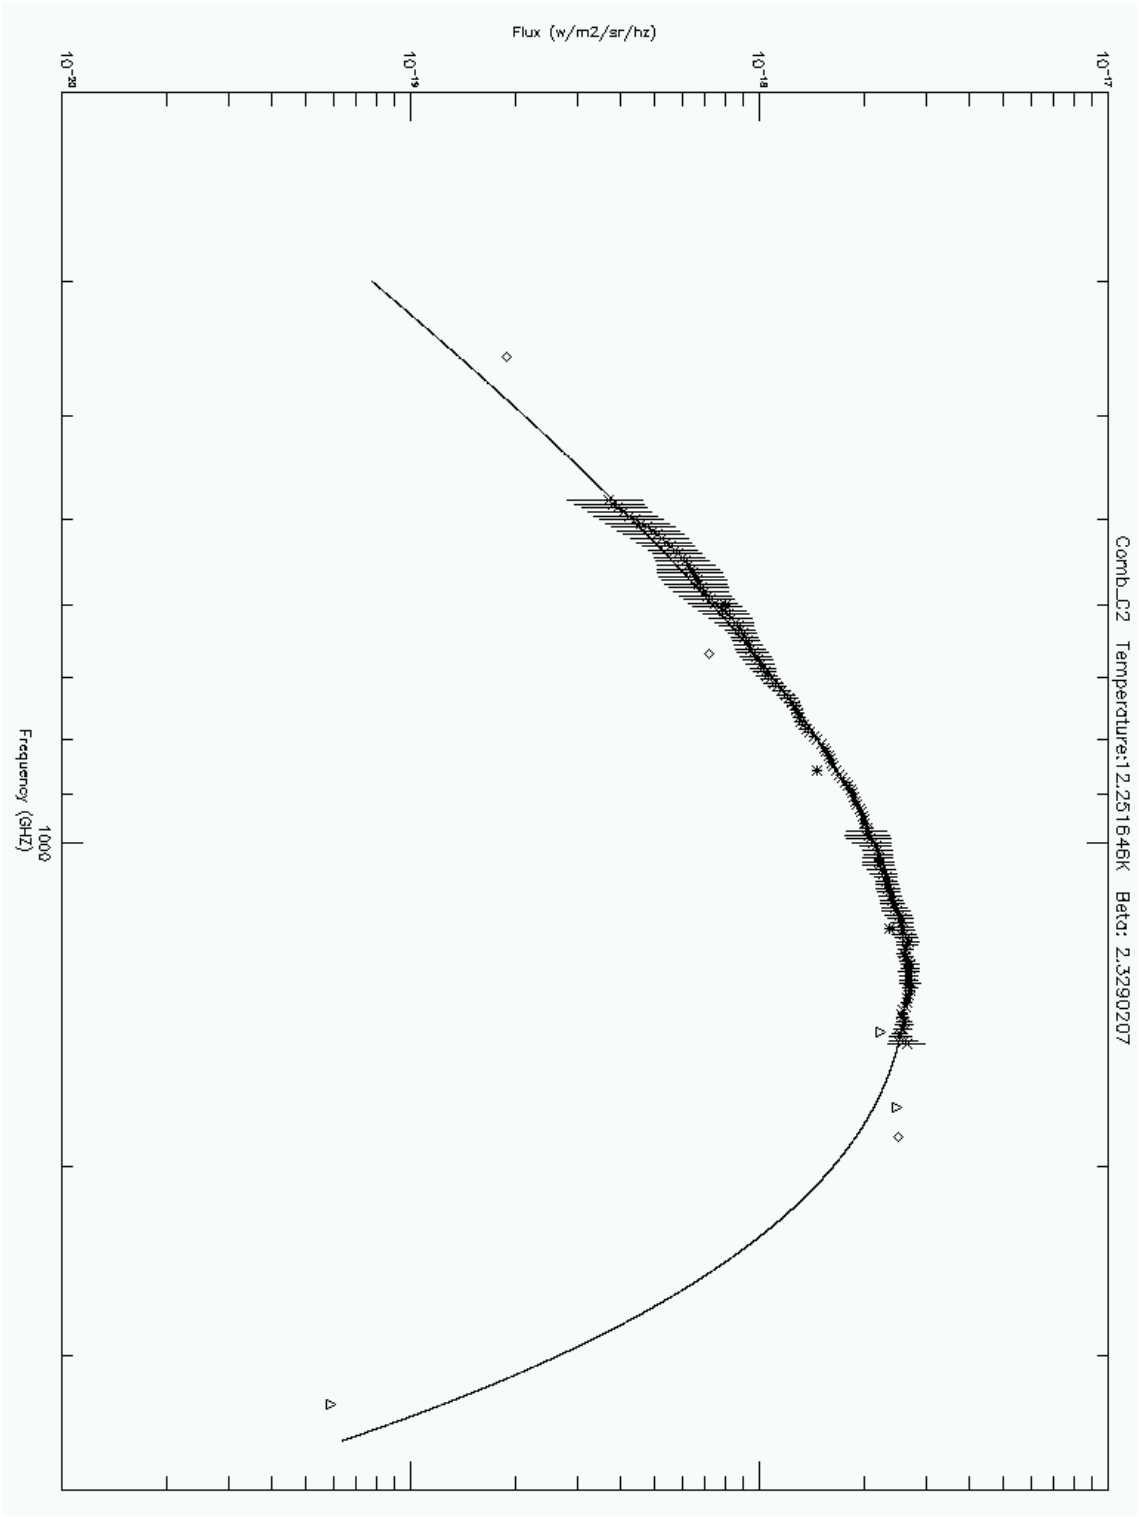

Figure 6.14: L1689B Greybody Fits - pixel C2.

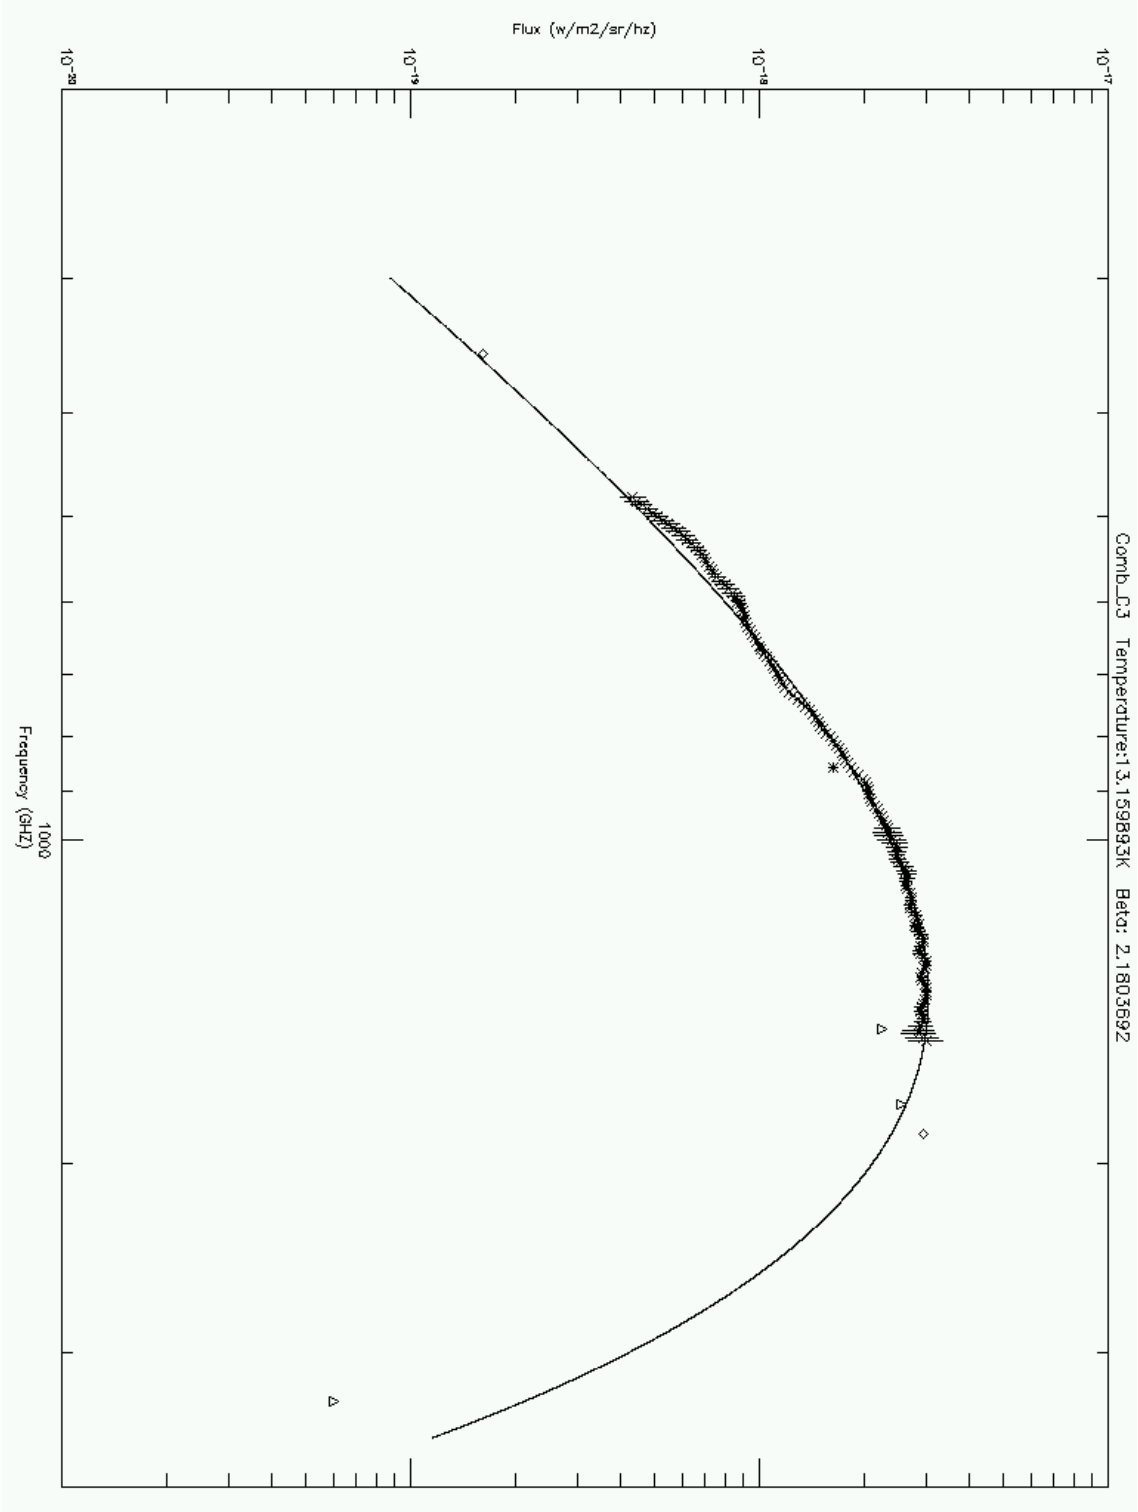

Figure 6.15: L1689B Greybody Fits - pixel C3.

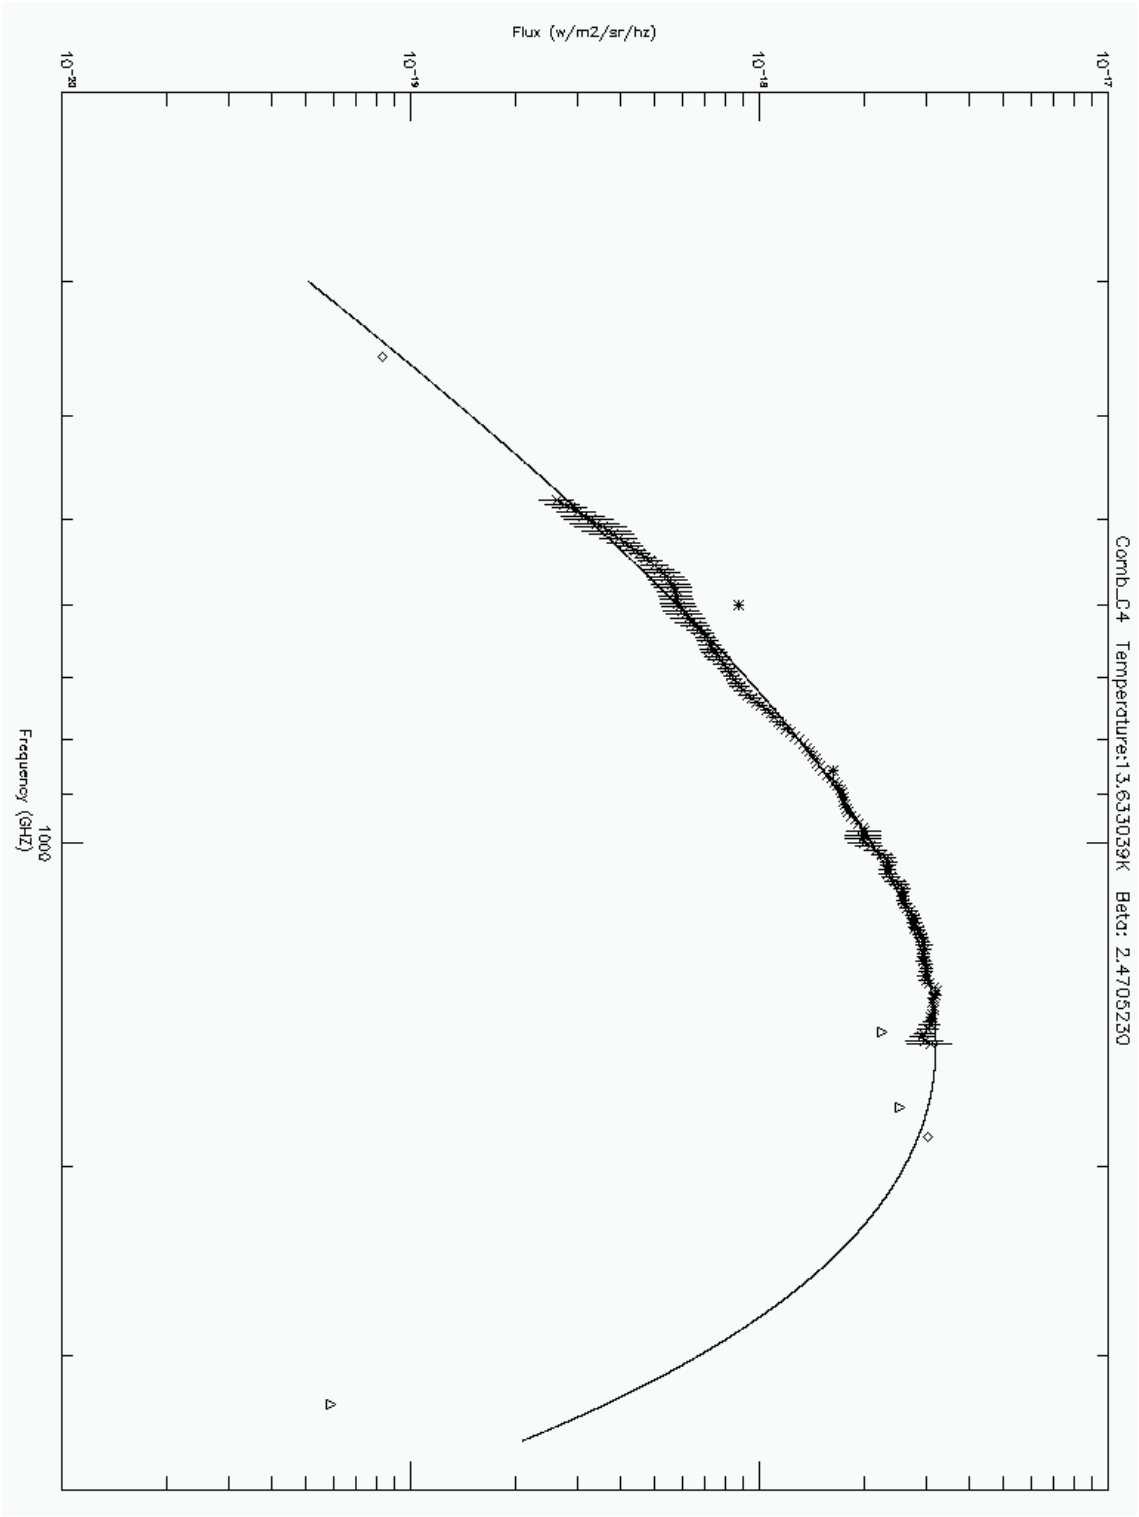

Figure 6.16: L1689B Greybody Fits - pixel C4.

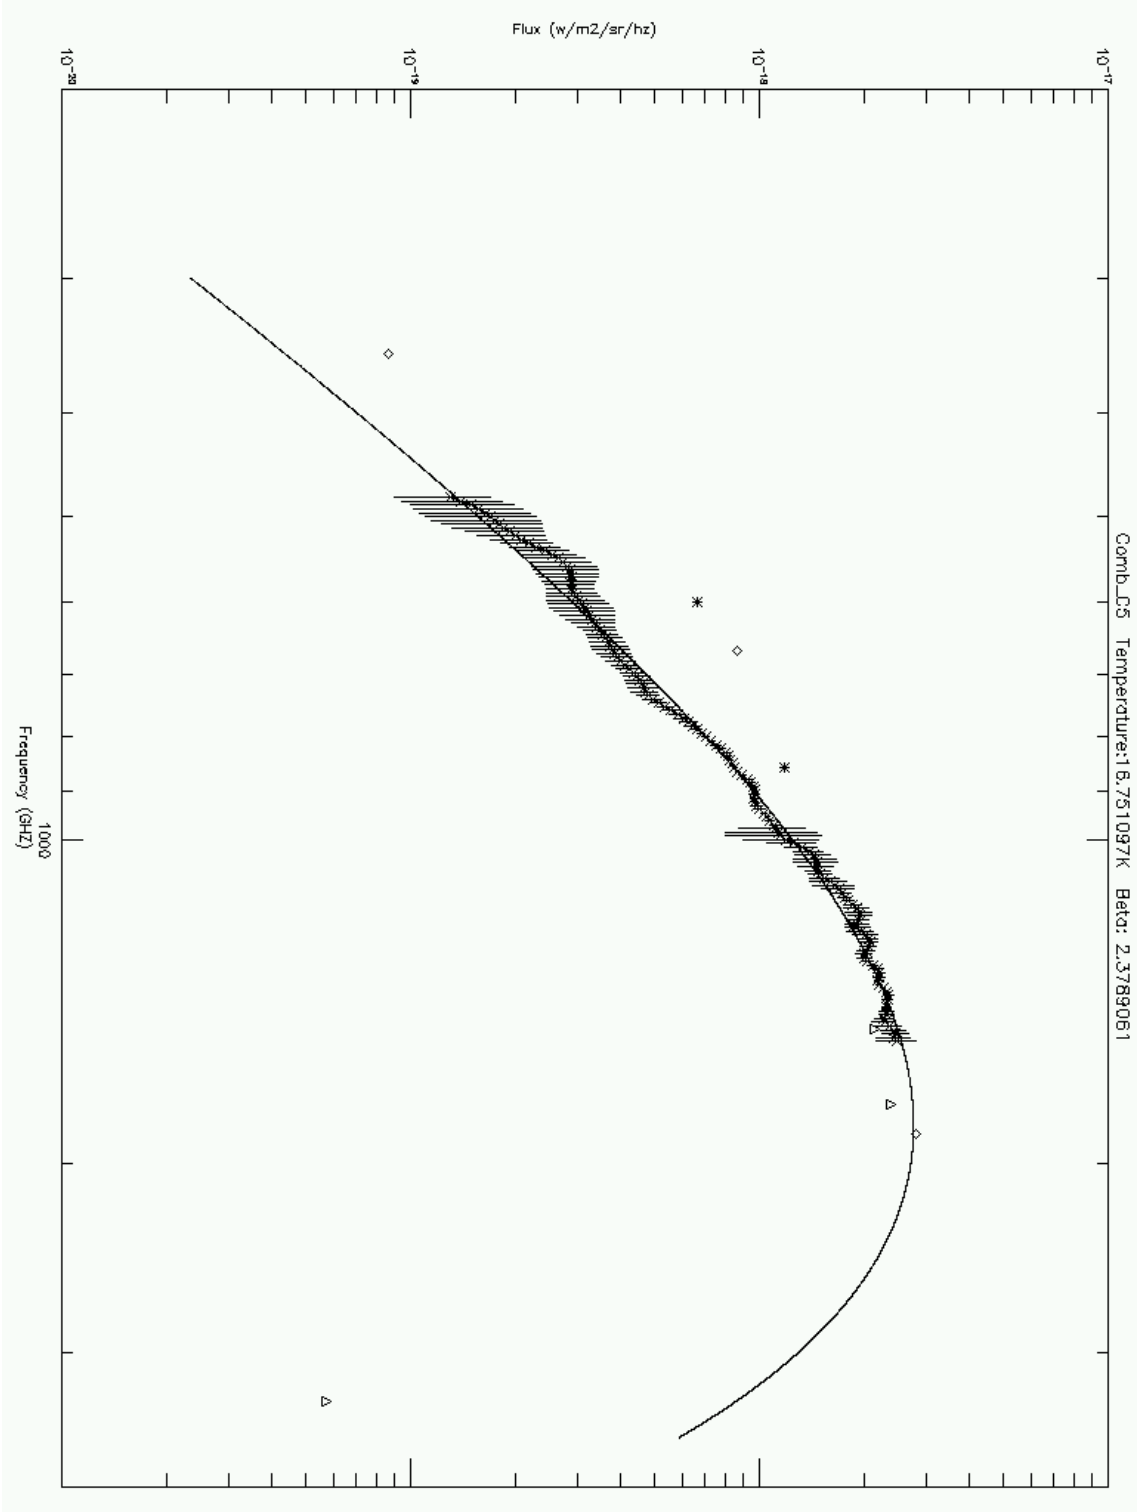

Figure 6.17: L1689B Greybody Fits - pixel C5.

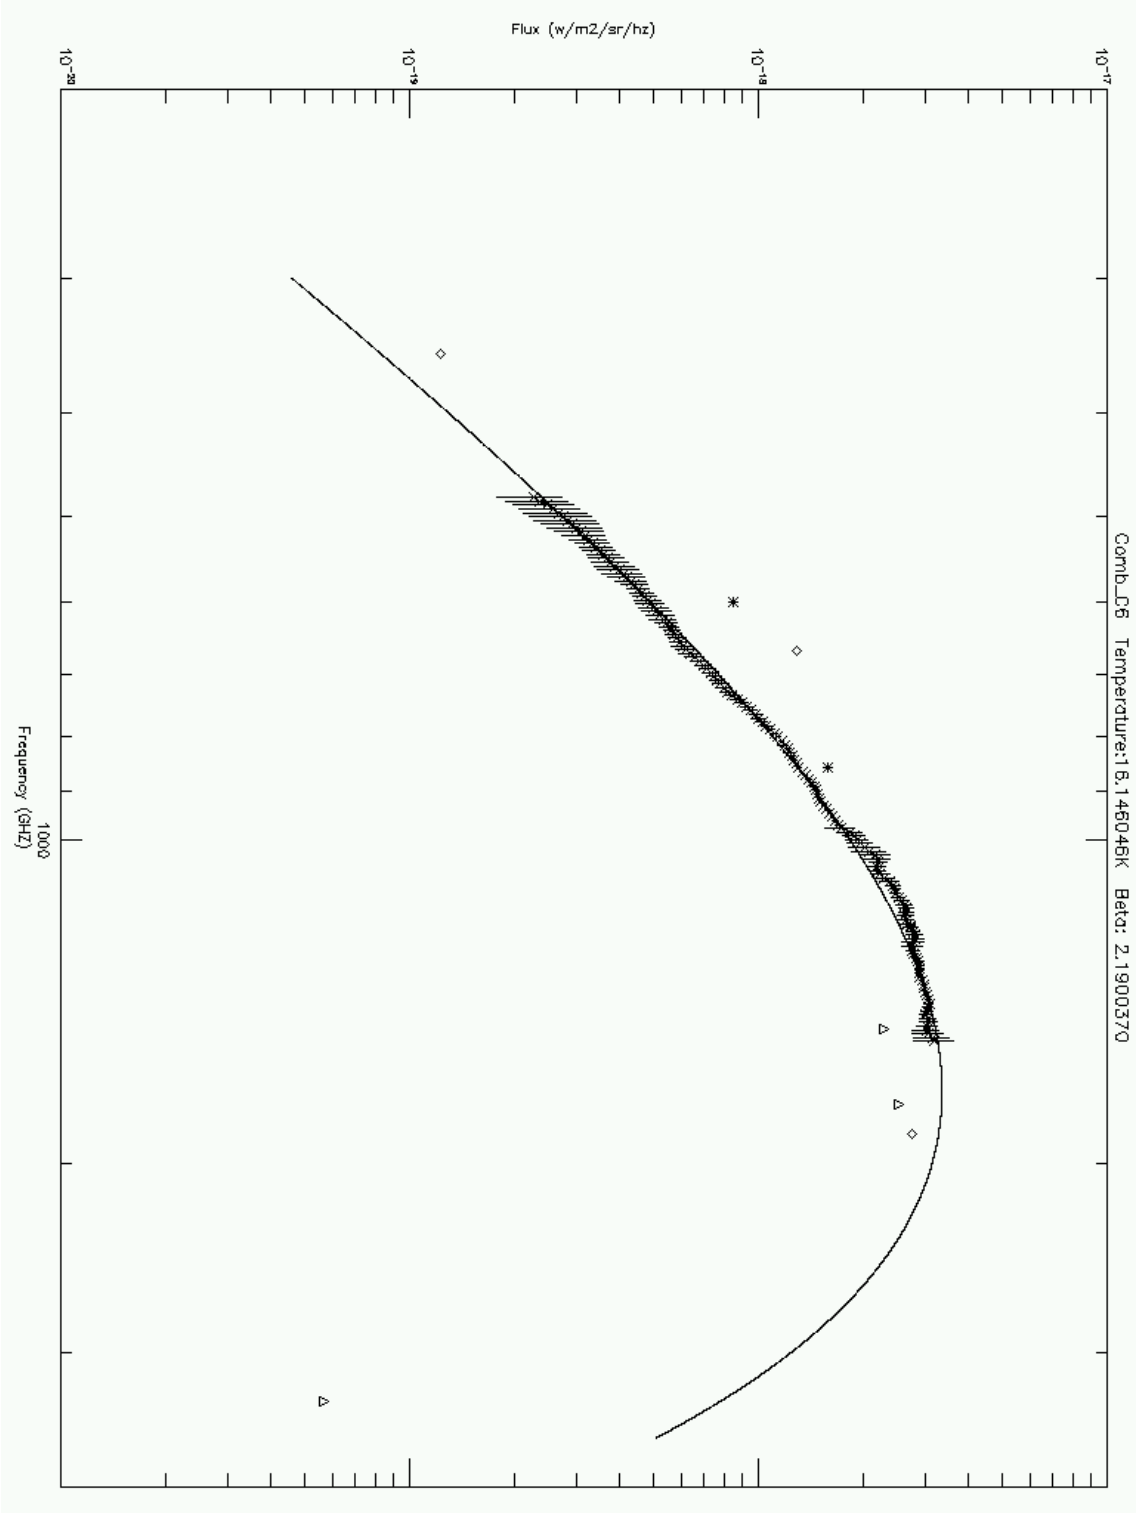

Figure 6.18: L1689B Greybody Fits - pixel C6.

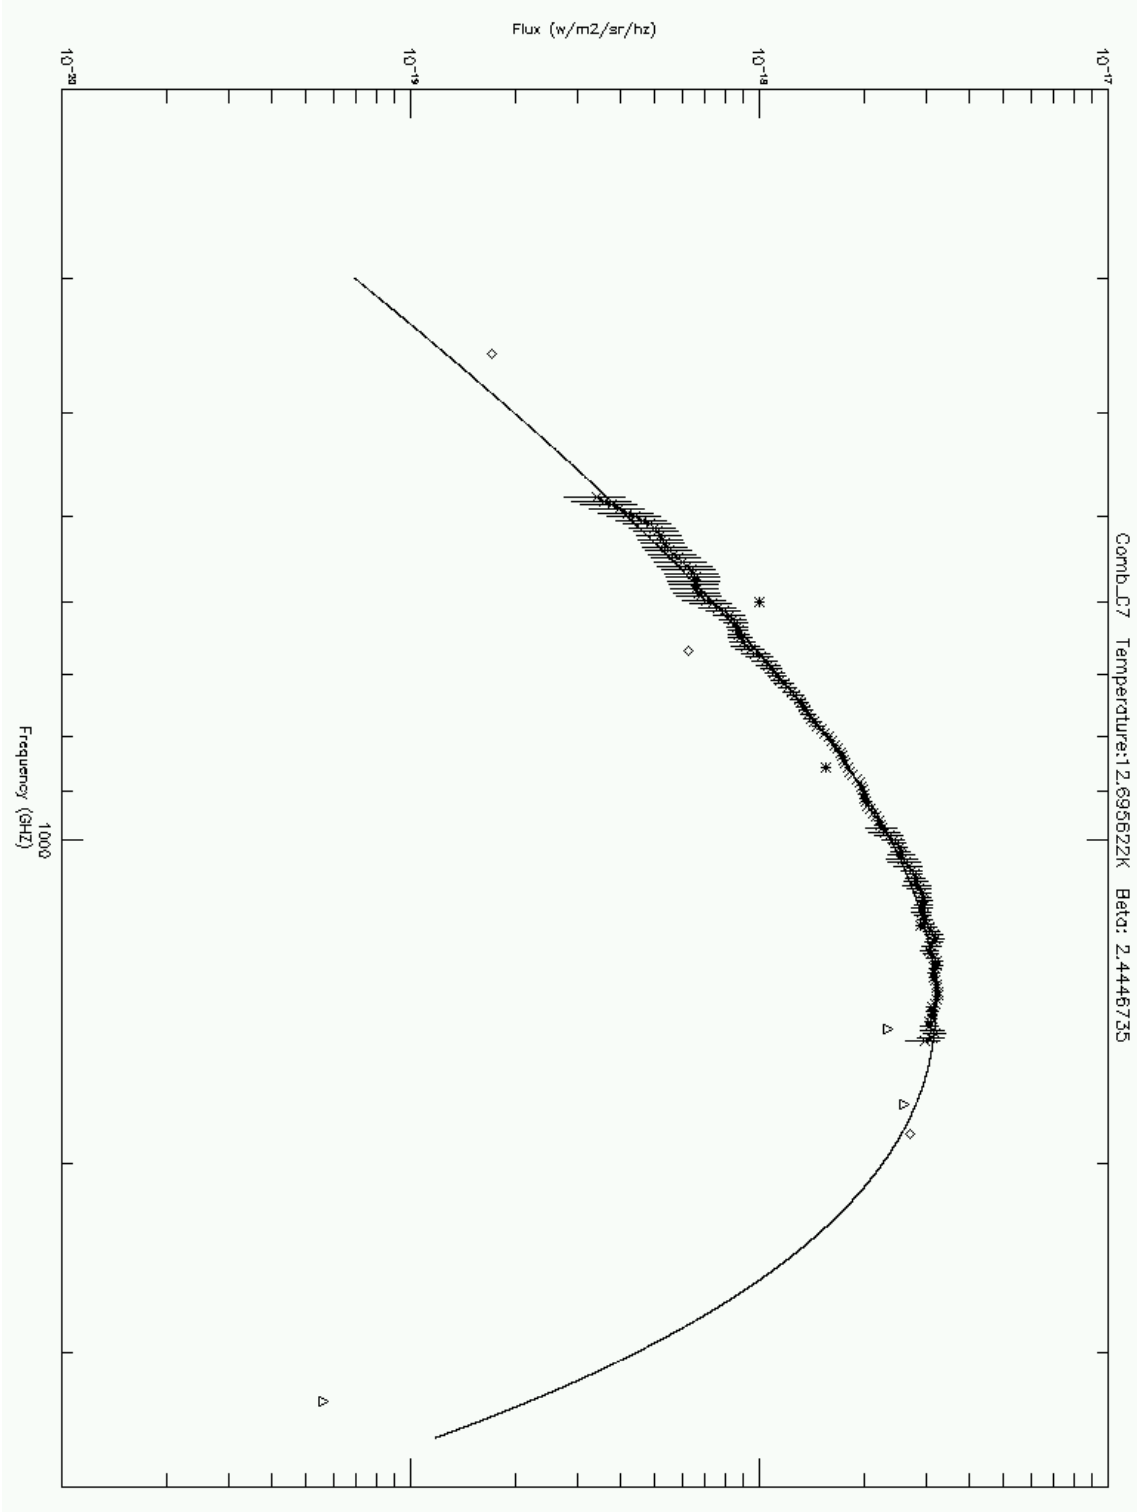

Figure 6.19: L1689B Greybody Fits - pixel C7.

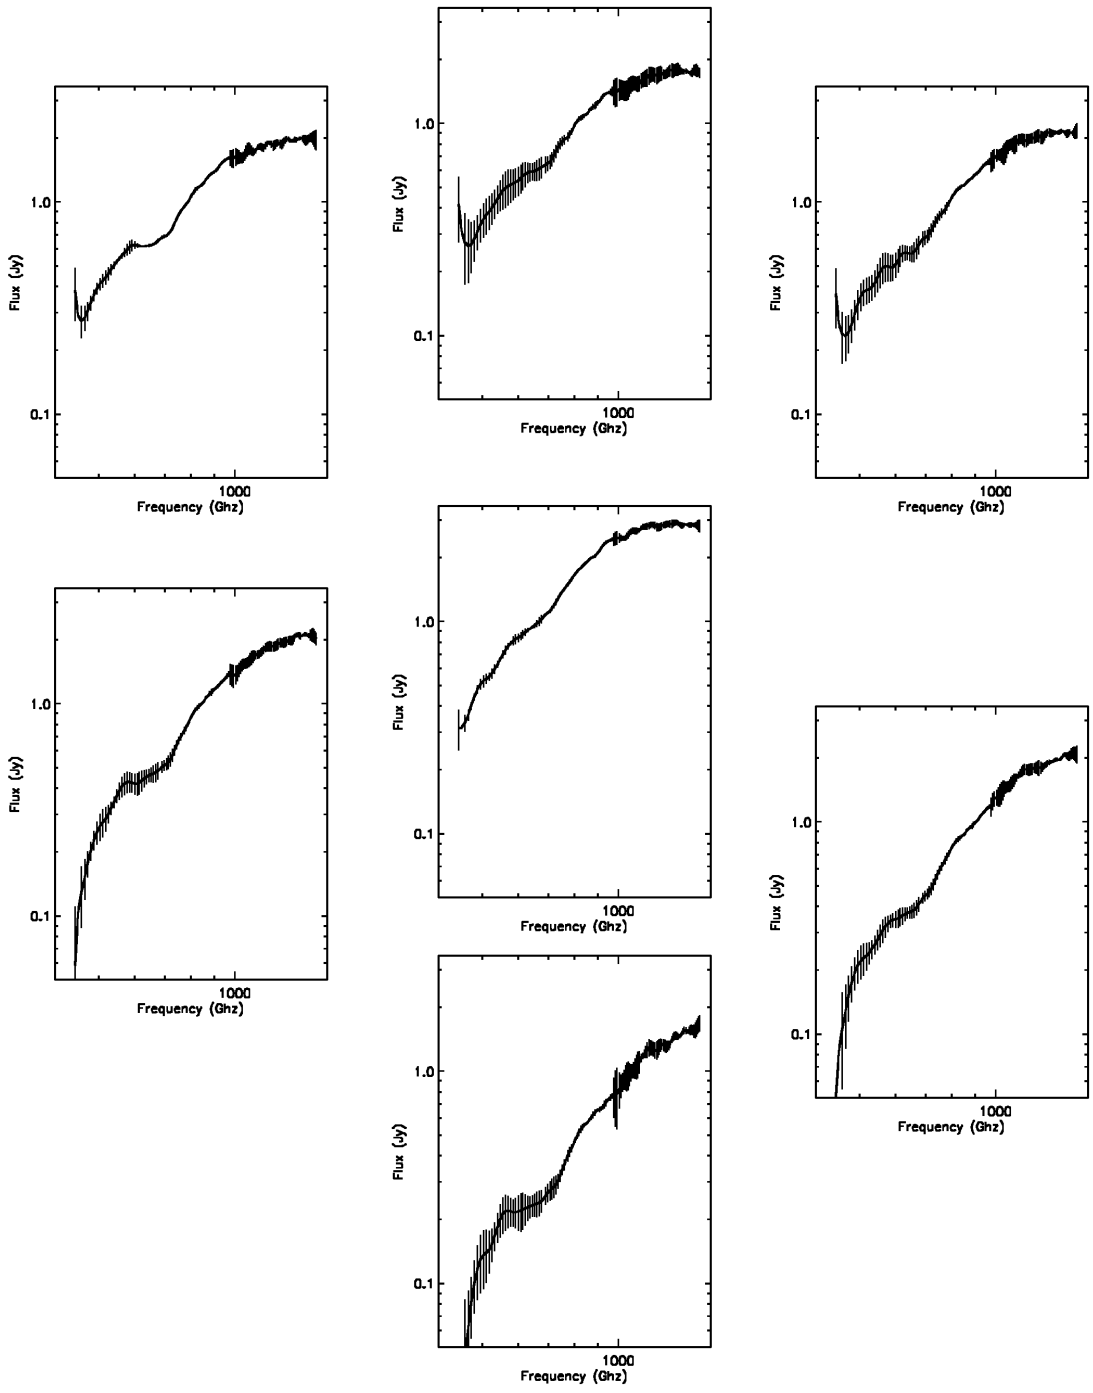

Figure 6.20: L1689B SED - in pixel position without ripple correction.

	Temperature	
	12.25	
13.16		12.70
	10.65	
13.63		16.15
	16.75	

Table 6.3: Fit results for temperature from L1689B, shown in their relative position, (c.f. Figure 6.10).

	β	
	2.33	
2.18		2.44
	2.88	
2.47		2.19
	2.38	

Table 6.4: Fit results for β from L1689B, shown in their relative position, (c.f. Figure 6.10).

Pixel Name	Temperature (K)	Error (K)	β	Error
C1	10.65	0.20	2.88	0.07
C2	12.25	0.12	2.33	0.04
C3	13.16	0.32	2.18	0.07
C4	13.63	0.27	2.47	0.07
C5	16.75	0.78	2.38	0.12
C6	16.15	0.45	2.19	0.08
C7	12.70	0.20	2.44	0.06

Table 6.5: Fit results for each pixel.

6.4.1 L1689B Temperature profile

Table 6.3 shows the greybody temperature distribution over the pixel array. This table shows that the central pixel is $2-3K$ cooler than the outlying pixels. This suggests that not only is the core shielded from the ISRF, but that there is no internal heating source detectable at $500\mu m$. From this we conclude that L1689B is a starless core, in agreement with previous work.

The differential temperature profile of L1689B was mapped by Ward-Thompson et al. (2002) using ISO data. In this study the colour temperature was mapped for two fluxes ($200\mu m$ and 170 μm). They discovered a temperature profile that increased from the centre in one direction, but decreased from the centre in the orthogonal direction. This was attributed to a lack of signal-to-noise (SN) in the orthogonal direction, or a lack of complete mapping of the entire core. Where the SN was high, the temperature profile showed a cooler inner region. This is consistent with our results as shown.

The pixel array from Figure 6.9 has a hexagonal shape with an obvious 3-axis symmetry. We have taken this triaxial symmetry, and plotted the temperature profile over the three axes. We plot the temperature profile in Figure 6.21. This figure shows that the central pixel is colder than the outer ring of pixels. All axes show an asymmetry, which is weaker along axis C and within the error bars on the points. In this figure, we have reversed the x-axis of the C plot so it plots $C7-C1-C4$. We see that there is some evidence for a directional source in the asymmetry of the total profile. Its position on a projected 2D plane would lie approximately on a line pointing from the central pixel and in a direction towards C6 to the south-west.

The existence of an external heating source in this region was discussed by Nutter et al. (2006), as an attempt to explain the discrepancy in star formation rates in the L1688 and L1689 regions in Ophiuchus. L1688 shows a more active star formation history than L1689. Elmegreen & Lada (1977) outlined a process where the formation of OB stars triggers a positive feedback cycle of local star formation. Nearby young massive stars interact with the molecular cloud and the result is that star formation propagates. Nutter et al. (2006) concluded that the differential star formation rate is because of the different distance of the two clouds to the Upper Scorpius OB association, and in particular to σ Sco (Loren, 1989). This is a very young, very massive hierarchial multiple OB system that lies at a projected distance of around 6 pc. We have found our temperature profile for L1689B suggests the source of external heating lies to the south-west.

We can therefore hypothesise that the measured temperature profile of the core is suggestive of the external heating effects of σ Sco or other members of the Upper Scorpius OB association. Shielding from the L1689 region may be the reason why the direction of the measured profile does not exactly match the direction of the association. Nutter et al. (2006) considered the extensions to the L1688 and L1689 regions, to the north-east of each cloud, as evidence of interaction between the association and the cloud. L1689B lies on the furthest edge of the L1689 cloud from σ Sco, and given the shape of the cloud it is feasible to expect a degree of shielding from the star. This shielding would affect the direction of the temperature profile.

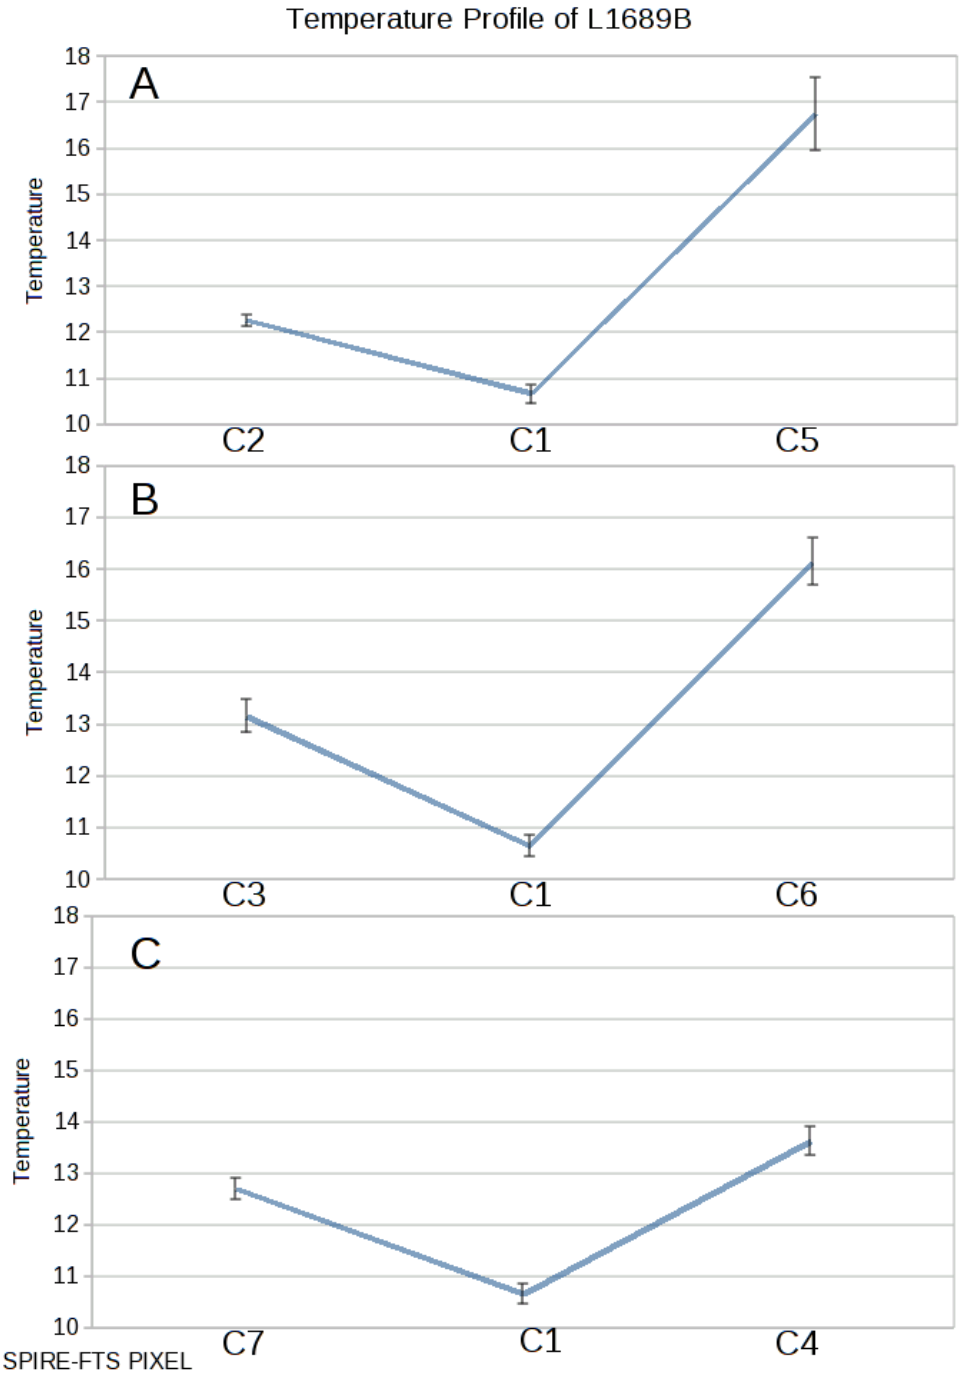

Figure 6.21: Triaxial temperature profile for L1689B.

6.4.2 L1689B Spectral Index profile

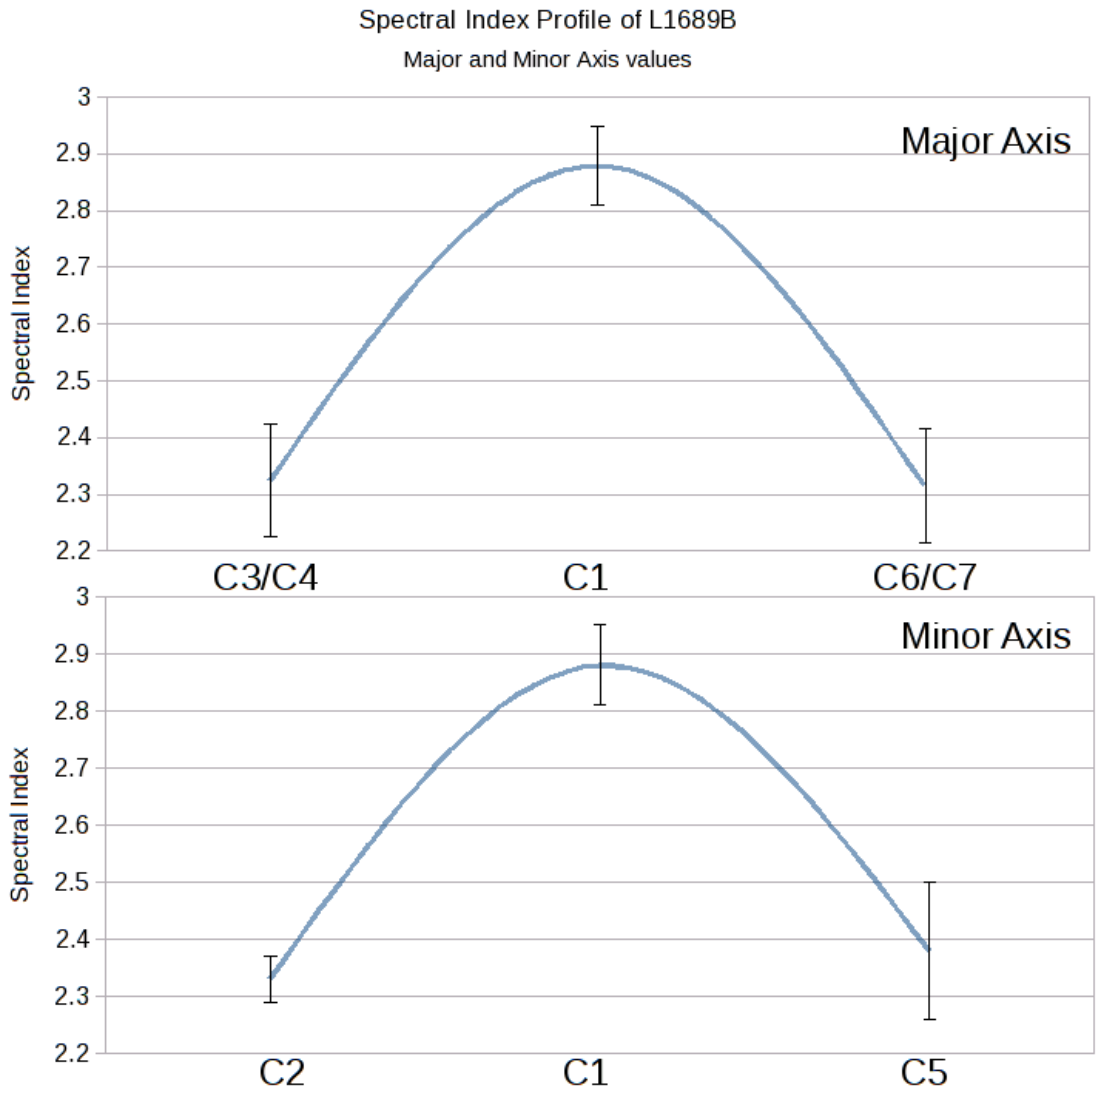

Figure 6.22: Spectral index profile for L1689B.

We can generate profiles of the spectral index β, where we are looking at the internal properties of the core. We have seen in Figure 6.10 that L1689B shows an elliptical shape, and so we generate a biaxial spectral index profile. We generate a minor axis profile of $C2 - C1 - C5$, and a major axis profile by taking the mean of the pixels either side of $C1$. This major axis profile is therefore $mean(C3, C4) -$

$C1 - mean(C6, C7)$. This is shown in Figure 6.22.

The profile is largely symmetric across the core. The profile shows a central, maximum value of β that decays smoothly with distance. This would suggest that there is some intrinsic property of the core we are mapping. If we consider the pixel array as a perfect hexagon centred on the core, the difference from the central pixel to our major and minor axes endpoints is around 10%. Taking this into account, there is excellent agreement in the profile of the spectral index over the major and minor axes. $\beta_{C1} \sim 2.9$, and over the whole core β is greater than 2.2. This is larger than the typical value of $\beta = 2$ considered for these cores (e.g. Ward-Thompson et al. (2002)).

From this profile, we need to understand:

1. What is the explanation for the high central value of β in this core?
2. Is the large value of β throughout the core indicative of an intrinsic property of the core?
3. Why does β decrease in all directions away from the centre?

Stepnik et al. (2003) showed that dust grain aggregation in molecular clouds could lead to an increase in the emissivity in the sub-millimetre by a factor of 3. They also argued that such a change would explain the colder temperatures found in such a region. They concluded that dust evolution in denser regions was responsible for such enhancement, and that coagulation of dust grains could be an important part of the evolutionary process.

Paradis, Bernard, & Mény (2009) further studied variations in the dust emissivity in molecular clouds in our Galaxy. The spectral index β was found to signifi-

cantly increase in cold molecular clouds in the far-infrared (defined as $\lambda < 600\mu m$) to values greater than 2. In regions where molecular phase dust is colder than dust in the atomic phase, there is a threefold increase in the emissivity in this region. This increase is attributed to the coagulation of large dust grains into fractal aggregations, from dust emission models by Meny et al. (2007). Paradis et al. (2009) concluded that this is a tracer of quiescence in a molecular region, as star formation activity in the cloud would prevent the formation of these grains.

We conclude that the spectral index profile we see is the result of a starless core, with no internal heating source. The increase in β towards the centre is the result of an increase in fractal grain density as dust grains coagulate. We believe this is in agreement with the density profile observed by Ward-Thompson et al. (1994). The increase in grain density and spectral index profile would also be in agreement with the CO depletion data as presented by Redman et al. (2002). The presence of fractal grains would also provide further evidence as to the starless nature of the core.

6.4.3 Conclusions

The temperature profile of this core leads us to conclude that L1689B is a starless core. We see no evidence of an internal source of heating. The profile suggests that the proximity of the core to the Upper Scorpius OB association is responsible for external heating. This is evident in the asymmetry of the temperature profile. The spectral index profile shows a clear symmetry between major and minor axes of the core. We conclude that an increase in fractal grain density towards the centre of the core is the explanation for the spectral index we see, which we believe is in agreement with the CO depletion, density profile and starless nature of the core.

6.5 Summary

We have selected the L1689B on-source SPIRE observations for analysis, as they contain the highest S/N of the observations provided. The core was processed by correcting for instrumental effects in the spectral domain, and a number of coincident pixels from both SLW and SSW were used to map the core over an extended region. Each pixel observation was fitted by a greybody, and from this we were able to map out the profile of temperature T and spectral index β over this region.

This core has been well studied in the literature, and the density profile and morphology have been discussed. There appears to be a two-tier density structure, and an oblate shape has been proposed for this core. There is also strong evidence for a large degree of CO depletion within the core. This is believed to be due to freezing out of CO molecules onto dust grains. Earlier studies have concluded that L1689B is rotating, and is in the process of forming a central protostar.

The SPIRE FTS observations of L1689B have revealed an asymmetric temperature profile, and a symmetric spectral index profile, over the mapped region of the core. The temperature profile has revealed the centre of the core is $\sim 3K$ colder than the outlying regions, and from this we conclude there is no central heating source seen within the core. This categorises L1689B as a starless core. The asymmetry in the temperature profile is suggestive of an external heating source to the core, and we conclude that this is due to the Upper Scorpius OB association, and most likely σ Sco. This lies within 6 pc of the core, and interaction between the association and the local cloud complex has already been suggested to explain the morphology of the region in earlier studies.

The spectral index profile reveals that the value of β is greater than 2.2 over the whole of the mapped region, and increases to a central maximum value of $\beta_{cent} \approx 2.9$. This is significantly larger than $\beta = 2$ commonly used in modelling prestellar cores. We note that we also see a dependence on the radius. We conclude this is due to an decrease in fractal grain density with radius, which would explain the spectral index profile we have found. This explanation would also agree with the CO depletion and density profile of this core, as well as our designation of this core as starless.

Chapter 7

Summary and Future Work

7.1 Summary

In Chapter 1, we outlined the process of low mass star formation through the characteristics of prestellar cores. From the definition of a prestellar core (Ward-Thompson et al., 1994) we have discussed core modelling through consideration of the density profile. Properties such as temperature, size and morphology have been examined. The mass and density of the prestellar cores are determined from observations of the dust emission. Absorption studies have also allowed study of density profiles in some cases. The lifetime of a prestellar core has been deduced from statistical studies and various lifetimes have been considered in the literature. We have also discussed the dynamical properties of cores, such as rotation and turbulence, as well as the effects of ambient magnetic fields on the evolution of the core. We have also considered the chemical composition of the core. From this we have seen how low mass star formation is affected by ambient conditions, and how these conditions are reflected in the composition and evolution of the core.

In Chapter 2, we examined the basic theory behind the FTS. From a discussion of the idealised instrument, we considered the advantages of the FTS over other instrument designs. A mathematical discussion of the instrument was presented. From the idealised instrument, we then compared how instrumental restrictions constrain the FTS. From these, we examined the observational limits such as spectral resolution and bandwidth, as well as the effects of a finite field of view, on these defining characteristics. In examining the effects of these limits, we noted that apodisation of the interferogram to remove sidelobe effects from the ILS is not a simple process. The choice of apodisation function used is important to balance the observational needs for each case. Finally, we considered how another feature of a real FTS may cause observational issues. A finite scan size for the movable mirror need not necessarily sample the point of zero path difference. If this point is missed, then processing the resultant interferogram needs to consider the phase error introduced as a result. We considered two methods to correct for this phase error.

In Chapter 3, we showed that it is possible to construct a portable, non-cryogenic infra-red interferometer using off-the-shelf components. A multiple pixel detector enabled us to examine some of the developmental issues that faced the design and construction of larger instruments such as SPIRE. Although some construction issues are still ongoing, the instrument has shown that it is possible to construct a low-cost interferometer that can be used to produce useful science. The design and construction of LoCI highlights many issues that are faced in astronomical detectors, such as housing design for the instrument and calibration and testing of any interpretational software.

In chapter 4 we described SLIDE, which was developed as an IDL tool to quickly and accurately identify line parameters in SPIRE SEDs, although it has

progressed far beyond the original line identification tool. As such, it is used in conjunction with the processing pipeline as opposed to being any replacement for it. A number of useful continuum fit functions were identified and tested, although SLIDE primarily uses continuum-fitting as a means to improve line-fitting rather than to perform detailed continuum analysis. In examination of SEDs produced by SPIRE and through continuum fitting, fringing proved to be problematic. The accurate removal of fringing to leave a smooth baseline is important to accurately identify line parameters. Issues with fringing of the SED as a result of artefacts of processing within the pipeline were identified and analysed. Tests were performed to identify the effect such fringing would have on line identification. 'Zones of avoidance' were found where the identification of weak lines would be increasingly sensitive to continuum subtraction. SLIDE was tested using noise-free and noise-added input SED's with various line profiles, and through this testing the continuum fit and line profile fit routines were developed to be as robust as possible. Issues with SCG line fitting below a specific Gaussian width component were identified, and a new 'Composite' fit routine was developed to surmount the problem. SLIDE was originally developed to perform simplistic line identification on text-file SEDs, but now SLIDE has a variety of continuum- and line-fitting routines to process SPIRE data and can accept text- or FITS-file input. It outputs graphical and textual information of the input SED. SLIDE was eventually merged into the HIPE tool.

In Chapter 5, we fitted the SPIRE photometer Relative Spectral Response Function (RSRF) into a set of polynomials. Using these on a range of astronomical SPIRE-FTS observations, we identified a realistic range of line contamination percentages that could be seen in photometer observations. Using the SPIRE photometer bands with other wavelengths, we produced a range of simulated grey-

bodies defined by their temperature and their spectral index. From these we generated a range of simulated observations by adding noise and line contamination and compared the recovered temperatures and spectral index to the original values. We considered line contamination added to either a single band, or a number of bands. In all cases, the amount of offset in either parameter was under 3σ and cannot be considered significant. We then identified two sources from our original list as extreme cases. For the two astronomical sources we considered, the effect of line contamination was not detectable. Adding line contamination to a single band, or a number of bands, did not result in a shift in T or β greater than 3 σ. This cannot be considered as a detectable change. Hence, we concluded that line contamination is probably not a serious problem for real astronomical sources.

In Chapter 6, we selected the L1689B on-source SPIRE-FTS observations for analysis. The source was processed by correcting for instrumental effects in the spectral domain, and a number of coincident pixels from both SLW and SSW were used to map the core over an extended region. Each pixel observation was fitted by a greybody, and from this we were able to map out the profile of temperature T and spectral index β over this region. The SPIRE FTS observations of L1689B revealed an asymmetric temperature profile, and a symmetric spectral index profile, over the mapped region of the core. The temperature profile showed that the centre of the core is $\sim 3K$ colder than the outlying regions, and from this we concluded there is no central heating source seen within the core. This confirms L1689B as a starless core. The asymmetry in the temperature profile is suggestive of an external heating source to the core, and we conclude that this is due to the Upper Scorpius OB association, and most likely σ Sco.

The spectral index profile reveals that the value of β is greater than 2.2 over the whole of the mapped region, and increases to a central maximum value of

$\beta_{cent} \approx 2.9$. This is significantly larger than the value of $\beta = 2$ commonly used in modelling prestellar cores. We note that we also see a dependence on the radius. We conclude this is due to an increase in fractal grain density growth towards the centre, which would explain the spectral index profile we have found. This explanation would also agree with the CO depletion and density profile of this core. This is consistent with dust grain growth models in prestellar cores.

7.2 Future work

For LoCI, the instrument will be used for undergraduate research projects. The instrument will give students the chance to use a small-scale, portable FTS. There is also the opportunity for the instrument to be upgraded, as the components used for the instrument were purchased in 2006. Future work on LoCI would involve updating the detector array.

For SLIDE, the future is that as well as integration within HIPE, it will remain as an IDL tool for SED analysis. SLIDE was designed as an interactive tool to allow the user the ability to try a variety of methods to extract line information. As data processing for SPIRE matures, it is hoped that SLIDE maintainence and development can continue.

The line contamination simulations we have performed were with a common percentage, as allowing for different percentages in each band would have increased simulation time significantly. It would be interesting to consider the effects of mixed percentages, which would accurately simulate the decrease in line contamination near the peak of the greybody as the continuum increases. This would require careful consideration of the relative percentages, examining real sources and creating

a set of useful simulated scenarios to consider.

Our study of L1689B confirmed this particular core as starless. Nutter et al. (2006) concluded that differential star formation rates between L1688 and L1689 can be explained by the proximity of the Upper Scorpius OB association. Using the SPIRE FTS to map further cores within this area would be useful to see whether temperature profiles for each core would yield a common direction that would also coincide with the Sco OB association. This would confirm our hypothesis.

Appendix A

Derivations and Tables

Derivations in Section A.1 to Section A.4.3 are collated from Whitworth (2006a) and Whitworth (2006b), course notes for undergraduate modules for which the author is eternally grateful.

A.1 Radiative transfer and the mass of a core

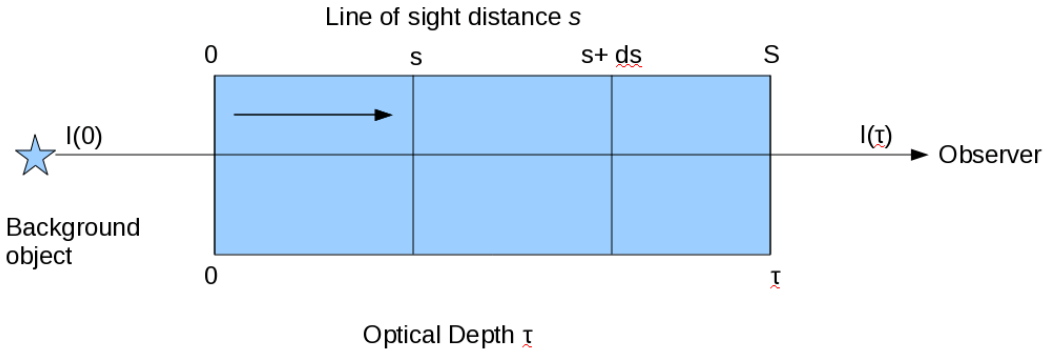

Figure A.1: Radiative transfer through an active medium.

We shall consider the path of radiation through an optically active medium, as in Figure A.1. Radiation of frequency ν from the background $I(0)$ enters the medium of path length S and is observed as $I(\tau)$ through the medium. The radiative transfer equation governing this system is given as

$$\frac{dI_\nu(s)}{ds} = -\kappa_\nu I_\nu(s) + j_\nu, \tag{A.1}$$

κ_ν is the inverse free path of the photon through the medium, and is an absorption coefficient. We can define the optical depth τ in terms of κ_ν.

$$\tau_\nu = \int_0^S \kappa_\nu ds. \tag{A.2}$$

If τ is less then unity, the medium is considered optically thin. When it is much greater then unity, the medium is optically thick and radiation from the medium is observed from a surface one mean free path thick. In essence, the optical depth is a measure of distance in the medium with units of the mean-free-path and is dimensionless. j_ν is the emission coefficient and is a measure of the spontaneous emission seen in the medium.

Dividing Equation A.1 by κ_ν, we note that

$$\frac{1}{\kappa_\nu}\frac{dI_\nu(s)}{ds} = -I_\nu(s) + \frac{\kappa_\nu}{j_\nu}. \tag{A.3}$$

Integrating Equation A.2 gives

$$d\tau_\nu = \kappa_\nu ds. \tag{A.4}$$

Substituting this into the previous equation, we find that

$$\frac{dI_\nu(s)}{d\tau_\nu} + I_\nu(s) = S_\nu = \frac{\kappa_\nu}{j_\nu}. \tag{A.5}$$

where we now define S_ν as the Source function. We also note that if this system is in thermodynamic equilibrium, the left hand side of this equation is zero and so,

$$I_\nu = S_\nu = B_\nu(T_M) \tag{A.6}$$

where $B_\nu(T_M)$ is the Planck blackbody function for the medium at temperature T_M. We can solve Equation A.5 by multiplying both sides by an integrating factor $\exp(\tau_\nu^1)$,

$$\exp(\tau_\nu^1)\frac{dI_\nu(s)}{d\tau_\nu} + \exp(\tau_\nu^1)I_\nu(s) = \exp(\tau_\nu^1)S_\nu \tag{A.7}$$

where the lefthand side is by the product rule

$$\frac{d}{d\tau}\left(I_\nu \exp(\tau_\nu^1)\right). \tag{A.8}$$

We can now integrate both sides by τ_ν^{11}, between the limits of $0, \tau$ and do some change of variables in the integration to avoid confusion. This becomes

$$I_\nu \exp(\tau_\nu^1) = \int_0^\tau S_\nu \exp(\tau_\nu^{11}) d\tau_\nu^{11} + CONSTANT. \tag{A.9}$$

We note that when $\tau = 0$, the constant of integration is merely the background intensity $I_\nu(0)$ and so we can simplify the above equation to

$$I_\nu = \int_0^\tau S_\nu \exp(\tau_\nu^{11} - \tau_\nu^1) d\tau_\nu^{11} + I_\nu(0)\exp(-\tau_\nu^1). \tag{A.10}$$

In a uniform medium, S_{nu} has no distance dependence and the integration is simply

$$I_\nu = S_\nu(1 - \exp(-\tau)) + I_\nu(0)\exp(-\tau_\nu^1). \tag{A.11}$$

If we assume there is no background intensity in this case, we can discard the final term as it becomes zero. If we substitute for the Planck Blackbody function as in Equation A.6 we then end up with

$$I_\nu = B_\nu(T_M)(1 - \exp(-\tau)). \tag{A.12}$$

Which we can compare to the greybody equation used previously throughout this work. To get the mass of a prestellar core from observations, we note that the mass can be found from the mass surface density of material Σ by integrating over the area of the core,

$$M = \int \Sigma dA = \mu_{H_2} m_H \int N(H_2) dA. \tag{A.13}$$

Here μ_{H_2} is the mean molecular weight of an hydrogen molecule, m_H is the mass of the hydrogen atom and $N(H_2)$ is the column density of molecular hydrogen. Here A is the area of the core, and we convert this to a solid angle $d\Omega$ by noting that $dA = D^2 d\Omega$ where D is the distance to the core. This means we now have

$$M = \mu_H m_H D^2 \int N(H_2) d\Omega. \tag{A.14}$$

Here we define optical depth as,

$$\tau = \int \kappa \rho ds \tag{A.15}$$

where κ is the dust opacity and a measure of the absorption per unit mass of the medium. ρ is the mass density. This means that

$$N(H_2) = \int n(H_2) ds = \int \frac{\rho}{\mu_H m_H} ds, \tag{A.16}$$

and $n(H_2)$ is the number density of molecular hydrogen.

If we consider **optically thin** emission then $\tau \ll 1$. Equation A.12 then simplifies to

$$I_\nu \approx B_\nu(T_M)\tau. \tag{A.17}$$

This means that

$$\tau = \frac{I_\nu}{B_\nu(T_M)}. \tag{A.18}$$

We then substitute this result into Equation A.16 to obtain that

$$N(H_2) = \frac{1}{\kappa \mu_H m_H} \int \kappa \rho ds = \frac{\tau_\nu}{\kappa \mu_H m_H} = \frac{I_\nu}{B_\nu(T_M) \kappa \mu_H m_H}. \tag{A.19}$$

We can now substitute this result for the column density into Equation A.14 to find that

$$M = \mu_H m_H D^2 \int \frac{I_\nu}{B_\nu(T_M)\kappa\mu_H m_H} d\Omega. \tag{A.20}$$

This simplifies to

$$M = \frac{D^2}{B_\nu(T_M)\kappa} \int I_\nu d\Omega. \tag{A.21}$$

Finally we define $S_\nu^I = \int I_\nu d\Omega$ as the integrated intensity over the solid angle of the core at a frequency ν to get the result that

$$M_\nu = \frac{S_\nu^I D^2}{B_\nu(T_M)\kappa}. \tag{A.22}$$

We define M_ν as the mass of the core, using the observations at a frequency ν over the solid angle of the core.

A.2 Modelling a prestellar core

If we consider a prestellar core as an isothermal spherically symmetric matter distribution, then for it to be in stable equilibrium all forces must balance. If we consider the balance between pressure forces acting outwards, and gravitational attraction acting inwards then we obtain the following equation

$$\frac{1}{\rho(r)}\frac{dP(r)}{dr} = -\frac{GM(r)}{r^2} = -\frac{G}{r^2}\int_0^r 4\pi\rho(r)r^2\ dr, \tag{A.23}$$

where $P(r) = \rho(r)a^2$. $P(r)$ is the pressure at a radius r, and $\rho(r)$ is the density at a radius r and G is the gravitational constant. a is the isothermal sound speed and is defined as

$$a = \sqrt{\frac{kT}{\bar{m}}} \tag{A.24}$$

where k is Boltzmann's constant, T is the temperature of the core and $\bar{m}$ is the mean mass of the particle. If we multiply both sides of A.23 then we obtain

$$\frac{r^2}{\rho(r)}\frac{dP(r)}{dr} = -4\pi G\int_0^r \rho(r)r^2\ dr, \tag{A.25}$$

and if we then derive both sides with respect to r we obtain

$$\frac{d}{dr}\left[\frac{r^2}{\rho(r)}\frac{dP(r)}{dr}\right] = -4\pi G\rho(r). \tag{A.26}$$

Dividing both sides by ar^2 we obtain (noting a is independent of r)

$$\frac{1}{r^2}\frac{d}{dr}\left[\frac{r^2}{a^2\rho(r)}\frac{dP(r)}{dr}\right] = -\frac{4\pi G\rho(r)}{a^2}. \tag{A.27}$$

On the left hand side, we can simplify the contents of the brackets to note that $dP(r) = a^2 d\rho(r)$. This means that

$$\frac{r^2}{a^2\rho(r)}\frac{dP(r)}{dr} \equiv \frac{a^2r^2}{a^2\rho(r)}\frac{d\rho(r)}{dr} \equiv r^2\frac{d\rho(r)}{\rho(r)dr} \equiv r^2\frac{d\ ln\ \rho(r)}{dr}. \tag{A.28}$$

leaving our equation of state, Poisson's Equation, in this case as

$$\frac{1}{r^2}\frac{d}{dr}\left[r^2\frac{d\ ln\ \rho(r)}{dr}\right] = -\frac{4\pi G\rho(r)}{a^2}. \tag{A.29}$$

A.3 A singular solution to the hydrostatic equation of state

If we consider a sphere of infinite radius, then we can consider a power law relationship between radius and density such that

$$\rho(r) = \rho_0 \left(\frac{r}{r_0}\right)^{-\alpha} \tag{A.30}$$

where r_0, ρ_0 are central values and α is the slope of the density. This solution is singular as $\rho \to \infty$ as $r \to 0$. Substituting this into our equation of state, we note that

$$ln\ \rho(r) = -\alpha\ ln\ \rho_0 \frac{r}{r_0}, \tag{A.31}$$

and by using the chain rule

$$\frac{d\ ln\ \rho(r)}{dr} = -\alpha * \frac{r_0}{r} * \frac{1}{r_0} = -\frac{\alpha}{r}, \tag{A.32}$$

meaning our equation of state is now

$$\frac{1}{r^2}\frac{d}{dr}\left(-\alpha r\right) = -\frac{\alpha}{r^2} = -\frac{4\pi G}{a^2}\rho_0 \left(\frac{r}{r_0}\right)^{-\alpha}. \tag{A.33}$$

If we wish to define the central density ρ_0 as indenpendent on the radius of the sphere, and solely on r_0 then this equation can only be true if $\alpha = 2$ as then we note all powers of r drop out of the equation and we end up with

$$\frac{2}{r^2} = \frac{4\pi G}{a^2}\rho_0 \frac{r_0^2}{r^2}, \tag{A.34}$$

which simplifies to

$$\rho_0 = \frac{a^2}{2\pi G r_0^2} \tag{A.35}$$

which is singular, as by substitution into Equation A.30, we obtain

$$\rho(r) = \frac{a^2}{2\pi G r^2}, \tag{A.36}$$

and $\rho(r) \to \infty$ as $r \to 0$. We can find the mass inside a radius r by noting that by substitution into our earlier definition for the mass within a radius, we find that

$$M(r) = \int_0^r 4\pi r^2 \rho(r) dr = \frac{2a^2}{G} \int_0^r dr = 2\frac{a^2 r}{G} \tag{A.37}$$

A.4 The Jeans Mass and Radius

If we consider a spherical cloud of isothermal, ideal gas within a medium of pressure P_{ext}, with a mass of M, a radius R and a uniform density of ρ. The mass is simply

$$M = \frac{4}{3}\pi R^2 \rho, \tag{A.38}$$

$$\therefore \rho = \frac{3M}{4\pi R^3} \tag{A.39}$$

and we can define the isothermal sound speed as before as

$$a = \sqrt{\frac{kT}{\bar{m}}}, \tag{A.40}$$

where k is Boltzmann's constant, T is the temperature and $\bar{m}$ is the mean particle mass. This means the internal pressure P_{int} is simply

$$P_{int} = a^2 \rho = \frac{3a^2 M}{4\pi R^3}. \tag{A.41}$$

For a uniform density, we note that within a radius r the particle velocity within this is given by

$$\frac{dr}{dt} = \frac{r}{R}\frac{dR}{dt}, \tag{A.42}$$

where dR/dt is the overall velocity of expansion or contraction of the cloud. Since kinetic energy is generally given by the expression $1/2mv^2$, we can integrate this general expression over the whole cloud to get the total kinetic energy of the cloud as

$$E_{ke} = \frac{1}{2}\int_0^R \rho 4\pi r^2 \left(\frac{r}{R}\frac{dR}{dt}\right)^2 dr, \tag{A.43}$$

$$\therefore E_{ke} = \frac{2\pi\rho R^3}{5}\left(\frac{dR}{dt}\right)^2 = \frac{3M}{10}\left(\frac{dR}{dt}\right)^2. \tag{A.44}$$

We can consider the potential energy of the cloud as its gravitational potential energy E_{pe}, by integrating the general expression $E_{pe} = -GMm/r$ over all points of the sphere. We define $m(r)$ as the mass of the sphere interior to a radius r therefore

$$m(r) = \int_0^r 4\pi r^2 \rho dr \tag{A.45}$$

and the general expression becomes

$$E_{pe} = -\int_0^R G\frac{4\pi\rho r^3}{3} 4\pi r^2 \rho dr \tag{A.46}$$

$$\therefore E_{pe} = -\frac{16}{15}\pi^2\rho^2 G\int_0^R r^4 dr = -\frac{16}{15}(\pi\rho)^2 G R^5. \tag{A.47}$$

We now substiture for M as given in the beginning of this section to obtain the expression for the sphere's gravitational potential energy as

$$E_{pe} = -\frac{3GM^2}{5R}. \tag{A.48}$$

We can now use the law of the conservation of energy to derive an equation of state for this cloud. We do so by balancing the rate of change of total energy of the cloud against the work done $P\ dV/dt$ in changing the volume of the cloud in expansion or contraction, leading us to the expression

$$\frac{d}{dt}\left(E_{ke} + E_{pe}\right) \equiv \left[P_{int} - P_{ext}\right]\frac{dV}{dt}. \tag{A.49}$$

The left hand side is simply

$$\frac{d}{dt}\left(\frac{3M}{10}\left(\frac{dR}{dt}\right)^2 - \frac{3GM^2}{5R}\right) \tag{A.50}$$

$$= \frac{3M}{5}\frac{dR}{dt}\frac{d^2R}{dt^2} + \frac{3GM^2}{5R^2}\frac{dR}{dt} \tag{A.51}$$

and the right hand side is

$$[P_{int} - P_{ext}]4\pi R^2\frac{dR}{dt}. \tag{A.52}$$

We have already defined P_{int} earlier, and so we substitute that into the right hand side and equate A.51 and A.51 to get the expression

$$\frac{3M}{5}\frac{dR}{dt}\frac{d^2R}{dt^2} + \frac{3GM^2}{5R^2}\frac{dR}{dt} = \left[\frac{3a^2M}{4\pi R^3} - P_{ext}\right]4\pi R^2\frac{dR}{dt}. \tag{A.53}$$

We can rearrange this expression algebraically and cancel out common terms to obtain the simpler form

$$\frac{d^2R}{dt^2} = -\frac{GM}{R^2} + \frac{5a^2}{R} - 20\frac{\pi P_{ext}R^2}{3M}. \tag{A.54}$$

We note that this energy equation can be integrated over the radius of the cloud to obtain a potential function which characterises the expansion of the contraction of the cloud. If we define

$$-\frac{d\phi}{dR} = \frac{d^2R}{dt^2} = -\frac{GM}{R^2} + \frac{5a^2}{R} - 20\frac{\pi P_{ext}R^2}{3M}, \tag{A.55}$$

then integrating with respect to R yields our required potential function, which contains three terms. These terms cover the clouds own self-gravity, internal and external pressure respectively with respect to an arbitary radius R_0 which is the constant of integration that sets a zero to this derived potential function.

$$\phi(R) = \phi_{self\ gravity}(R) + \phi_{internal\ pressure}(R) + \phi_{external\ pressure}(R) \tag{A.56}$$

$$\phi(R) = -GM\left(\frac{1}{R} - \frac{1}{R_0}\right) - 5a^2 ln\ \frac{R}{R_0} + \frac{20\pi P_{ext}}{9M}(R^3 - R_0^3) \tag{A.57}$$

A.4.1 Freefall collapse times

If we ignore pressure forces, then we can simplify our expression to become simply

$$\frac{d^2R}{dt^2} = -\frac{GM}{R^2}, \tag{A.58}$$

and this can be integrated as follows, first by multiplying by an integrating factor on both sides

$$2\frac{dR}{dt}\frac{d^2R}{dt^2} = -2\frac{dR}{dt}\frac{GM}{R^2}, \tag{A.59}$$

and is integrated to obtain

$$\left(\frac{dR}{dt}\right)^2 = 2\frac{GM}{R} + Constant. \tag{A.60}$$

We can remove the constant of integration by setting a radius R_0 which at time $t = 0$ is constant, and hence $dR/dt = 0$. This means that

$$\left(\frac{dR}{dt}\right) = -\sqrt{2GM\left(\frac{1}{R} - \frac{1}{R_0}\right)}. \tag{A.61}$$

As it is a square root, the sign is chosen as the negative to denote inward collapse. Solving this differential equation requires a substitution of the form

$$R = R_0 \cos^2\theta, \tag{A.62}$$

which gives us

$$\frac{dR}{dt} = -2R_0 \cos\theta \sin\theta \frac{d\theta}{dt}, \tag{A.63}$$

and by substitution the right hand side becomes

$$-\sqrt{\frac{2GM}{R_0}}\tan\theta. \tag{A.64}$$

$$\therefore 2R_0\cos\theta\sin\theta\frac{d\theta}{dt} = \sqrt{\frac{2GM}{R_0}}\tan\theta. \tag{A.65}$$

Rearranging terms, and cancelling out trigonometrical items leaves

$$2\cos^2\theta\frac{d\theta}{dt} = (1+\cos 2\theta)\frac{d\theta}{dt} = \sqrt{\frac{2GM}{R_0^3}}. \tag{A.66}$$

This expression is integrated with respect to time to obtain

$$\theta + \frac{sin2\theta}{2} = \sqrt{\frac{2GM}{R_0^3}}t, \tag{A.67}$$

or

$$\theta + \sin\theta\cos\theta = \sqrt{\frac{2GM}{R_0^3}}t. \tag{A.68}$$

We can now re-substitute our definition of R_0 in terms of θ by noting that

$$\theta = \cos^{-1}\sqrt{\frac{R}{R_0}} \tag{A.69}$$

$$\frac{R}{R_0} = \cos^2\theta, \tag{A.70}$$

$$\therefore \cos\theta = \sqrt{\frac{R}{R_0}}, \tag{A.71}$$

$$\sin^2\theta = 1 - \cos^2, \tag{A.72}$$

$$\theta \therefore \sin\theta = \sqrt{1 - \frac{R}{R_0}}. \tag{A.73}$$

Substituting these gives us

$$\sqrt{\frac{2GM}{R_0^3}}t = \cos^{-1}\sqrt{\frac{R}{R_0}} + \sqrt{\frac{R}{R_0}}\sqrt{1 - \frac{R}{R_0}}. \tag{A.74}$$

To define a freefall time, we seek a time t_{ff} such that at $t = 0$, $R = R_0$ and at $t = t_{ff}$, $R = 0$. This would give us a time for collapse of the cloud when only gravitational forces are at work. If we define

$$t_{ff} = \sqrt{\frac{\pi^2 R_0^3}{8GM}} \equiv \sqrt{\frac{3\pi}{32G\rho_0}} \tag{A.75}$$

then Equation A.74 can be rewritten by substitution and rearrangement into

$$t = \frac{2t_{ff}}{\pi}\left(\cos^{-1}\sqrt{\frac{R}{R_0}} + \sqrt{\frac{R}{R_0}}\sqrt{1 - \frac{R}{R_0}}\right), \tag{A.76}$$

where we note substituting $R = R_0$ yields $t = 0$, and $R = 0$ yields $t = t_{ff}$ as required.

A.4.2 The Jeans Mass

For a cloud that is within a medium, the pressure P_{ext} is defined, if there is no difference in density between cloud and medium, as

$$P_{ext} = \rho a^2, \tag{A.77}$$

if we define the Jeans length as

$$R_{Jeans} = a\sqrt{\frac{15}{4\pi G\rho}}, \tag{A.78}$$

then substituting these into Equation A.57, and rearranging yields

$$\frac{\phi(R)}{a^2} = -\left(\frac{R_0}{R_{Jeans}}\right)^2\left(\frac{R_0}{R} - 1\right) - ln\frac{R}{R_0} + \frac{1}{3}\left(\left(\frac{R}{R_0}\right)^3 - 1\right) \tag{A.79}$$

This is important, as we can consider three cases. For a large radius greater than the Jeans length the pressure terms are negligible and the cloud is gravitationally unstable. For a much smaller radius than the Jeans length the cloud is stable gravitationally, and it exists in equilibrium between external and internal pressure forces. When the radius is of order of the Jeans length, the three terms are comparable and all three forces balance, and this is the radius of clouds that are just able to condense out of the medium. The mass can be obtained by

$$M_{Jeans} = \frac{4}{3}\pi R_{Jeans}^3\rho \equiv \sqrt{\frac{375}{4\pi}}\frac{a^3}{\sqrt{G^3\rho}}. \tag{A.80}$$

The timescale for a core to condense out of the medium is given by

$$t_c = \frac{t_{ff}}{\sqrt{1 - \left(\frac{R_{Jeans}}{R_0}\right)^2}}, \tag{A.81}$$

where t_{ff} is the freefall time. We note for a core of mass and radius much greater then the Jeans values, the cloud will condense out in a freefall timescale. For

one where the values are comparable. the timescale lengthens considerably. As a result, these values can be considered the lowest values at which a core can form.

We note from the review by Evans (1999) that the freefall time and Jeans mass can be rewritten by substituting for the various constants as

$$t_{ff} = \sqrt{\frac{3\pi}{32G\rho_0}} \equiv \frac{3.4 \times 10^7}{\sqrt{n}} \; years \tag{A.82}$$

where n is the total particle density. Similarly,

$$M_{Jeans} = 18 \times M_{\odot} \frac{T^{1.5}}{\sqrt{n}} \tag{A.83}$$

where T is the cloud temperature.

A.4.3 The Virial Mass

The Virial theorem states that for a stable, self gravitating spherical distribution the total kinetic energy E_{ke} is related to the total potential energy by

$$E_{ke} = -\frac{1}{2} E_{pe}, \tag{A.84}$$

and the kinetic energy for the sphere is given by

$$E_{ke} = \frac{1}{2} M \Delta V^2, \tag{A.85}$$

where ΔV is the FWHM velocity of the particles within the sphere and we recall the total gravitational potential energy was derived by

$$E_{pe} = -\frac{3GM^2}{5R}. \tag{A.86}$$

Substituting and solving for M yields a value for the Virial Mass, and we note that the we convert into the dispersion of the velocities σ by a $2\sqrt{2ln\ 2}$ factor therefore

$$M_{vir} = \frac{5R\Delta V^2}{3G} \approx \frac{5R\sigma^2}{G} \tag{A.87}$$

which we term the Virial mass. Typically, the potential energy term is often written as

$$E_{pe} = -a\frac{3GM^2}{5R}, \tag{A.88}$$

where a is a numerical term, of order unity, which is included to allow for the non-uniform density of the core. Enoch et al. (2008) quoted the Virial mass in this case as

$$M_{vir} \approx \frac{5R\sigma^2}{aG}. \tag{A.89}$$

The ratio of the core mass to the Virial mass is introduced in Bertoldi & McKee (1992) as

$$\alpha = \frac{5R\sigma^2}{GM_{core}}, \tag{A.90}$$

where $\alpha = 1$ for an uniform gas sphere in virial balance, and 2 for a sphere that is self-gravitating.. This term is called the 'Virial Parameter' and is a measure of the ratio of the kinetic and gravitational energy of the core. It should be noted that σ is the velocity distribution of the average particle within the core. This is usually given a molecular mass of $\mu = 2.33$ to reflect a specific composition of molecular hydrogen and atomic helium. This is calculated from the observed dispersion as

$$\sigma^2 = \sigma_T^2 + \sigma_{NT}^2 = \sigma_{obs}^2 + \frac{kT}{m_H}\left(\frac{1}{\mu} - \frac{1}{\mu_{observed}}\right), \tag{A.91}$$

from Myers, Ladd, & Fuller (1991). Here we note that σ_T, σ_{NT} are the thermal and non-thermal velocity dispersions. They are defined as

$$\delta V_T^2 = 8ln\ 2\frac{kT}{m} \tag{A.92}$$

$$\delta v_{NT}^2 = \delta v_{obs}^2 - 8ln\ 2\frac{kT}{m_{obs}} \tag{A.93}$$

where m_{obs} is the molecular mass of the observed molecule. The true velocity is then calculated in quadrature from the thermal and non thermal components as followed.

$$\therefore \delta_v^2 = \delta v_T^2 + \delta v_{NT}^2 = \delta v_{obs}^2 + 8ln\ 2\frac{kT}{m}\left(\frac{1}{\mu} - \frac{1}{\mu_{obs}}\right) \tag{A.94}$$

and again converting from the observed FWHM into velocity dispersions, we obtain Equation A.91.

A.5 RSRF Polynomial Coefficients

Filter	Wavenumber Range	x^6	x^5	x^4	x^3
PLW					
	$15.32 - 16.88 cm^{-1}$	-2.52242864	241.04408908	-9595.11797987	203653.63667422
	$16.88 - 23.23 cm^{-1}$	-0.0005109094	0.06047379 41	-2.9718358236	77.6075860072
	$23.23 - 26.06 cm^{-1}$	0	0	-0.04824892	4.89200885
PMW					
	$23.62 - 27.23 cm^{-1}$	0.0133273379	-2.0308141737	128.8838728378	-4360.55227404
	$27.23 - 31.53 cm^{-1}$	0.0014389717	-0.2540214957	18.6747946333	-731.8398144601
	$31.53 - 37.19 cm^{-1}$	0.001870247 15	-0.3869626748	33.32979299747	-1529.61403545378
PSW					
	$33.09 - 35.53 cm^{-1}$	0	0	0	-0.1651385566
	$35.53 - 38.55 cm^{-1}$	0	0.00609282342	-1.14898705743	86.59858661053
	$38.55 - 42.17 cm^{-1}$	-0.00133755835	0.31856149414	-31.59289834793	1669.96163640184
	$42.17 - 51.73 cm^{-1}$	-0.00007266859	0.02018677428	-2.33299964310	143.58329323505
Filter	Wavenumber Range	x^2	x^1	x^0	
PLW					
	$15.32 - 16.88 cm^{-1}$	-2430790.69891723	15470144.21521960	-41013248.97317070	
	$16.88 - 23.23 cm^{-1}$	-1135.8097275101	8832.4371195473	-28508.7097146106	
	$23.23 - 26.06 cm^{-1}$	-185.67700112	3126.40111630	-19702.11927600	
PMW					
	$23.62 - 27.23 cm^{-1}$	82951.7143712965	-841256.6038914770	3553385.6980547500	
	$27.23 - 31.53 cm^{-1}$	16124.0062058862	-189365.7274972160	926170.5519777940	
	$31.53 - 37.19 cm^{-1}$	39449.33686673770	-542098.66300486400	3100879.63238794000	
PSW					
	$33.09 - 35.53 cm^{-1}$	17.0523675622	-586.2926186609	6712.2567476506	
	$35.53 - 38.55 cm^{-1}$	-3260.81207226290	61343.65895160610	-461254.61901518700	
	$38.55 - 42.17 cm^{-1}$	-49620.72827387060	785833.11899702300	-5181941.43146082000	
	$42.17 - 51.73 cm^{-1}$	-4963.21991479341	91364.02330386770	-699735.50789079200	

Table A.1: RSRF Polynomial Coefficients.

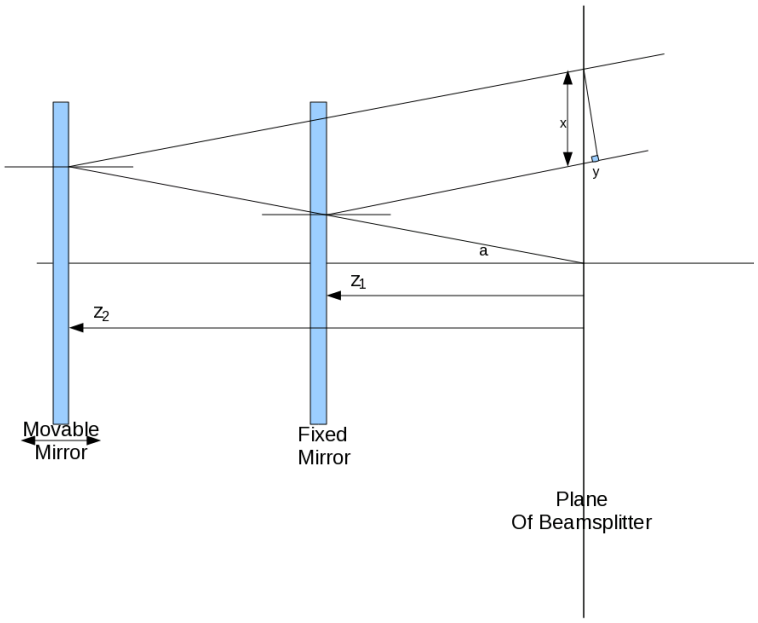

Figure A.2: Block diagram of off-axis ray in FTS.

A.6 Extended field of view for a single pixel FTS

For an off-axis ray in an extended field-of-view FTS, the optical path is shown in Figure A.2. This diagram is stylized, folding both optical paths into the same plane for ease of view and showing the plane of the beamsplitter which is the point the off-axis ray exits the beamsplitter.

The on-axis system has an optical path difference given by $z = 2(z_2 - z_1)$, and for the off-axis system we can similarly define z^o where a is the off-axis angle. The superscript ^{o} defines an off-axis ray. For the ray that travels to the movable mirror, the optical path z_2^o is simply equivalent to:

$$z_2^o = \frac{2z_2}{\cos a}. \tag{A.95}$$

For z_1^0 we note there is a small extra path y such that:

$$z_1^o = \frac{2z_1}{\cos a} + y. \tag{A.96}$$

We can find y by using the symmetry of the system, by noting that the distance x is given by:

$$x = 2z_2 \tan a - 2z_1 \tan a. \tag{A.97}$$

Here we note the distance from the optical plane and the point of contact with the movable mirror is given by $z_2 \tan a$, and likewise is $z_1 \tan a$ for the fixed mirror. When the beam returns to the line marked *Plane of Beamsplitter* this distance is doubled. Again, by symmetry, we note that:

$$\sin a = \frac{y}{x} \rightarrow y = x \sin a \tag{A.98}$$

and we can simply the expression for y such that:

$$y = 2(z_2 - z_1) \tan a \sin a \equiv 2(z_2 - z_1)\frac{\sin^2 a}{\cos a} \tag{A.99}$$

and by substitution back into our oblique ray path difference expression:

$$z_1^o = \frac{2z_1}{\cos a} + 2(z_2 - z_1)\frac{\sin^2 a}{\cos a}. \tag{A.100}$$

Since we can define the total oblique optical path difference as:

$$z^o = (z_2^o - z_1^o), \tag{A.101}$$

we find that:

$$z^o = \frac{2z_2}{\cos a} - \left(\frac{2z_1}{\cos a} + 2(z_2 - z_1)\frac{\sin^2 a}{\cos a}\right), \tag{A.102}$$

and by collecting like terms, we find that:

$$z^o = \frac{2(z_2 - z_1) - 2\sin^2 a(z_2 - z_1)}{\cos a}. \tag{A.103}$$

This further simplifies to become:

$$z^o = \frac{2(z_2 - z_1)(1 - \sin^2 a)}{\cos a}. \tag{A.104}$$

We note from earlier that $z = 2(z_2 - z_1)$ and apply a trigonometric substitution that $(1 - \sin^2 a) \equiv \cos^2 a$ and get the final result that:

$$z^o = z \cos a \tag{A.105}$$

A.7 Derivation of integral for an interferometer

An input signal to the interferometer will fall upon the beamsplitter, and be split into the two paths. A path difference will cause the two light paths to create interference when recombined.

For a path difference z, there is an associated phase difference P given by

$$P = \frac{2\pi z}{\lambda} \equiv 2\pi\sigma z. \tag{A.106}$$

defining σ as the wavenumber. If the incident beam has an amplitude A then the resultant amplitude for both paths is given by

$$A_D = \frac{A}{2} + \frac{A \exp 2\pi i\sigma z}{2}, \tag{A.107}$$

for the static and movable mirror path respectively. The intensity through the system is given by

$$I_{sys}(z) = A_D^* A_D = \frac{A^2}{4}(1 + \exp 2\pi i\sigma z)(1 + \exp -2\pi i\sigma z), \tag{A.108}$$

which we can simplify to, by using Euler's formula and defining θ as

$\theta \equiv 2\pi\sigma z$

$\exp i\theta = \cos\theta + i\sin\theta$

$\exp -i\theta = \cos\theta - i\sin\theta$

$\therefore I_{sys}(z) = \frac{A^2}{4}(1 + \exp i\theta)(1 + \exp -i\theta)$

$= \frac{A^2}{4}(1 + \cos\theta + i\sin\theta)(1 + \cos\theta - i\sin\theta)$

$= \frac{A^2}{4}(1 + \cos\theta - i\sin\theta + \cos\theta + \cos^2\theta - i\sin\theta\cos\theta + i\sin\theta + i\sin\theta\cos\theta - i^2\sin^2\theta)$

$= \frac{A^2}{4}(1 + 2\cos\theta + \cos^2\theta + \sin^2\theta)$

$= \frac{A^2}{4}(2 + 2\cos\theta)$

$I_0 \equiv A^2$

$$\therefore I_{sys}(z) = \frac{I_0}{2}(1 + \cos 2\pi\sigma z), \tag{A.109}$$

and this resultant intensity is an interferogram taken over a range of path difference z. For a general input $S(\sigma)$ we modify the output interferogram to take account the sum of all the monochromatic inputs and this becomes

$$I_{sys}(z) = \frac{1}{2}\int_0^\infty S(\sigma)(1 + \cos 2\pi\sigma z)d\sigma \equiv \frac{I_0}{2} + \frac{1}{2}\int_0^\infty S(\sigma)\cos 2\pi k\sigma d\sigma, \tag{A.110}$$

which we can inverse Fourier transform to get the result that

$$S(\sigma) \equiv 4\int_0^\infty \left(I_{syz}(z) - \frac{I_0}{2}\right)\cos 2\pi kzdz. \tag{A.111}$$

Bell (1972) yields a more complete examination of the subject.

A.8 Beamsplitter offset calculation

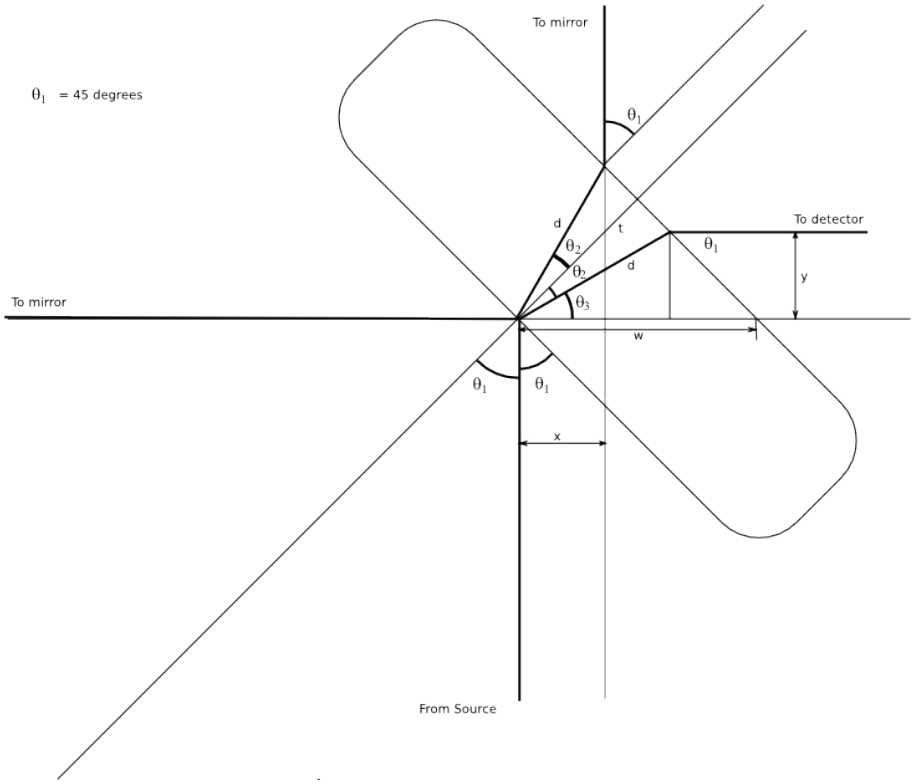

Figure A.3: Beamsplitter Offset calculations.

The beamsplitter we have used is of thickness $8mm$ and has a refractive index in the mid band of $n_{11\mu m}$ of 1.52. For the exterior, which will ultimately be vacuum sealed, we assume $n = 1$.

We can calculate the vertical offset x and the horizontal offset y and show they are identical.

$\theta_1 = 45^o$ and is the angle of the beamsplitter to the optical path. The beamsplitter surface is that which the source ray first contacts in this derivation.

We note that from Snell's law, where $n_1 = 1, n_2 = 1.52$:

$$n_1 \sin\theta_1 = n_2 \sin\theta_2 \tag{A.112}$$

$$1.0 * \frac{\sqrt{2}}{2} = 1.52 \sin\theta_2 \tag{A.113}$$

$$\theta_2 = \arcsin \frac{\sqrt{2}}{3.04} = 27.72^o \tag{A.114}$$

We start by finding w, by examining one part of the diagram:

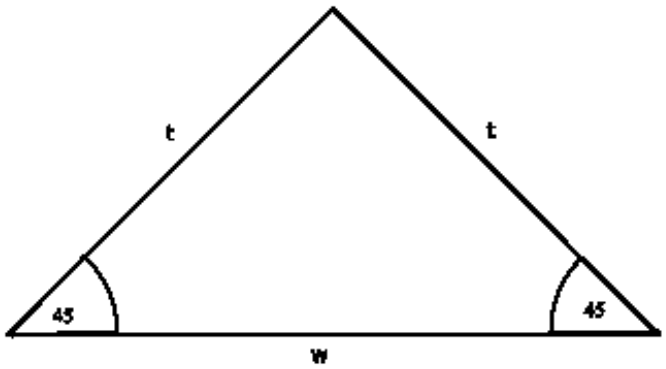

Figure A.4: To find w.

Here $t = 8mm$ is the thickness of the beamsplitter, and for this right-angled triangle $w = \sqrt{d^2 + d^2} = 11.31mm$. We can now find d by examining the θ_2 triangle.

$$\begin{aligned} \cos\theta_2 &= \frac{t}{d} \\ \therefore d &= \frac{t}{\cos\theta_2} \\ = \frac{8.0}{0.885} mm &= 9.04mm \end{aligned} \tag{A.115}$$

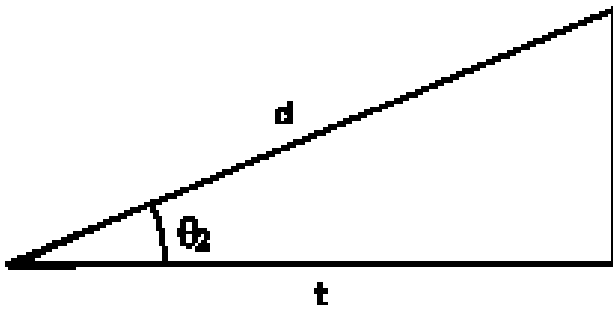

Figure A.5: To find d.

From d the θ_3 triangle gives us y, the horizontal offset. Here we note that $\theta_3 = 45^o - \theta_2 = 17.28^o$

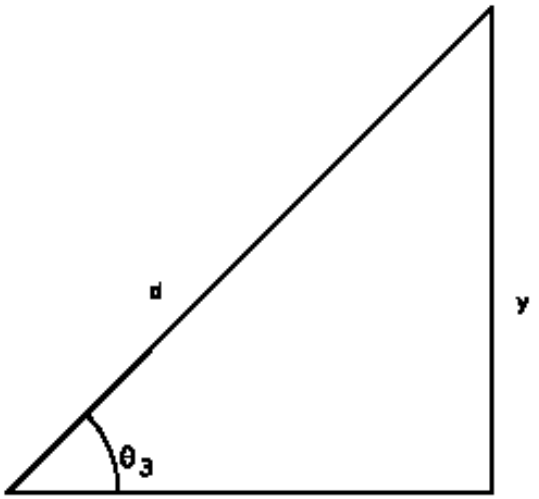

Figure A.6: To find y.

$$y = d \sin 17.28^o = 9.04 * 0.297mm = 2.68mm \tag{A.116}$$

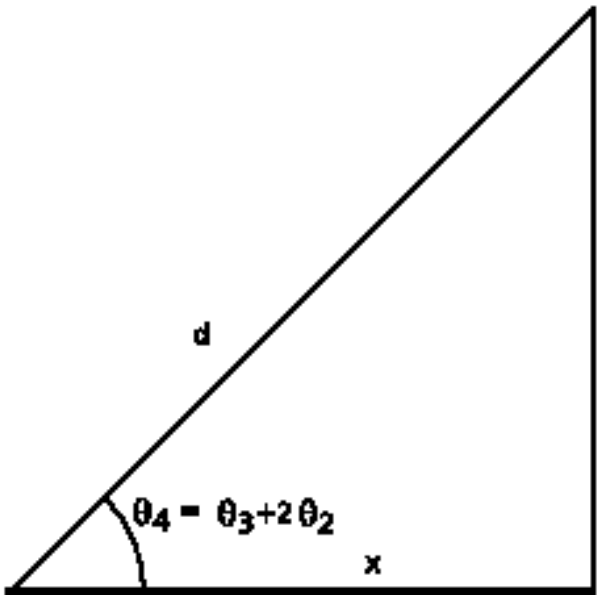

Figure A.7: To find x.

For symmetry, we expect that x, the horizontal offset is the same. Here we note that $\theta_4 = \theta_3 + 2\theta_2 = 45^o + \theta_2 = 72.72^o$.

$$\begin{aligned} \cos 72.72^o &= \frac{x}{d} \\ \therefore x &= 9.04 * \cos 72.72^o \\ &= 2.68mm \end{aligned} \tag{A.117}$$

Q.E.D

A.9 SLIDE Usage and code

SLIDE has been involved in data reduction for the following:

A.9.1 Published work

van der Werf et al. (2010b), From the abstract:

We present a full high resolution SPIRE FTS spectrum of the nearby ultraluminous infrared galaxy Mrk 231. In total 25 lines are detected, including CO J = 5-4 through J = 13-12, 7 rotational lines of H^2O, 3 of OH^+ and one line each of H_2O^+, CH^+, and HF. We find that the excitation of the CO rotational levels up to J = 8 can be accounted for by UV radiation from star formation. However, the approximately flat luminosity distribution of the CO lines over the rotational ladder above J = 8 requires the presence of a separate source of excitation for the highest CO lines. We explore X-ray heating by the accreting supermassive black hole in Mrk 231 as a source of excitation for these lines, and find that it can reproduce the observed luminosities. We also consider a model with dense gas in a strong UV radiation field to produce the highest CO lines, but find that this model strongly overpredicts the hot dust mass in Mrk 231. Our favoured model consists of a star forming disk of radius 560 pc, containing clumps of dense gas exposed to strong UV radiation, dominating the emission of CO lines up to J = 8. X-rays from the accreting supermassive black hole in Mrk 231 dominate the excitation and chemistry of the inner disk out to a radius of 160 pc, consistent with the X-ray power of the AGN in Mrk 231. The extraordinary luminosity of the OH+ and H2O+ lines reveals the signature of X-ray driven excitation and chemistry in this region.

Panuzzo et al. (2010b), from the abstract:

We present the observations of the starburst galaxy M82 taken with the Herschel SPIRE Fourier-transform spectrometer. The spectrum (194-671 micron) shows a prominent CO rotational ladder from J = 4-3 to 13-12 emitted by the central region of M82. The fundamental properties of the gas are well constrained by the high J lines observed for the first time. Radiative transfer modeling of these high-S/N ^{12}CO and ^{13}CO lines strongly indicates a very warm molecular gas component at $\approx$ 500 K and pressure of $\approx 3\times10^6 K cm^{-3}$, in good agreement with the H2 rotational lines measurements from Spitzer and ISO. We suggest that this warm gas is heated by dissipation of turbulence in the interstellar medium (ISM) rather than X-rays or UV flux from the straburst. This paper illustrates the promise of the SPIRE FTS for the study of the ISM of nearby galaxies.

Herschel is an ESA space observatory with science instruments provided by

European-led Principal Investigator consortia and with important participation from NASA.

González-Alfonso et al. (2010b), from the abstract:

The Ultra luminous infrared galaxy (ULIRG) Mrk 231 reveals up to seven rotational lines of water (H_2O) in emission, including a very high-lying (Eupper = 640 K) line detected at a 4σ level, within the Herschel/SPIRE wavelength range (190 ¡ (Wavelength in microns) ¡ 640), whereas PACS observations show one H2O line at 78 micron in absorption, as found for other H_O lines previously detected by ISO. The absorption/emission dichotomy is caused by the pumping of the rotational levels by far-infrared radiation emitted by dust, and subsequent relaxation through lines at longer wavelengths, which allows us to estimate both the column density of H_2O and the general characteristics of the underlying far-infrared continuum source. Radiative transfer models including excitation through both absorption of far-infrared radiation emitted by dust and collisions are used to calculate the equilibrium level populations of H2O and the corresponding line fluxes. The highest-lying H2O lines detected in emission, with levels at 300-640 K above the ground state, indicate that the source of far-infrared radiation responsible for the pumping is compact (radius = 110-180 pc) and warm (Tdust = 85-95 K), accounting for at least 45% of the bolometric luminosity. The high column density, N(H_2O) $\approx 5\times10^{17} cm^{-2}$, found in this nuclear component, is most probably the consequence of shocks/cosmic rays, an XDR chemistry, and/or an undepleted chemistry where grain mantles are evaporated. A more extended region, presumably the inner region of the 1-kpc disk observed in other molecular species, could contribute to the flux observed in low-lying H_2O lines through dense hot cores, and/or shocks. The H_2O 78 micron line observed with PACS shows hints of a blue-shifted wing seen in absorption, possibly indicating the occurrence of H_2O in the prominent outflow detected in OH (Fischer et al. 2010, AA, 518, L41). Additional PACS/HIFI observations of H2O lines are required to constrain the kinematics of the nuclear component, as well as the distribution of H_2O relative to the warm dust.

Herschel is an ESA space observatory with science instruments provided by European-led Principal Investigator consortia and with important participation from NASA.

Rangwala et al. (2011), from the abstract:

We present Herschel SPIRE-FTS observations of Arp-220, a nearby ULIRG. The FTS continuously covers 190 – 670 microns, providing a good measurement of the continuum and detection of several molecular and atomic species. We detect luminous CO (J = 4-3 to 13-12) and water ladders with comparable total luminosity; very high-J HCN absorption; OH+, H_2O^+, and HF in absorption; and CI and NII. Modeling of the continuum yields warm dust, with T = 66 K, and an unusu-

ally large optical depth of ≈5 at 100 microns. Non-LTE modeling of the CO shows two temperature components: cold molecular gas at T ≈ 50 K and warm molecular gas at T ≈ 1350 K. The mass of the warm gas is 10% of the cold gas, but dominates the luminosity of the CO ladder. The temperature of the warm gas is in excellent agreement with H_2 rotational lines. At 1350 K, H_2 dominates the cooling ($\approx 20 L_\odot/M_\odot$) in the ISM compared to CO ($\approx 0.4 L_\odot/M_\odot$). We found that only a non-ionizing source such as the mechanical energy from supernovae and stellar winds can excite the warm gas and satisfy the energy budget of $\approx 20 L_\odot/M_\odot$. We detect a massive molecular outflow in Arp 220 from the analysis of strong P-Cygni line profiles observed in OH+, H_2O^+, and H_2O. The outflow has a mass $> 10^7 M_\odot$ and is bound to the nuclei with velocity $< 250 km/s$. The large column densities observed for these molecular ions strongly favor the existence of an X-ray luminous AGN ($10^{44} ergs/s$) in Arp 220.

Sacchi et al. (2011), Abstract:

We present the 450-1550 GHz spectra of three nearby Seyfert galaxies (NGC1068, NGC7130 and NGC7582) taken with the Herschel SPIRE FTS. For the case of NGC1068 we reconstruct the nuclear spectral line energy distribution (SLED) of the CO lines, applying nonLTE radiative transfer and a Bayesian likelihood analysis to estimate the physical properties of the molecular gas in the circumnuclear region. Groundbased observations of the low-J transitions with high (few arcsec) angular resolution are required to reconstruct the nuclear SLED avoiding contamination from colder molecular gas on larger galactic scales. We find evidence for a very warm molecular gas component with a density $\approx 10^{3.9} cm^{-3}$, similar to that found in previous works (Papadopoulos, Seaquist 1999, Usero et al. 2004, Kamenetzky et al. 2011), but with a much higher temperature (≈ 550 K instead of 20-160 K). The higher-J transitions of CO are compatible with being excited in X-ray dissociation regions (XDR). However, in order to explain the entire CO SLED a comparable contribution from photodissociation regions (PDR) is required.

A.9.2 Work in Progress

2011 - A team based in University College London and lead by Mikako Matsuura and Mike Barlow are using SLIDE as part of a paper on VY CMa, studying the large number of lines using PACS data. It will involve about 1000 lines across, and the majority of lines are due to H_2O.

A.10 AJ Rykala Refereed Publication List

Davies et al. (2010) On the origin of M81 group extended dust emission.

Pohlen et al. (2010) Radial distribution of gas and dust in spiral galaxies. The case of M 99 (NGC 4254) and M 100 (NGC 4321).

Roussel et al. (2010) SPIRE imaging of M 82: Cool dust in the wind and tidal streams.

Bendo et al. (2010) The Herschel Space Observatory view of dust in M81.

Sauvage et al. (2010) The central region of spiral galaxies as seen by Herschel. M 81, M 99, and M 100.

Cortese et al. (2010) Herschel-SPIRE observations of the disturbed galaxy NGC 4438.

Eales et al. (2010) Mapping the interstellar medium in galaxies with Herschel/SPIRE.

Boselli et al. (2010a) FIR colours and SEDs of nearby galaxies observed with Herschel.

O'Halloran et al. (2010) Herschel photometric observations of the low metallicity dwarf galaxy NGC 1705.

Galametz et al. (2010) Herschel photometric observations of the nearby low metallicity irregular galaxy NGC 6822.

Gomez et al. (2010) The dust morphology of the elliptical Galaxy M 86 with SPIRE.

González-Alfonso et al. (2010c) Herschel observations of water vapour in Markarian 231.

van der Werf et al. (2010c) Black hole accretion and star formation as drivers of gas excitation and chemistry in Markarian 231.

Panuzzo et al. (2010c) Probing the molecular interstellar medium of M82 with Herschel-SPIRE spectroscopy.

Swinyard et al. (2010) In-flight calibration of the Herschel-SPIRE instrument.

Raymond et al. (2010) The Effectiveness of Mid IR / Far IR Blind, Wide Area, Spectral Surveys in Breaking the Confusion Limit.

Boselli et al. (2010b) The Herschel Reference Survey.

Bibliography

Andre P., Montmerle T., From T Tauri stars to protostars: Circumstellar material and young stellar objects in the rho Ophiuchi cloud, 1994, ApJ, 420, 837

Andre P., Ward-Thompson D., Barsony M., From Prestellar Cores to Protostars: the Initial Conditions of Star Formation, 2000, Protostars and Planets IV, 59

Andre P., Ward-Thompson D., Motte F., Probing the initial conditions of star formation: the structure of the prestellar core L 1689B., 1996, A&A, 314, 625

Bacmann A., André P., Puget J., Abergel A., Bontemps S., Ward-Thompson D., An ISOCAM absorption survey of the structure of pre-stellar cloud cores, 2000, A&A, 361, 555

Ballesteros-Paredes J., Klessen R. S., Mac Low M.-M., Vazquez-Semadeni E., Molecular Cloud Turbulence and Star Formation, 2007, Protostars and Planets V, 63

Beichman C. A., Myers P. C., Emerson J. P., Harris S., Mathieu R., Benson P. J., Jennings R. E., Candidate solar-type protostars in nearby molecular cloud cores, 1986, ApJ, 307, 337

Bell R. J., 1972, Introductory Fourier Transform Spectroscopy. Academic Press

Bendo G. J. et al., The Herschel Space Observatory view of dust in M81, 2010, A&A, 518, L65

Benson P. J., Myers P. C., A survey for dense cores in dark clouds, 1989, ApJS, 71, 89

Bertoldi F., McKee C. F., Pressure-confined clumps in magnetized molecular clouds, 1992, ApJ, 395, 140

Bland-Hawthorn J., Cecil G. N., Classical Spectroscopy, 1997, ArXiv Astrophysics e-prints

Bonnell I. A., Bate M. R., Clarke C. J., Pringle J. E., Accretion and the stellar mass spectrum in small clusters, 1997, MNRAS, 285, 201

Bonnor W. B., Boyle's Law and gravitational instability, 1956, MNRAS, 116, 351

Boselli A. et al., FIR colours and SEDs of nearby galaxies observed with Herschel, 2010a, A&A, 518, L61

Boselli A. et al., The Herschel Reference Survey, 2010b, PASP, 122, 261

Bourke T. L., Hyland A. R., Robinson G., Studies of star formation in isolated small dark clouds - I. A catalogue of southern BOK globules: optical and IRAS properties, 1995, MNRAS, 276, 1052

Bourke T. L., Hyland A. R., Robinson G., James S. D., Wright C. M., Studies of star formation in isolated small dark clouds - II. A southern ammonia survey, 1995, MNRAS, 276, 1067

Bracewell R. N., 1965, The Fourier Transform and its applications. McGraw-Hill

Caselli P., van der Tak F. F. S., Ceccarelli C., Bacmann A., Abundant H_2D^+ in the pre-stellar core L1544, 2003, A&A, 403, L37

Caux E., Bottinelli S., Vastel C., Glorian J. M., 2011, in IAU Symposium, Vol. 280, IAU Symposium, p. 120P

Chase D. B., Phase correction in ft-ir, 1982, Appl. Spectrosc., 36, 240

Ciolek G. E., Mouschovias T. C., Ambipolar Diffusion, Interstellar Dust, and the Formation of Cloud Cores and Protostars. IV. Effect of Ultraviolet Ionization and Magnetically Controlled Infall Rate, 1995, ApJ, 454, 194

Clemens D. P., Barvainis R., A catalog of small, optically selected molecular clouds - Optical, infrared, and millimeter properties, 1988, ApJS, 68, 257

Clemens D. P., Yun J. L., Heyer M. H., BOK globules and small molecular clouds - Deep IRAS photometry and (C-12)O spectroscopy, 1991, ApJS, 75, 877

Colbert J. W., Malkan M. A., Clegg P. E., Cox P., Fischer J., Lord S. D., Luhman M., Satyapal S., Smith H. A., Spinoglio L., Stacey G., Unger S. J., ISO LWS Spectroscopy of M82: A Unified Evolutionary Model, 1999, ApJ, 511, 721

Cortese L. et al., Herschel-SPIRE observations of the disturbed galaxy NGC 4438, 2010, A&A, 518, L63

Crapsi A., Caselli P., Walmsley C. M., Myers P. C., Tafalla M., Lee C. W., Bourke T. L., Probing the Evolutionary Status of Starless Cores through N_2H^+ and N_2D^+ Observations, 2005, ApJ, 619, 379

Crystran Ltd , 2010, Datasheet for ZnSe optical properties, Internet URL:

Davies J. I. et al., On the origin of M81 group extended dust emission, 2010, MNRAS, 409, 102

Dejonghe H., A completely analytical family of anisotropic Plummer models, 1987, MNRAS, 224, 13

di Francesco J., Evans N. J., II, Caselli P., Myers P. C., Shirley Y., Aikawa Y., Tafalla M., An Observational Perspective of Low-Mass Dense Cores I: Internal Physical and Chemical Properties, 2007, Protostars and Planets V, 17

Dohlen K., Origne A., Pouliquen D., Swinyard B. M., 2000, in Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series, Vol. 4013, J. B. Breckinridge & P. Jakobsen , ed, Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series, p. 119

Draine B. T., On the Submillimeter Opacity of Protoplanetary Disks, 2006, ApJ, 636, 1114

Dullemond C., 2005, The formation of stars and planets: Self-gravitating hydrostatic gas spheres, www.mpia.de/homes/dullemon/lectures/starplanet/chap_05_cloudcores.ppt

Eales S. A. et al., Mapping the interstellar medium in galaxies with Herschel/SPIRE, 2010, A&A, 518, L62

Ebert R., Über die Verdichtung von H I-Gebieten. Mit 5 Textabbildungen, 1955, ZAp, 37, 217

Elmegreen B. G., Lada C. J., Sequential formation of subgroups in OB associations, 1977, ApJ, 214, 725

Enoch M. L., Evans N. J., II, Sargent A. I., Glenn J., Rosolowsky E., Myers P., The Mass Distribution and Lifetime of Prestellar Cores in Perseus, Serpens, and Ophiuchus, 2008, ApJ, 684, 1240

European Space Agency , 2010, Spire observers manual, Internet URL:

Evans N. J., II, Physical Conditions in Regions of Star Formation, 1999, ARA&A, 37, 311

Evans N. J., II, Rawlings J. M. C., Shirley Y. L., Mundy L. G., Tracing the Mass during Low-Mass Star Formation. II. Modeling the Submillimeter Emission from Preprotostellar Cores, 2001, ApJ, 557, 193

Filler A., Apodization and Interpolation in Fourier-Transform Spectroscopy, 1964, J.Opt.Soc.Am, 54, 762

Galametz M. et al., Herschel photometric observations of the nearby low metallicity irregular galaxy NGC 6822, 2010, A&A, 518, L55

Genest J., Tremblay P., Instrument Line Shape of Fourier Transform Spectrometers: Analytic Solutions for Nonuniformly Illuminated Off-Axis Detectors, 1999, Appl. Opt., 38, 5438

Gomez H. L. et al., The dust morphology of the elliptical Galaxy M 86 with SPIRE, 2010, A&A, 518, L45

González-Alfonso E. et al., Herschel observations of water vapour in Markarian 231, 2010a, A&A, 518, L43

González-Alfonso E. et al., Herschel observations of water vapour in Markarian 231, 2010b, A&A, 518, L43

González-Alfonso E. et al., Herschel observations of water vapour in Markarian 231, 2010c, A&A, 518, L43

Goodwin S. P., Nutter D., Kroupa P., Ward-Thompson D., Whitworth A. P., The relationship between the prestellar core mass function and the stellar initial mass function, 2008, A&A, 477, 823

Goodwin S. P., Ward-Thompson D., Whitworth A. P., Why do starless cores appear more flattened than protostellar cores?, 2002, MNRAS, 330, 769

Griffin M. et al., 2006, in Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series, Vol. 6265, Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series

Griffin M. J. et al., The Herschel-SPIRE instrument and its in-flight performance, 2010, A&A, 518, L3

Griffiths D. H., 2007, Fourier Transform Infrared Spectrometry. Wiley Interscience, p. 85

Hecht E., 1998, Optics (Third Edition ed.). Addison-Wesley, p. Section 9.4

Henning T., Michel B., Stognienko R., Dust opacities in dense regions, 1995, Planet. Space Sci., 43, 1333

Ino , 2004, 160x120 uncooled bolometer fpa datasheet, Technical report, Ino

Jarrett T. H., Chester T., Cutri R., Schneider S. E., Huchra J. P., The 2MASS Large Galaxy Atlas, 2003, AJ, 125, 525

Jeans J. H., 1928, Astronomy and cosmogony. Cambridge [Eng.] The University press

Jessop N. E., Ward-Thompson D., A far-infrared survey of molecular cloud cores, 2000, MNRAS, 311, 63

Jessop N. E., Ward-Thompson D., The initial conditions of isolated star formation - IV. C18O observations and modelling of the pre-stellar core L1689B, 2001, MNRAS, 323, 1025

Jones S., Naylor D., Gom B., Spencer L., 2009, in Bridging Excellence, The 30th Canadian Symposium on Remote Sensing, p. 556

Kirk J. M., Ward-Thompson D., André P., The initial conditions of isolated star formation - VI. SCUBA mappingof pre-stellar cores, 2005, MNRAS, 360, 1506

Kirk J. M., Ward-Thompson D., André P., The initial conditions of isolated star formation - VII. Spitzer mapping of pre-stellar cores, 2007, MNRAS, 375, 843

Ladjal D. et al., Herschel PACS and SPIRE imaging of CW Leonis, 2010, A&A, 518, L141

Larson R. B., Numerical calculations of the dynamics of collapsing proto-star, 1969, MNRAS, 145, 271

Lee C. W., Myers P. C., A Catalog of Optically Selected Cores, 1999, ApJS, 123, 233

Lee C. W., Myers P. C., Internal Motions in Starless Dense Cores, 2011, ApJ, 734, 60

Lee C. W., Myers P. C., Plume R., A Survey for Infall Motions toward Starless Cores. III. CS (3-2) and DCO^+ (2-1) Observations, 2004, ApJS, 153, 523

Lee C. W., Myers P. C., Tafalla M., A Survey of Infall Motions toward Starless Cores. I. CS (2-1) and N_2H^+ (1-0) Observations, 1999, ApJ, 526, 788

Lee C. W., Myers P. C., Tafalla M., A Survey for Infall Motions toward Starless Cores. II. CS (2-1) and N_2H^+ (1-0) Mapping Observations, 2001, ApJS, 136, 703

Loren R. B., The cobwebs of Ophiuchus. I - Strands of (C-13)O - The mass distribution, 1989, ApJ, 338, 902

Markwardt C. B., 2009, in Astronomical Society of the Pacific Conference Series, Vol. 411, D. A. Bohlender, D. Durand, & P. Dowler , ed, Astronomical Society of the Pacific Conference Series, p. 251

Mauron N., Huggins P. J., Multiple shells in the circumstellar envelope of IRC+10216, 1999, A&A, 349, 203

Meny C., Gromov V., Boudet N., Bernard J.-P., Paradis D., Nayral C., Far-infrared to millimeter astrophysical dust emission. I. A model based on physical properties of amorphous solids, 2007, A&A, 468, 171

Mertz L., 1965, Transformations in optics. Wiley Interscience

Mestel L., Spitzer L., Jr., Star formation in magnetic dust clouds, 1956, MNRAS, 116, 503

Mouschovias T. C., Magnetic braking, ambipolar diffusion, cloud cores, and star formation - Natural length scales and protostellar masses, 1991, ApJ, 373, 169

Mouschovias T. C., Spitzer L., Jr., Note on the collapse of magnetic interstellar clouds, 1976, ApJ, 210, 326

Myers P. C., Benson P. J., Dense cores in dark clouds. II - NH3 observations and star formation, 1983, ApJ, 266, 309

Myers P. C., Fuller G. A., Goodman A. A., Benson P. J., Dense cores in dark clouds. VI - Shapes, 1991, ApJ, 376, 561

Myers P. C., Ladd E. F., Fuller G. A., Thermal and nonthermal motions in dense cores, 1991, ApJ, 372, L95

Myers P. C., Linke R. A., Benson P. J., Dense cores in dark clouds. I - CO observations and column densities of high-extinction regions, 1983, ApJ, 264, 517

Nakano T., Quasistatic Contraction of Magnetic Protostars due to Magnetic Flux Leakage - Part One - Formulation and an Example, 1979, PASJ, 31, 697

Nakano T., Star Formation in Magnetic Clouds, 1998, ApJ, 494, 587

Naylor D., 2009, Control code for laboratory FTS as used by Lethbridge Instrument, Private Communication

Naylor D., 2010, Line fitting code integration into HIPE, produced by University of Lethbridge, Canada, Private Communication

Naylor M., D.A. Tahic, Apodizing functions for Fourier transform spectroscopy, 2007, J.Opt.Soc.Am, 24

Norton R., R.H. Beer, New Apodizing Functions for Fourier Spectrometry, 1976, J.Opt.Soc.Am, 66, 259

Nutter D., Ward-Thompson D., André P., A SCUBA survey of L1689 - the dog that didn't bark, 2006, MNRAS, 368, 1833

O'Halloran B. et al., Herschel photometric observations of the low metallicity dwarf galaxy NGC 1705, 2010, A&A, 518, L58

Ossenkopf V., Dust coagulation in dense molecular clouds: The formation of fluffy aggregates, 1993, A&A, 280, 617

Ossenkopf V., Henning T., Dust opacities for protostellar cores, 1994, A&A, 291, 943

Ott S., 2010, in Astronomical Society of the Pacific Conference Series, Vol. 434, Y. Mizumoto, K.-I. Morita, & M. Ohishi , ed, Astronomical Data Analysis Software and Systems XIX, p. 139

Panuzzo P. et al., Probing the molecular interstellar medium of M82 with Herschel-SPIRE spectroscopy, 2010a, A&A, 518, L37

Panuzzo P. et al., Probing the molecular interstellar medium of M82 with Herschel-SPIRE spectroscopy, 2010b, A&A, 518, L37

Panuzzo P. et al., Probing the molecular interstellar medium of M82 with Herschel-SPIRE spectroscopy, 2010c, A&A, 518, L37

Paradis D., Bernard J.-P., Mény C., Dust emissivity variations in the Milky Way, 2009, A&A, 506, 745

Pilbratt G. L., Riedinger J. R., Passvogel T., Crone G., Doyle D., Gageur U., Heras A. M., Jewell C., Metcalfe L., Ott S., Schmidt M., Herschel Space Observatory. An ESA facility for far-infrared and submillimetre astronomy, 2010, A&A, 518, L1

Plummer H. C., On the problem of distribution in globular star clusters, 1911, MNRAS, 71, 460

Pohlen M. et al., Radial distribution of gas and dust in spiral galaxies . The case of M 99 (NGC 4254) and M 100 (NGC 4321), 2010, A&A, 518, L72

Polehampton E., 2010, Processing SPIRE FTS Observations using the dark and observation as separate processed files, Private Communcation

Priddey R. S., McMahon R. G., The far-infrared-submillimetre spectral energy distribution of high-redshift quasars, 2001, MNRAS, 324, L17

Rangwala N. et al., Observations of Arp 220 using Herschel-SPIRE: An Unprecedented View of the Molecular Gas in an Extreme Star Formation Environment, 2011, ArXiv e-prints

Raymond G., Isaak K. G., Clements D., Rykala A., Pearson C., The Effectiveness of Mid IR / Far IR Blind, Wide Area, Spectral Surveys in Breaking the Confusion Limit, 2010, PASJ, 62, 697

Redman M. P., Keto E., Rawlings J. M. C., Williams D. A., Rotation of the pre-stellar core L1689B, 2004, MNRAS, 352, 1365

Redman M. P., Rawlings J. M. C., Nutter D. J., Ward-Thompson D., Williams D. A., Molecular gas freeze-out in the pre-stellar core L1689B, 2002, MNRAS, 337, L17

Roussel H. et al., SPIRE imaging of M 82: Cool dust in the wind and tidal streams, 2010, A&A, 518, L66

Sacchi N. et al., 2011, in IAU Symposium, Vol. 280, IAU Symposium, p. 322P

Sakai H., Vanasse G. A., Forman M. L., Spectral Recovery in Fourier Spectroscopy, 1968, Journal of the Optical Society of America (1917-1983), 58, 84

Sauvage M. et al., The central region of spiral galaxies as seen by Herschel. M 81, M 99, and M 100, 2010, A&A, 518, L64

Savini G., 2009, Work on categorising fringing functions for SPIRE high resolution observations, Private Communcation

Shu F. H., Self-similar collapse of isothermal spheres and star formation, 1977, ApJ, 214, 488

Shu F. H., Adams F. C., Lizano S., Star formation in molecular clouds - Observation and theory, 1987, ARA&A, 25, 23

Simpson R. J., Johnstone D., Nutter D., Ward-Thompson D., Whitworth A. P., The initial conditions of isolated star formation - X. A suggested evolutionary diagram for prestellar cores, 2011, ArXiv e-prints

Simpson R. J., Nutter D., Ward-Thompson D., The initial conditions of star formation - VIII. An observational study of the Ophiuchus cloud L1688 and implications for the pre-stellar core mass function, 2008, MNRAS, 391, 205

Skrutskie M. F. et al., The Two Micron All Sky Survey (2MASS), 2006, AJ, 131, 1163

Sohn J., Lee C. W., Park Y.-S., Lee H. M., Myers P. C., Lee Y., Probing Inward Motions in Starless Cores Using the HCN(J = 1-0) Hyperfine Transitions: A Pointing Survey toward Central Regions, 2007, ApJ, 664, 928

Stamatellos D., Whitworth A. P., Ward-Thompson D., The dust temperatures of the pre-stellar cores in the ρ Oph main cloud and in other star-forming regions: consequences for the core mass function, 2007, MNRAS, 379, 1390

Stepnik B., Abergel A., Bernard J.-P., Boulanger F., Cambrésy L., Giard M., Jones A. P., Lagache G., Lamarre J.-M., Meny C., Pajot F., Le Peintre F., Ristorcelli I., Serra G., Torre J.-P., Evolution of dust properties in an interstellar filament, 2003, A&A, 398, 551

Swinyard B. M. et al., In-flight calibration of the Herschel-SPIRE instrument, 2010, A&A, 518, L4

van der Werf P. P. et al., Black hole accretion and star formation as drivers of gas excitation and chemistry in Markarian 231, 2010a, A&A, 518, L42

van der Werf P. P. et al., Black hole accretion and star formation as drivers of gas excitation and chemistry in Markarian 231, 2010b, A&A, 518, L42

van der Werf P. P. et al., Black hole accretion and star formation as drivers of gas excitation and chemistry in Markarian 231, 2010c, A&A, 518, L42

Walmsley C. M., Flower D. R., Pineau des Forêts G., Complete depletion in prestellar cores, 2004, A&A, 418, 1035

Walters A. D., Klotz A., Caux E., Crovisier J., 2005, in IAU Symposium, Vol. 235, IAU Symposium, p. 155P

Ward-Thompson D., André P., Crutcher R., Johnstone D., Onishi T., Wilson C., An Observational Perspective of Low-Mass Dense Cores II: Evolution Toward the Initial Mass Function, 2007, Protostars and Planets V, 33

Ward-Thompson D., André P., Kirk J. M., The initial conditions of isolated star formation - V. ISOPHOT imaging and the temperature and energy balance of pre-stellar cores, 2002, MNRAS, 329, 257

Ward-Thompson D. et al., A Herschel study of the properties of starless cores in the Polaris Flare dark cloud region using PACS and SPIRE, 2010, A&A, 518, L92

Ward-Thompson D., Motte F., Andre P., The initial conditions of isolated star formation - III. Millimetre continuum mapping of pre-stellar cores, 1999, MNRAS, 305, 143

Ward-Thompson D., Scott P. F., Hills R. E., Andre P., A Submillimetre Continuum Survey of Pre Protostellar Cores, 1994, MNRAS, 268, 276

Whitworth A. P., Ward-Thompson D., An Empirical Model for Protostellar Collapse, 2001, ApJ, 547, 317

Whitworth P. A., 2006a, PX3212 - course notes for Stellar Physics, Cardiff University Course Notes

Whitworth P. A., 2006b, PX4106 - course notes for The Interstellar Medium, Star Formation and Cosmic Gas Dynamics, Cardiff University Course Notes

Wood D. O. S., Myers P. C., Daugherty D. A., IRAS images of nearby dark clouds, 1994, ApJS, 95, 457