ROM SAF Report 30
Ref: SAF/ROM/METO/REP/RSR/030
Web: www.romsaf.org
Date: 27 March 2018

EUMETSAT
ROM SAF
RADIO OCCULTATION METEOROLOGY

ROM SAF Report 30

# A first look at the feasibility of assimilating single and dual frequency bending angles

I D Culverwell[1] and S B Healy[2]

[1]Met Office, Exeter, UK
[2]ECMWF, Reading, UK

### Document Author Table

| | ***Name*** | ***Function*** | ***Date*** | ***Comments*** |
|---|---|---|---|---|
| ***Prepared by:*** | I D Culverwell | ROM SAF Project Team | 27 March 2018 | |
| ***Reviewed by:*** | M Forsythe, J Eyre | Met Office review | 6 Mar 2018 | |
| ***Reviewed by:*** | C Piccolo | Met Office approval | 14 Mar 2018 | |
| ***Approved by:*** | K B Lauritsen | ROM SAF Project Manager | 26 Mar 2018 | |

### Document Change Record

| ***Issue/Revision*** | ***Date*** | ***By*** | ***Description*** |
|---|---|---|---|
| 0.1 | 25 Oct 2017 | IDC | 1st draft |
| 0.2 | 21 Dec 2017 | IDC | 2nd draft |
| 0.3 | 16 Mar 2018 | IDC | 3rd draft, after review |
| 1.0 | 27 Mar 2018 | IDC | Final version, following approval |

### ROM SAF

The Radio Occultation Meteorology Satellite Application Facility (ROM SAF) is a decentralised processing centre under EUMETSAT which is responsible for operational processing of GRAS radio occultation (RO) data from the Metop satellites and RO data from other missions. The ROM SAF delivers bending angle, refractivity, temperature, pressure, and humidity profiles in near-real time and offline for NWP and climate users. The offline profiles are further processed into climate products consisting of gridded monthly zonal means of bending angle, refractivity, temperature, humidity, and geopotential heights together with error descriptions.

The ROM SAF also maintains the Radio Occultation Processing Package (ROPP) which contains software modules that will aid users wishing to process, quality-control and assimilate radio occultation data from any radio occultation mission into NWP and other models.

The ROM SAF Leading Entity is the Danish Meteorological Institute (DMI), with Cooperating Entities: i) European Centre for Medium-Range Weather Forecasts (ECMWF) in Reading, United Kingdom, ii) Institut D'Estudis Espacials de Catalunya (IEEC) in Barcelona, Spain, and iii) Met Office in Exeter, United Kingdom. To get access to our products or to read more about the project please go to: http://www.romsaf.org

### Intellectual Property Rights

**Abstract**

The report describes the results of some numerical experiments designed to investigate the feasibility of carrying out 1D variational retrievals using dual or single frequency radio occultation bending angles, rather than the usual ionospherically corrected 'neutral' bending angles. A detailed analysis of the results of carrying out these experiments on two example radio occultation profiles is followed by a statistical analysis of the results of doing so on 3500 consecutive profiles observed by the GRAS instrument on Metop-A. The retrieved ionospheric parameters are briefly reviewed. Further work is suggested.

# Contents

# 1 Introduction

This report is a preliminary study into the feasibility of using both RO bending angle profiles — 'L1', at $f_1$ = 1.57542 GHz, and 'L2', at $f_2$ = 1.22760 GHz — in NWP data assimilation. In operational variational retrievals, which assimilate ionosphere-free, 'neutral' bending angles, a linear combination ('LC') of the two bending angles is taken, the weights of which are chosen to eliminate the ionospheric contribution to the bending (see [12]).

The L2 bending angle, however, frequently drops out somewhere below about 20 km, and this prevents LC from being constructed. Sometimes this problem is solved by extrapolating the L1–L2 difference from above, but the best way to do this is open to question. It is also possible that useful information about the (mean) state of the ionosphere along the ray path(s) is being thrown away by this procedure. Equally, information contained in the L1 signal at lower altitudes may be lost for want of an L2 counterpart at these levels. There are also occasions where the L2 signal is present but not readily usable. For example, during most of the GPS/MET feasibility experiment ([10]) 'anti-spoofing' was turned on, which made the L2 signal too noisy to be processed by the usual means. The L1 signal, however, was unaffected by the anti-spoofing, and therefore if this alone could be used in retrievals then many more of the GPS/MET profiles might be used — for example, in reanalyses.

It is instructive, therefore, to consider the possibility of carrying out variational retrievals using single or dual frequency bending angles instead of LC. This requires the ionospheric contribution to the ray bending to be included in the forward model. NWP centres typically assimilate bending angles below 80 km, and there is unlikely to be enough information in them to say much about the ionosphere (the bulk of which — the F2-layer — is usually sited around 300 km above the ground) such as the electron density profile or its horizontal gradients. We have therefore assumed the ionospheric bending to be caused by an idealised, spherically symmetric single 'Chapman layer' ionosphere, defined by its height, width and peak electron density. These three parameters are included in the 'state vector', on which the forward model acts, and can therefore be adjusted during the 1D variational retrieval, in the same way as the temperature and the humidity (which only affect the neutral bending angles), to produce overall (neutral plus ionospheric) bending angles which can be fitted to the observed L1 and L2.

It should be noted that the model of ionospheric bending employed in this report has been used by EUMETSAT to extrapolate L1–L2 bending angles where the latter disappear ([8]). It has also led to promising developments in the reduction of the so-called 'residual' error that remains after the usual ionospheric correction ([3], [5]). A simplified version of the bending angle formula has also found use in the extraction of bending angle data from excess phase data detected by the GNOS instrument carried by the FY-3C satellite ([7]). Its proposed use here is therefore consistent with some existing operational and research practice.

The retrieved ionospheric parameters, and the implied electron density profile derived from them, also allow some inferences about the ionosphere to be drawn, and these are briefly discussed. It should be stressed, however, that this is not the principal aim of this work, which is to examine the feasibility of using single and dual frequency bending angles, rather than ionospherically corrected ones, in variational retrievals of meteorological fields of the neutral atmosphere.

## 2 Theory

Most of the theory is covered in an earlier ROM SAF Report 17 ([1]). The conclusion is that the L1 and L2 bending angles at impact parameter $a$ can be written as

$$\alpha_{Lj}(a) = \alpha_{LC}(a) + \alpha_{LIj}(a) \quad (j = 1,2) \tag{2.1}$$

where $\alpha_{LC}$ is the usual neutral bending angle, given by the Abel transform of the neutral pressure-, temperature- and humidity-dependent refractivity $N$ (eg [6]), and $\alpha_{LIj}$ is the frequency-dependent bending angle induced by a spherically symmetric Chapman layer ionosphere. If this is taken to be centred at $r_0$ and to have a width $H$ and a peak electron density $n_e^{\max}$, so that its electron density distribution is given by

$$n_e(r) = n_e^{\max} \exp\left((1-u-e^{-u})/2\right), \quad \text{where } u = (r-r_0)/H, \tag{2.2}$$

then RSR 17 shows that $\alpha_{LIj}(a)$ then takes the form

$$\alpha_{LIj}(a) = (k_4/f_j^2) n_e^{\max} [4e r_0^2 a^2 / H(r_0+a)^3]^{1/2} Z(l), \tag{2.3}$$

where

$$Z(l) = \int_{-l}^{\infty} \frac{(e^{-u}-1)\exp[(1/2)(1-u-e^{-u})]}{\sqrt{u+l}}\, du \tag{2.4}$$

$$= 2\sqrt{eg}\int_0^{\infty} \left(ge^{-w^2}-1\right) \exp\left[-\frac{1}{2}(w^2+ge^{-w^2})\right] dw, \quad \text{defining } g = e^l. \tag{2.5}$$

In Eqn (2.3) $k_4 = q_e^2/8\pi^2 m_e \varepsilon_0 = 40.3$ m$^3$ s$^{-2}$ is a universal constant related to the plasma frequency of the electrons, and $f_j$ is the frequency of the Lj signal. $Z(l)$ is a dimensionless function of the dimensionless distance below the electron density peak height $l = (r_0 - a)/H$ ([1]). It is of order unity in magnitude, and has a characteristic vertically reflected Z-shape.

RSR 17 ([1]) contains many approximations to $Z(l)$, but in practice a Padé approximation

$$Z(l) = \sum_{i=0}^{3} p_i \theta^i / \sum_{i=0}^{5} q_i \theta^i, \text{ where } \theta = \sinh^{-1}(e^l/2) \tag{2.6}$$

is numerically robust, easy to encode and accurate to within 2.2% for all $l$. (The coefficients $\{p_i\}$ and $\{q_i\}$ are given in Appendix A of RSR 17.)

This simple ionospheric forward model was incorporated in the forward model and 1dvar retrieval modules of ROPP8.0, which was released in February 2014. Full descriptions of the algorithms and examples of their use can be found in the ROPP User Guides ([11]). The relevant features here are:

- The state vector, which is the input to the forward model, is augmented by the ionospheric parameters:
$$\mathbf{x} = \{T, q, p*\} \mapsto \{T, q, p*, n_e^{\max}, r_0, H\} \tag{2.7}$$
- The observation vector, against which the forward modelled state vector is compared in the retrieval algorithm, is doubled in length:
$$\mathbf{y} = \{\alpha_{LC}\} \mapsto \{\alpha_{L1}, \alpha_{L2}\} \tag{2.8}$$

- The covariance matrices $\mathbf{B}$ and $\mathbf{R}$ that appear in the cost function $J(\mathbf{x}) = (1/2)(\mathbf{x}-\mathbf{b})^T\mathbf{B^{-1}}(\mathbf{x}-\mathbf{b}) + (1/2)(\mathbf{H}(\mathbf{x})-\mathbf{y})^T\mathbf{R^{-1}}(\mathbf{H}(\mathbf{x})-\mathbf{y})$, which is minimised in the retrieval algorithm, therefore need updating. How best to do this is open to question. The results in this report were calculated by assuming:
  - The errors in $\{n_e^{\max}, r_0, H\}$ are uncorrelated from each other and from those of $\{T, q, p*\}$;
  - $\sigma(n_e^{\max}) = 2\times10^{11}$ m$^{-3}$; $\sigma(r_0) = 150$ km; $\sigma(H) = 25$ km;
  - b$(n_e^{\max})$ = profile-dependent; b$(r_0) = 300$ km; b$(H) = 75$ km;
  - The errors in $\alpha_{L1}$ are uncorrelated from those of $\alpha_{L2}$;
  - $\sigma(\alpha_{L1}) = \max(\alpha_{LC}, 10\ \mu\text{rad})$; $\sigma(\alpha_{L2}) = \max(\alpha_{LC}, 30\ \mu\text{rad})$.

There are grounds for using $\sigma(\alpha_{L2}) = 3\sigma(\alpha_{L1})$ for GRAS data, in which the L2/P carrier phase noise is around three times larger than that of L1/CA ([4]). And in the neutral atmosphere, where $\alpha_{L1} = \alpha_{L2} = \alpha_{LC}$, $\sigma_{L1}$ should equal $\sigma_{L2}$.

The profile-dependent estimate for the first guess $n_e^{\max}$ is described in the ROPP_1DVAR User Guide ([11]). It amounts to making b$(n_e^{\max})$ proportional to the average difference of $\alpha_{L2} - \alpha_{L1}$ above 30 km.[1] Without this first guess, the number of iterations that the minimiser needs to converge, and the final cost function, are often much higher. We currently set the background $n_e^{\max}$ equal to this first guess. This is in fact not valid, because the background should be independent of the observations. The costs of rewriting ROPP to allow different backgrounds and first guesses, however, led us to neglect this inconsistency in a first investigation of the problem. Further consequences of this choice will be discussed in the report and the summary.

The choices described above deliver results whose quality and usefulness readers will be able to judge for themselves. They do at least allow retrievals to be made on a high percentage of GRAS RO profiles. But they should not be considered the final word.

[1] Specifically, b$(n_e^{\max}) = K\langle\alpha_{L2} - \alpha_{L1}\rangle$, where $K = 2.4\times10^{16}$ m$^{-3}$rad$^{-1}$.

# 3 Results

We examine the results of passing 247 level GRAS bending angles and colocated 91 level ECMWF background fields through the 1D–Var bending angle retrieval component of ROPP, using the parameters listed in the previous section.

## 3.1 Single profiles

The profiles fall into two categories, depending on whether $\alpha_{L2}$ is greater or less than $\alpha_{L1}$ at high altitude (say 50 km). These two categories are called 'regular' and 'anomalous' respectively, because in a spherically symmetric ionosphere, for which Eqn (2.1) is valid, $\alpha_{L2}(a)$ should exceed $\alpha_{L1}(a)$. Over 85% of cases are regular.

### 3.1.1 A regular profile

Fig 3.1 shows the results of carrying out a 1D–Var retrieval using the LC bending angles calculated for a regular occultation (`atm20120611_101706_M02_2150454543_N0019_XXXX.nc`). Not unusually, the L2 bending angle drops out below 15 km, and therefore so does LC, even though there is some useful tropospheric information in L1 below this height. L2–L1 is roughly constant at $\sim$30 $\mu$rad (positive because this is a regular occultation). Temperature increments of $\sim \pm 1$ K arise from the retrieval procedure. Tropospheric specific humidity increments of $\sim -0.1$ g/kg are found. The cost function falls to about two thirds of its initial value, until $2J/m$ is about 2.8, which is not particularly small.

Fig 3.2 shows the results of carrying out a 1D–Var retrieval using the L1 and L2 bending angles from the same occultation. The forward modelled L1 and L2 bending angles, which extend through the troposphere, fit the observed single frequency bending angles reasonably well, and also, therefore, their difference. The temperature increments are similar in size to those resulting from the LC retrieval. There are bigger differences to the humidity, which is not surprising since LC does not extend into the troposphere. Note that this procedure delivers a retrieved ionospheric electron density 'profile' (really, just the three parameters that specify it). The analysed total electron content (TEC) is 55.6 TEC, which is reasonable for a midlatitude, midsummer, night-time profile. The overall fit is not as good as it was for the LC retrieval: the final $2J/m$ is about 3.1. This is of course sensitive to the highly uncertain errors that were used in the retrieval (see Sec 2). Closer examination of this particular case reveals that a large part of the total cost function comes from a small 'jump' in the bending angles between 15 and 18 km, which is just about visible in Fig 3.2.

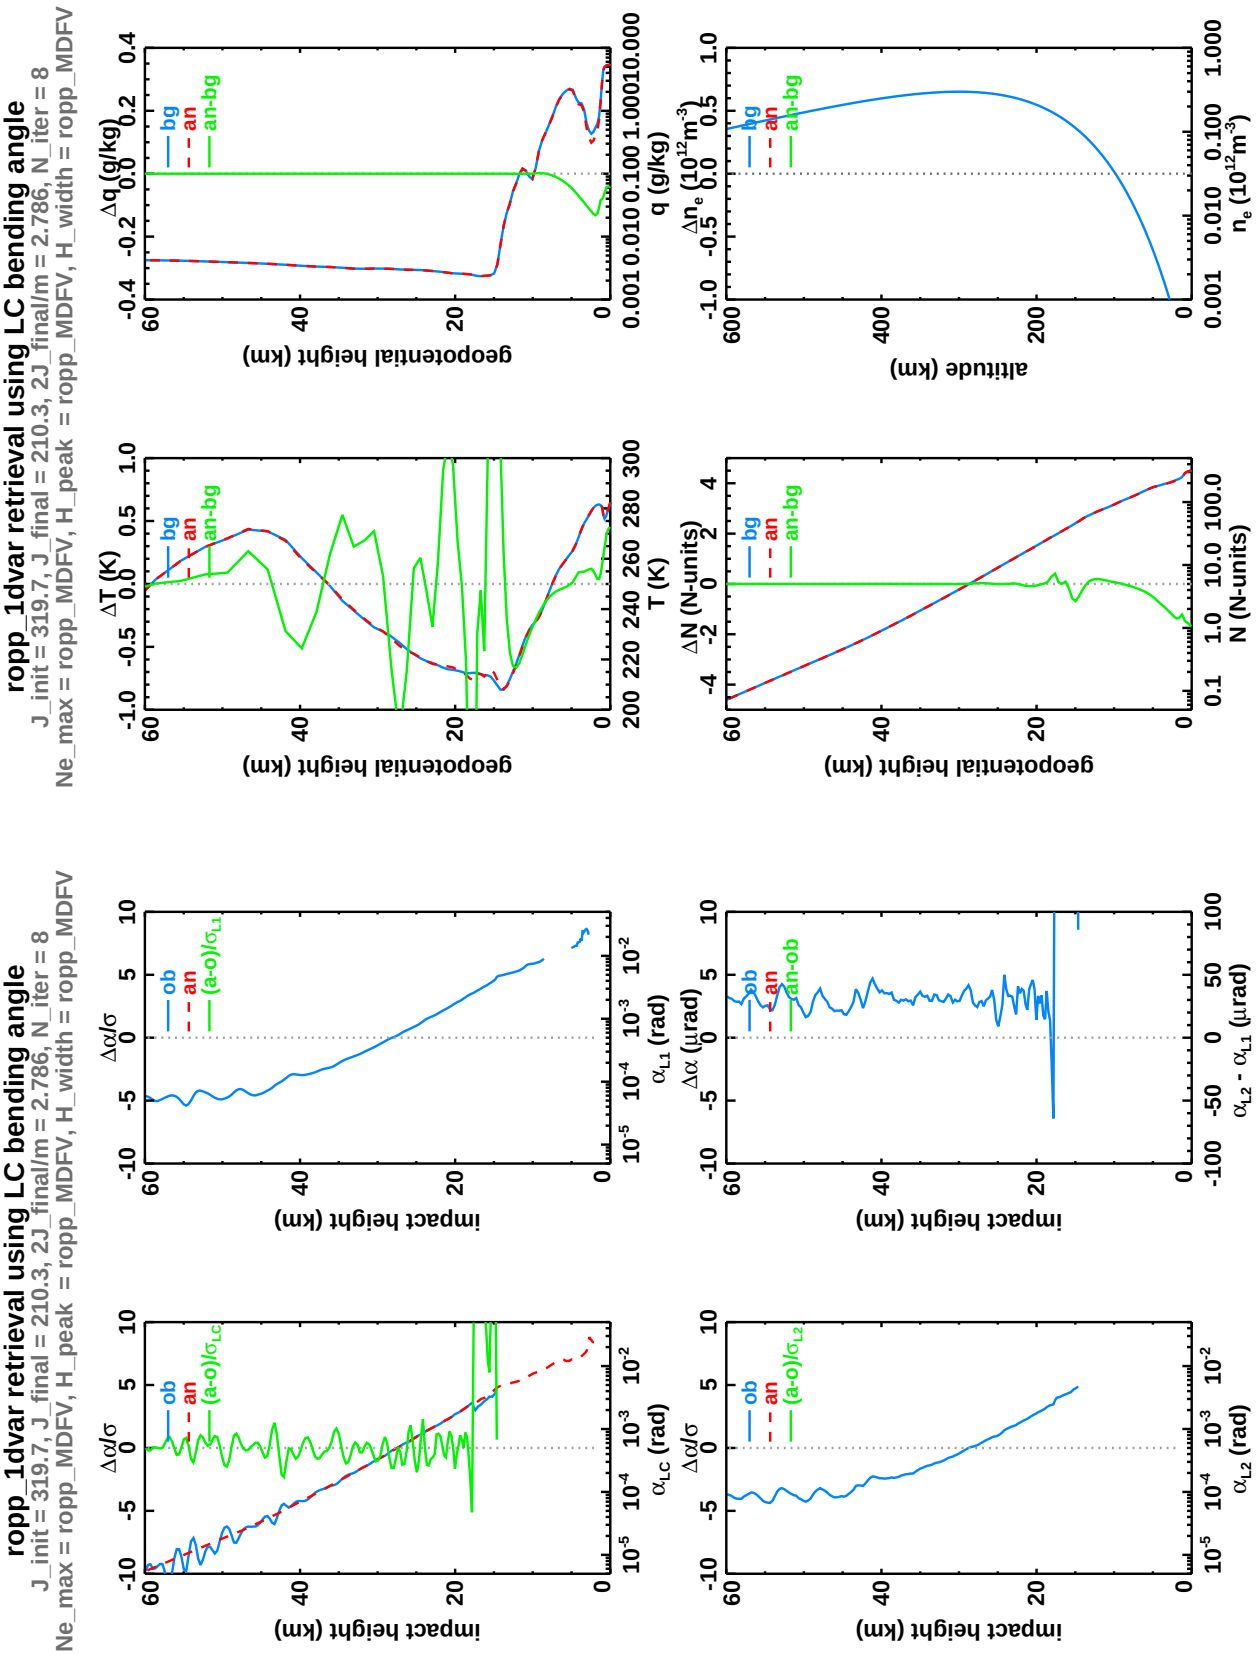


**Figure 3.1**: 1D–Var retrieval of a regular GRAS profile using LC bending angles (current operational procedure). Left hand panel: the observed, analysed and departures for LC, L1, L2 and L2–L1 bending angles. Right hand panel: the background, analysis and increments of $T$, $q$, $N$ and $n_e$.

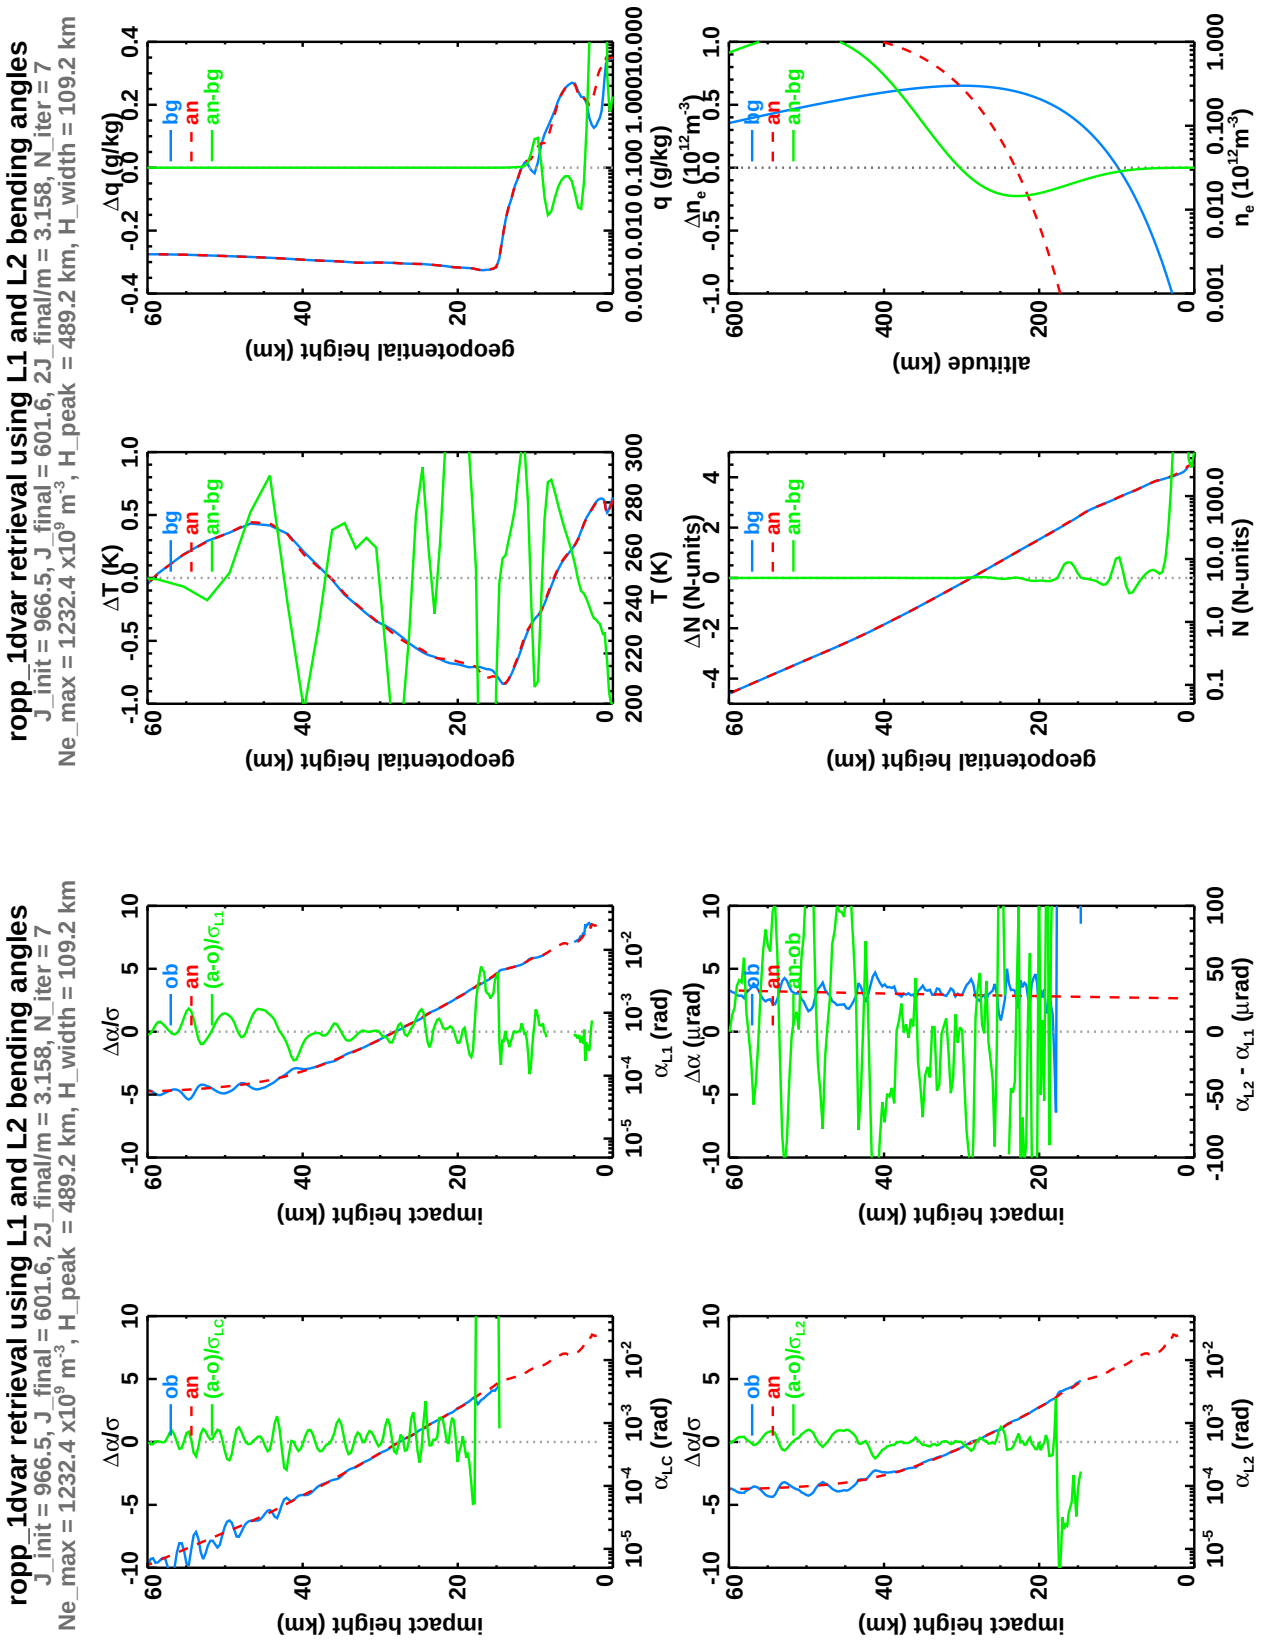


**Figure 3.2**: As for Fig 3.1, but using L1 and L2 bending angles, and direct modelling of the ionosphere in the forward model.

The next experiment, an extreme case of the second, is one in which L2 is completely absent. In this case it would be impossible to construct LC anywhere, and traditional data assimilation would therefore fail. But if we forward model, however crudely, the ionospheric bending, then a retrieval can still be attempted. Note that the absence of L2–L1 anywhere means that the first guess $n_e^{\max}$ cannot be tailored to the individual profile, as discussed in Sec 2, but must simply default to $3.0 \times 10^{11}$ m$^{-3}$.

The results of doing this on the same profile as before (but with L2 nullified) are shown in Fig 3.3. If anything the results are slightly better (in this example) than are obtained from using L1 and L2: the L1 O–B's, and the increments to the state vector variables, are about the same size as before. The final $2J/m$ of 1.4 is smaller, even though $m$ — the number of (valid) observations — is about half as big (218 c.f. 381). This improved fit is largely because the jump in the L1 bending angle between 15 and 18 km, referred to above, is fitted better now. The retrieved electron density is closer to the background, presumably because the initial value $(3.0 \times 10^{11}\text{m}^{-3})$ is very different from that of the L1/L2 case $(7.3 \times 10^{11}\text{m}^{-3})$.

The final experiment gets around the problem of missing L2 data by simply using L1 as a proxy for LC. In other words, the ionospheric component of L1 is ignored. This is unlikely to be a good approximation above about 40 km, but if it were to work it would be a simple solution, and indeed the only possible way to use just the L1 data without an ionospheric model. To get this experiment to succeed at all it was found necessary to increase the minimum $\sigma(\alpha_{LC})$, which applies above about 25 km, from 6 $\mu$rad to 10 $\mu$rad. Without this change, the initial cost function is too large to allow the minimiser to make any progress.

The results are shown in Fig 3.4. As expected, the L1 departures are very large at altitude. The increments in $T$, $q$ and $N$ are, however, similar to those in the other experiments. But the small (20%) reduction in the initial cost function, to a final value of $2J/m$ of 13.3, shows that not much fitting has taken place.

The statistics of a large number of retrievals run through these experiments will be studied in Sec 3.2. Before doing so we briefly turn to a study of the analagous results for a single 'anomalous' profile.

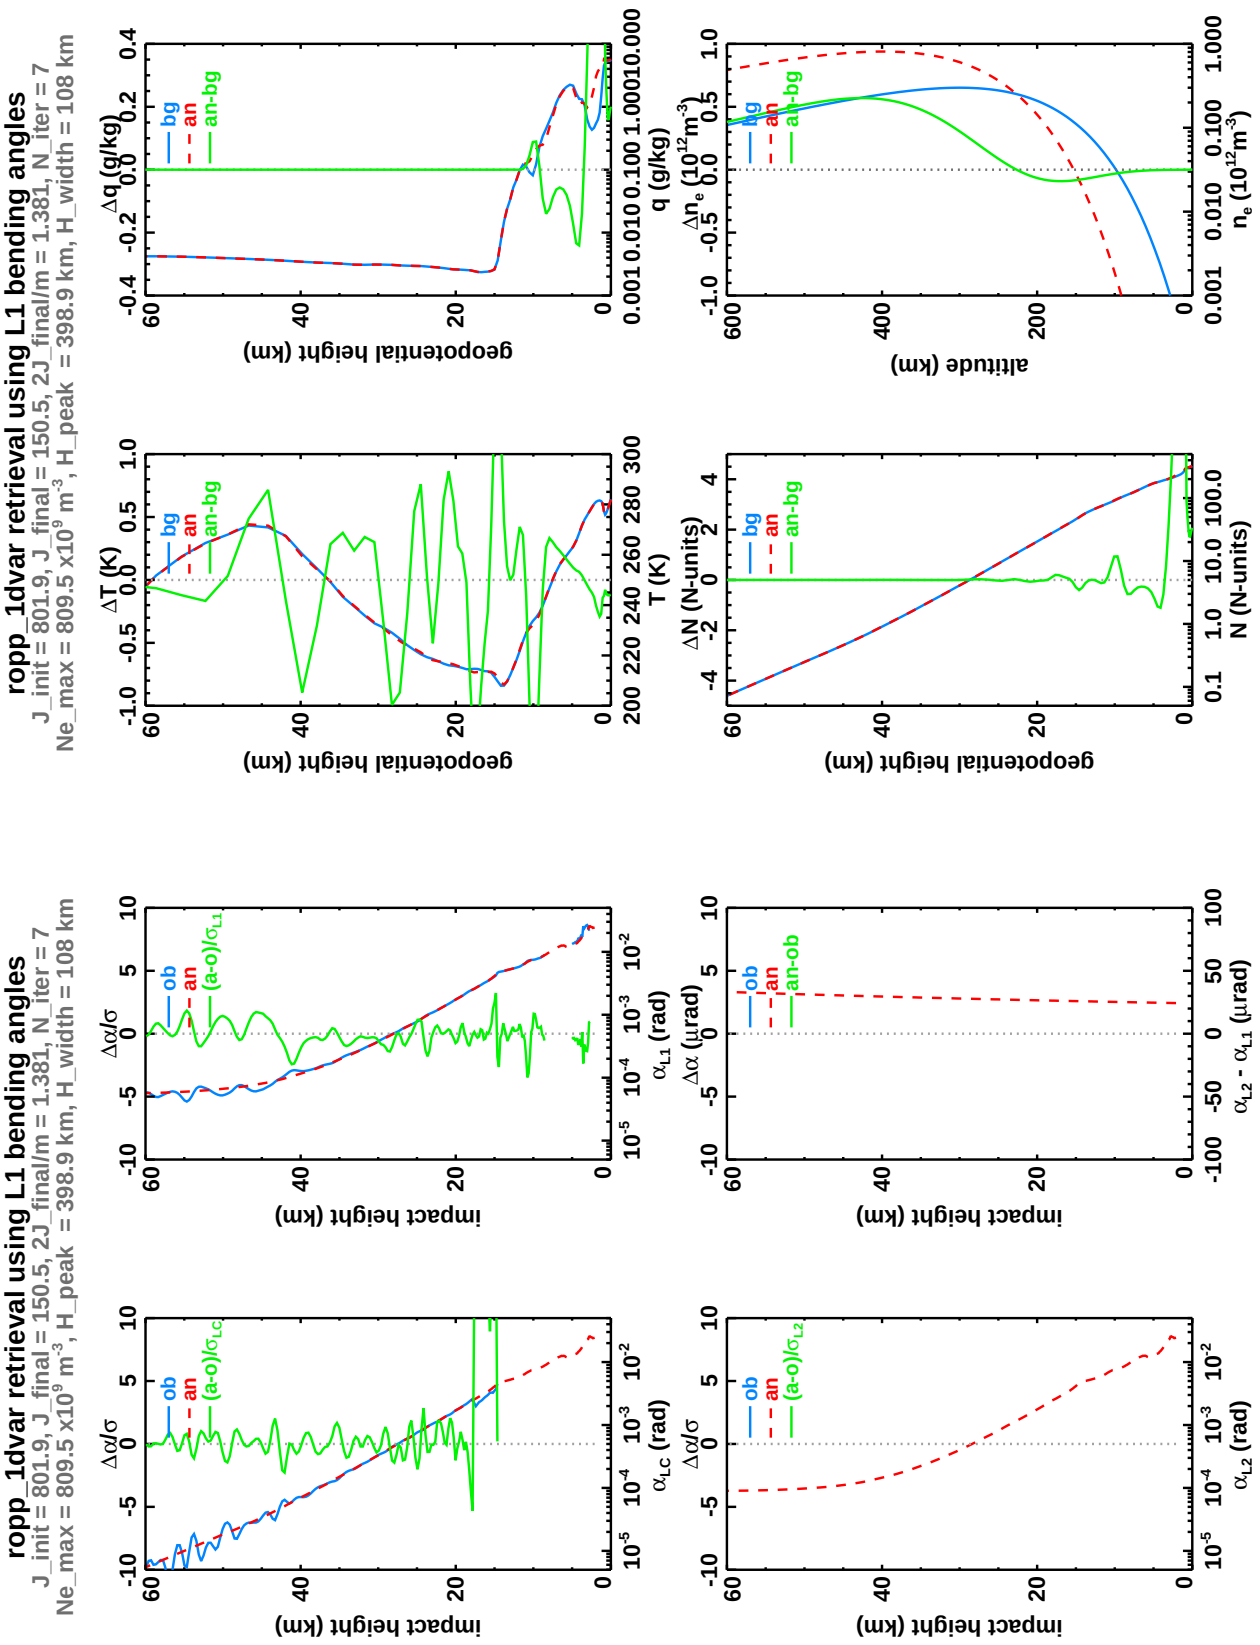


**Figure 3.3**: As for Fig 3.2, but using L1 bending angles only.

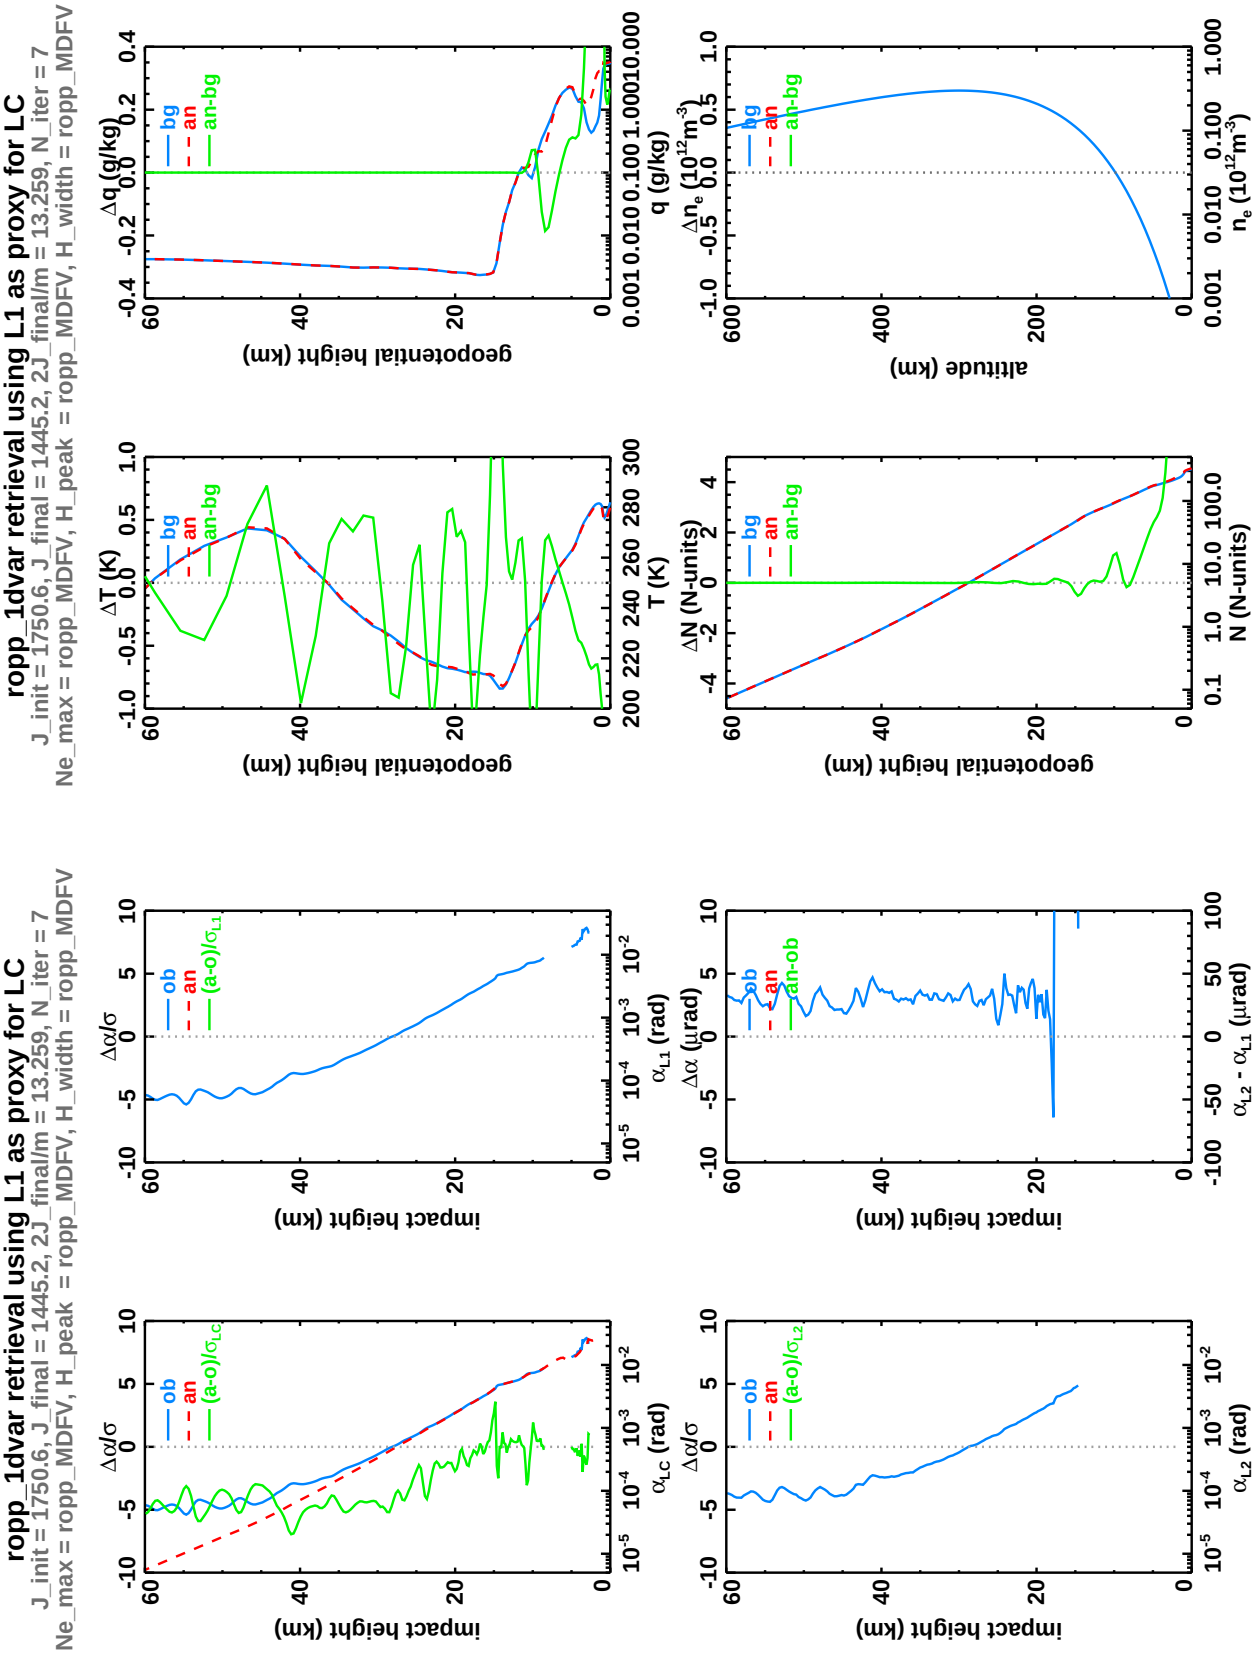


**Figure 3.4**: As for Fig 3.1, but using L1 as a proxy for LC.

### 3.1.2 An anomalous profile

Fig 3.5 shows the results of carrying out a 1D–Var retrieval using the LC bending angles calculated from an anomalous occultation (`atm20120611_095522_M02_2060454541_N0019_XXXX.nc`). In this case L2 drops out at an impact height of 10 km, and L1 at an impact height of about 3 km (which is around 1 km above the surface). Note that, despite the negative L1 and L2 bending angles above about 40 km, as is characteristic of an anomalous occultation, the ionospherically corrected LC bending angle looks reasonable, and appears to produce a reasonable retrieval: temperature increments of about $\pm$ 1 K, and very small humidity increments — perhaps because LC is only available above 10 km.

Fig 3.6 shows the results of carrying out a 1D–Var retrieval using L1 and L2 bending angles and the simple 3-parameter ionospheric model described in Sec 2. Good fits to L1, L2 and (therefore) L2–L1 are found, and the retrieved temperature and humidity look reasonable. Increments to the latter are larger in this case, perhaps because the tropospheric (L1) data can now be used in the retrieval. A similar distinction applies to the regular case (see Figs 3.1 and 3.2).

Eqns (2.1) and (2.3) show that

$$\alpha_{L2}(a) - \alpha_{L1}(a) \propto (f_2^{-2} - f_1^{-2})\, n_e^{\max} Z\left(\frac{r_0 - a}{H}\right), \tag{3.1}$$

which implies, since $f_2 < f_1$, that under the assumption of a spherically symmetric ionosphere the only way to get anomalous occultations (i.e. those having $\alpha_{L2} < \alpha_{L1}$) is for $Z(l)$ or $n_e^{\max}$ to be negative. The former can occur above the electron density peak ([1]), which is not physically feasible. Neither, of course, is the latter. But both possibilities are formally allowed in the retrieval scheme that has been incorporated into ROPP, and in this example it converges on a solution with negative $n_e^{\max}$. Despite the absurdity of this, the neutral atmosphere retrievals look fair enough, and the final $2J/m$ of 0.3 suggests convergence.

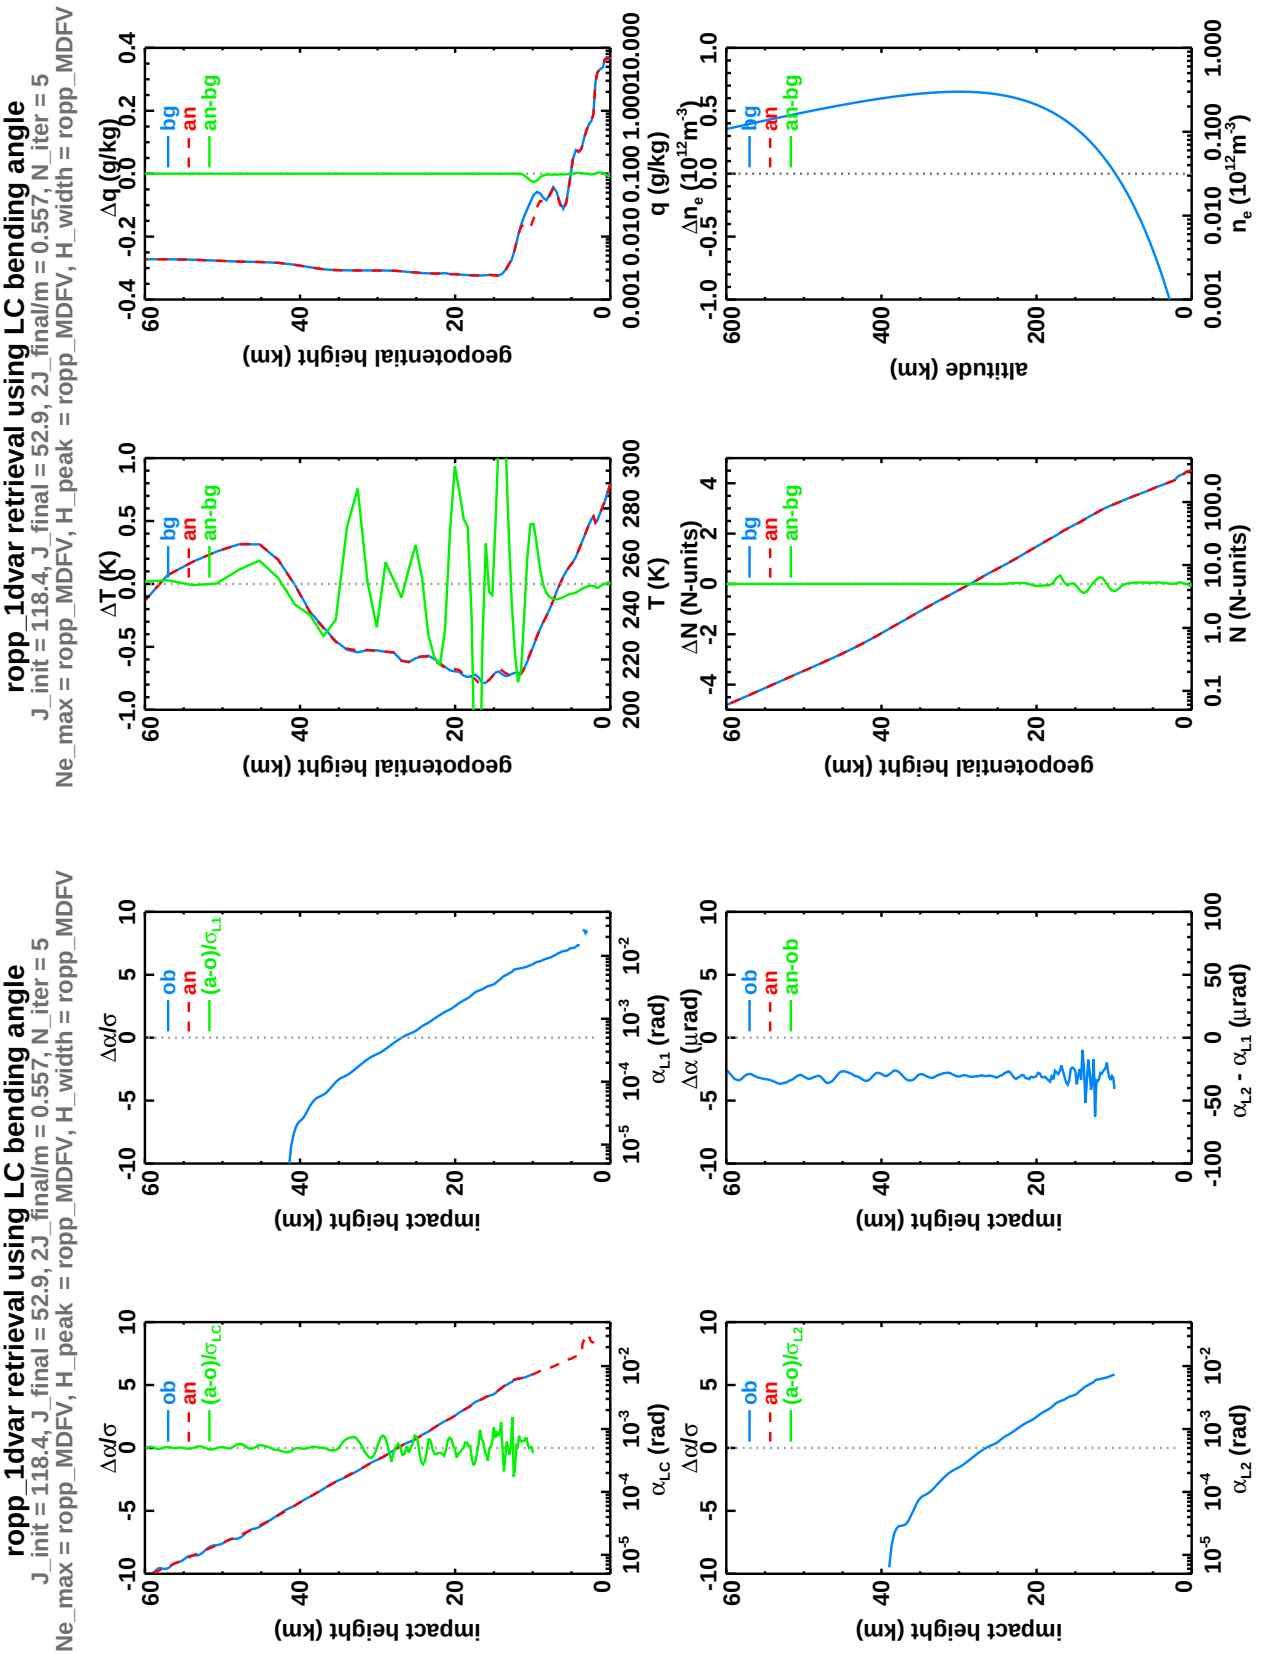


**Figure 3.5**: 1D–Var retrieval of an anomalous GRAS profile using LC bending angles (current operational procedure). Left hand panel: the observed, analysed and departures for LC, L1, L2 and L2–L1 bending angles. Right hand panel: the background, analysis and increments of $T$, $q$, $N$ and $n_e$.

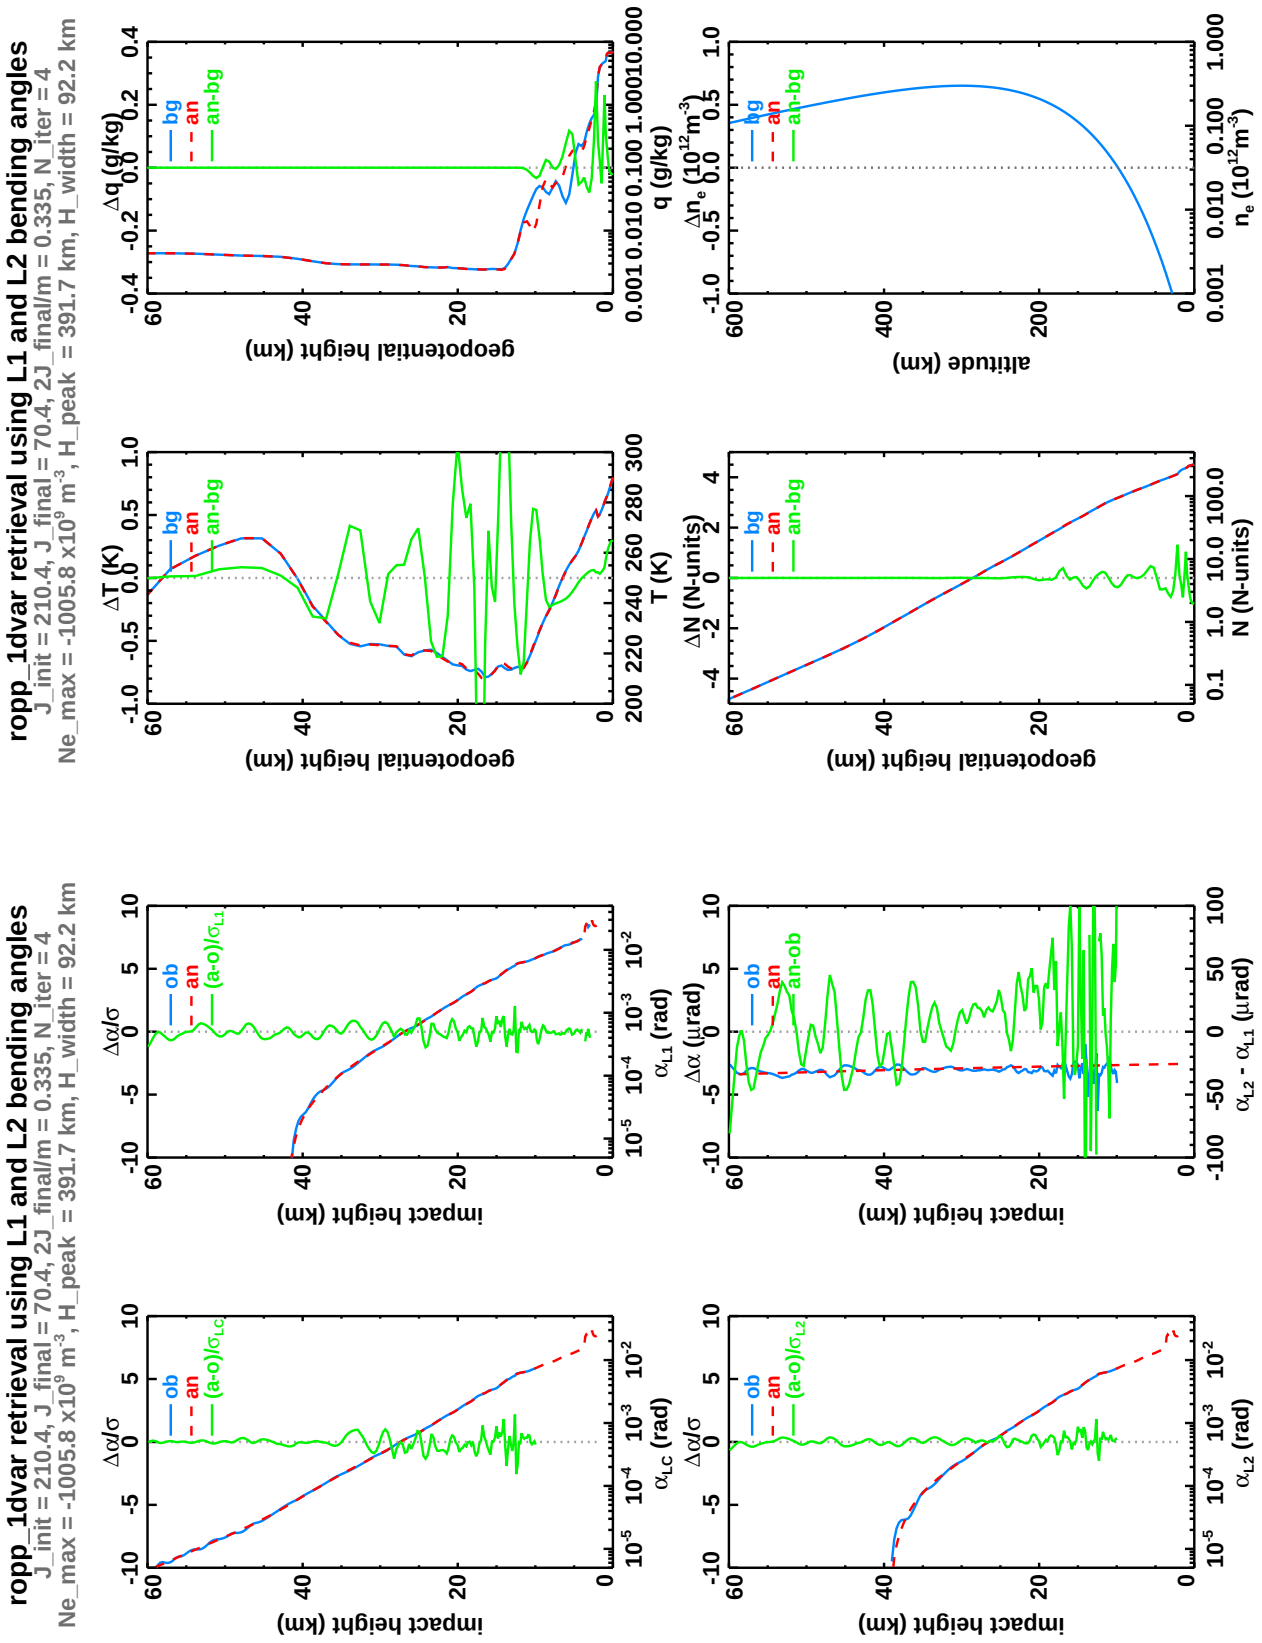


**Figure 3.6**: As for Fig 3.5, but using L1 and L2 bending angles, and direct modelling of the ionosphere in the forward model.

The results of the corresponding experiment in which L2 is fully nullified are shown in Fig 3.7. This is less successful: a poor fit to L1 is obtained, and the final $2J/m$ is 12.9. Detailed examination of the output shows that the solver found an ionosphere peak height $r_0$ that was less than 1% of the background (of 300 km), at which point it reset $r_0$ to that background value, the cost function became very large, and the solver gave up. In other words, it didn't converge at all. This contrasts with the analagous experiment with the 'regular' profile, for which the L1-only retrieval was, if anything, slightly better than that which used L1 and L2.

It is possible in this particular case to find a better solution by simply setting the background $n_e^{\max}$ to be $-3.0 \times 10^{11}$ m$^{-3}$ rather than $3.0 \times 10^{11}$ m$^{-3}$. The results are shown in Fig 3.8. The L1 O–Bs are smaller, as are the temperature increments above 35 km. The final normalised cost function, $2J/m$, is 0.5. This clearly illustrates that there is more work to do in refining the initial estimates of the background ionospheric parameters, and probably their errors too, before retrievals based on L1 and L2 can be considered a reliable alternative to using LC.

The results of the final experiment, in which L1 is used as a proxy for LC and the ionosphere is omitted from the forward model, are shown in Fig 3.9. (This experiment also applies the increase in the minimum $\sigma(\alpha_{LC})$ to 10 $\mu$rad, as discussed in Sec 3.1.1.) The fit to the observed LC is reasonable below 20 km, but obviously goes awry as the ionospheric contribution starts to dominate. As for the L1-only retrieval, the minimiser stopped because the cost function started to increase, which means it didn't really converge to a solution, as the final $2J/m$ value of 12.1 confirms.

In summary, reasonable ionospherically uncorrected 1D–Var retrievals using L1 and L2 bending angles, or just the L1 ones, can be obtained for regular occultations, which form the majority of cases. For the remaining anomalous occultations, ionospherically uncorrected 1D–Var retrievals are possible, at the expense of nonsensical implied electron density profiles. There is a suggestion that the retrievals in these cases are more sensitive to the choice of the background ionospheric parameters. Certainly, this aspect of the problem needs further examination.

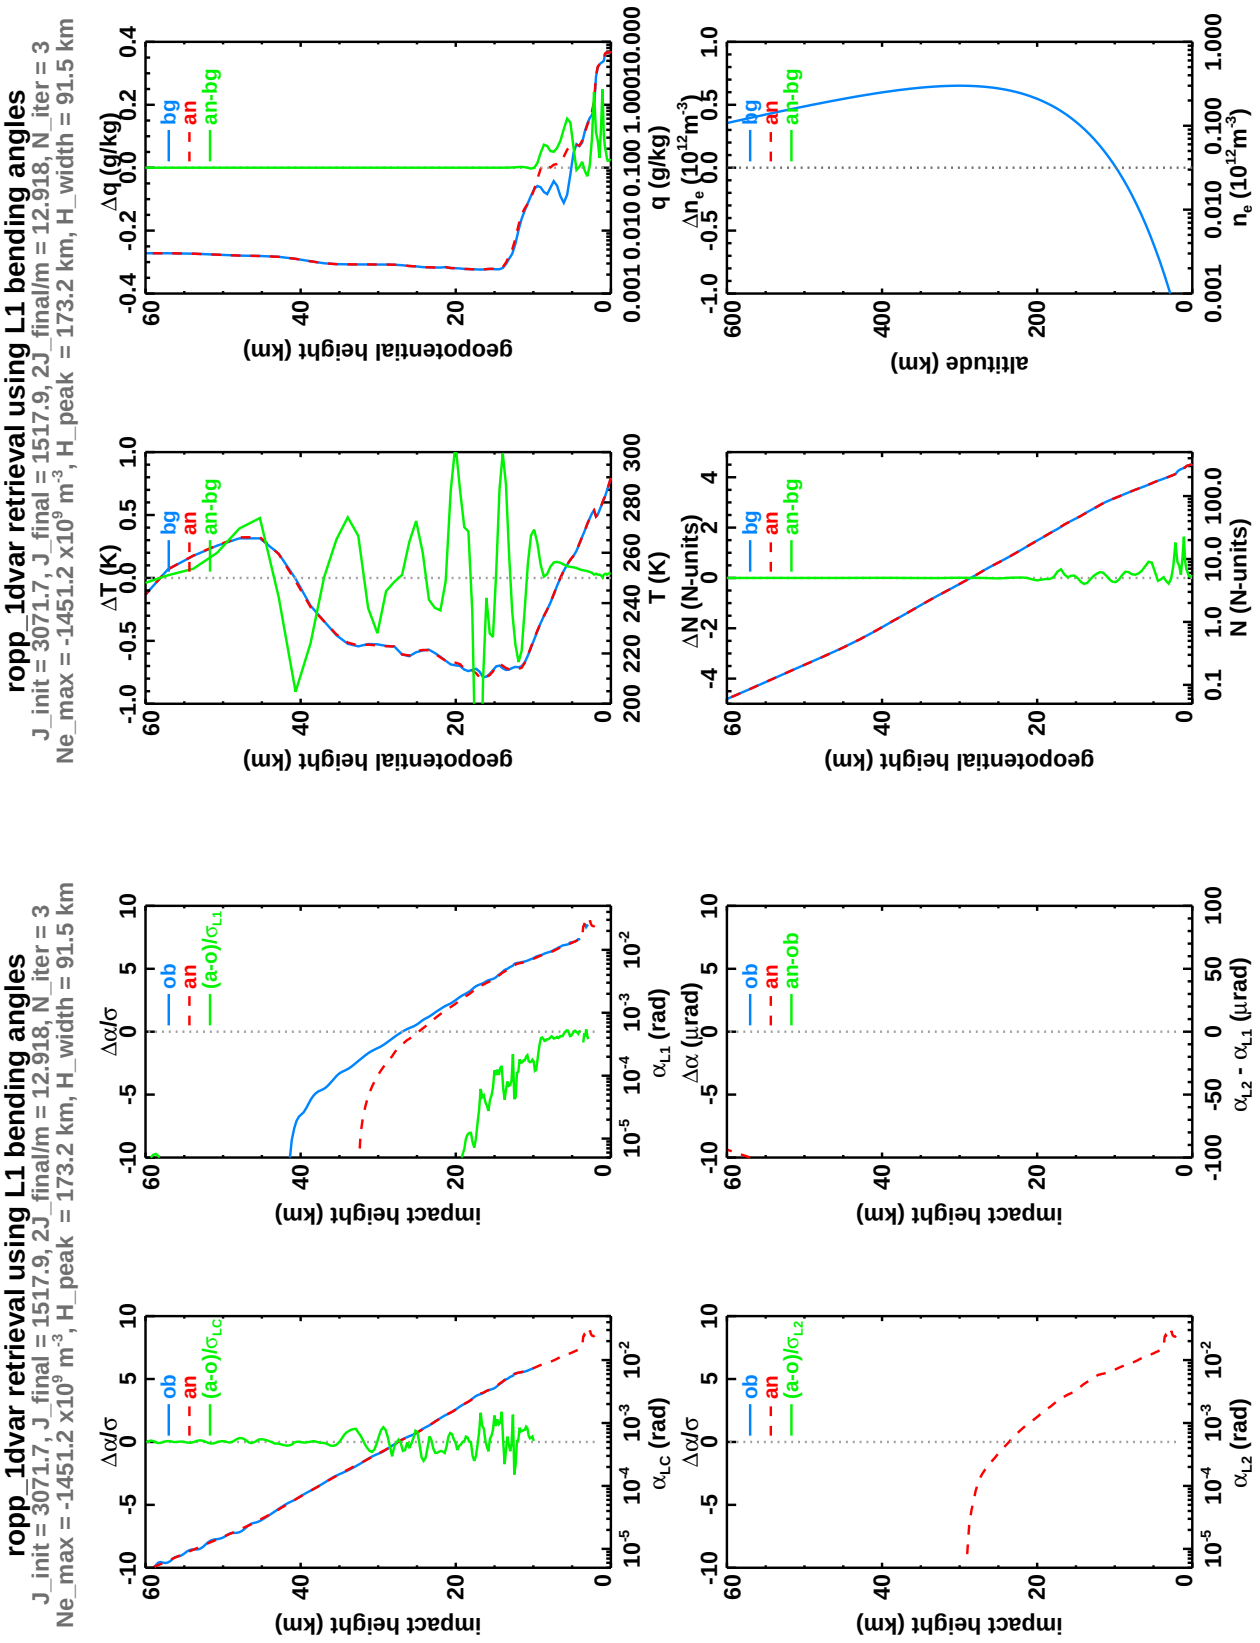


**Figure 3.7**: As for Fig 3.6, but using L1 bending angles only.

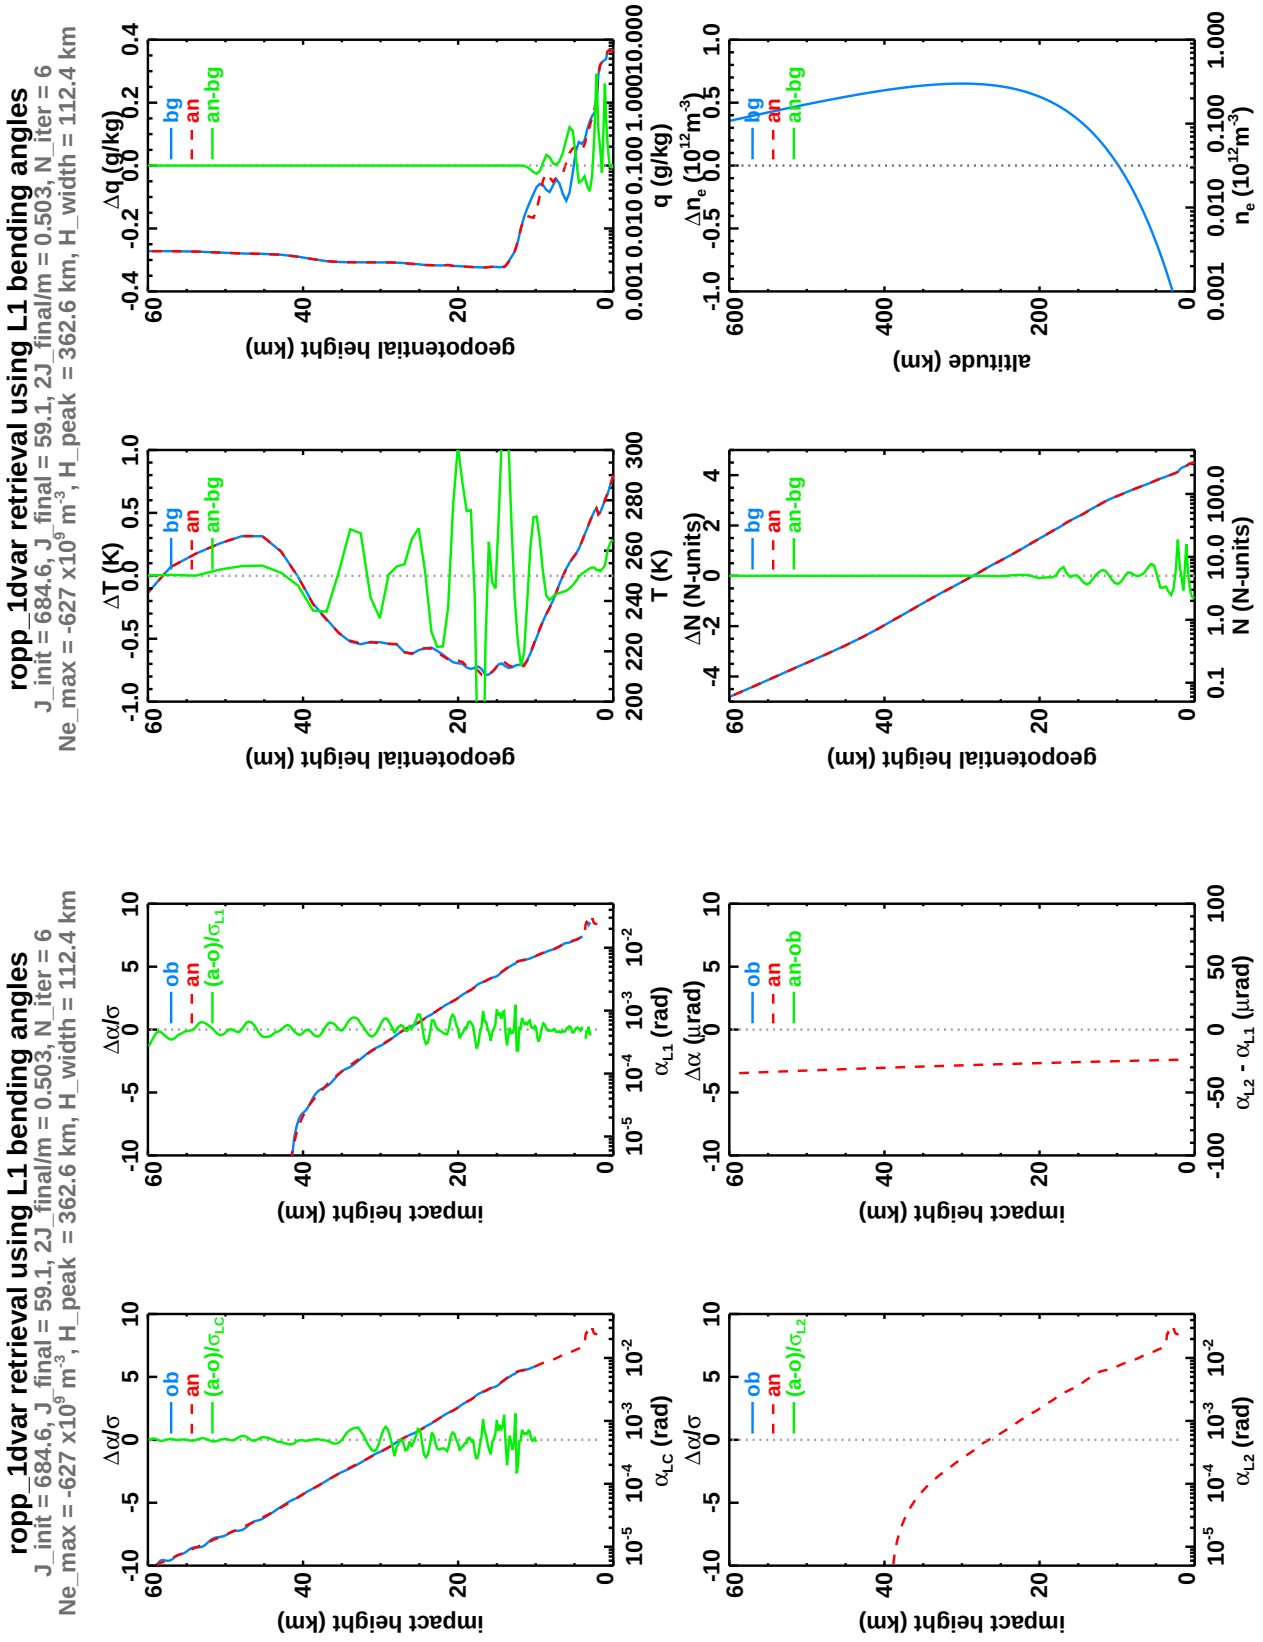


**Figure 3.8**: As for Fig 3.7, but using a background $n_e^{\max}$ of $-3.0 \times 10^{11}$ m$^{-3}$.

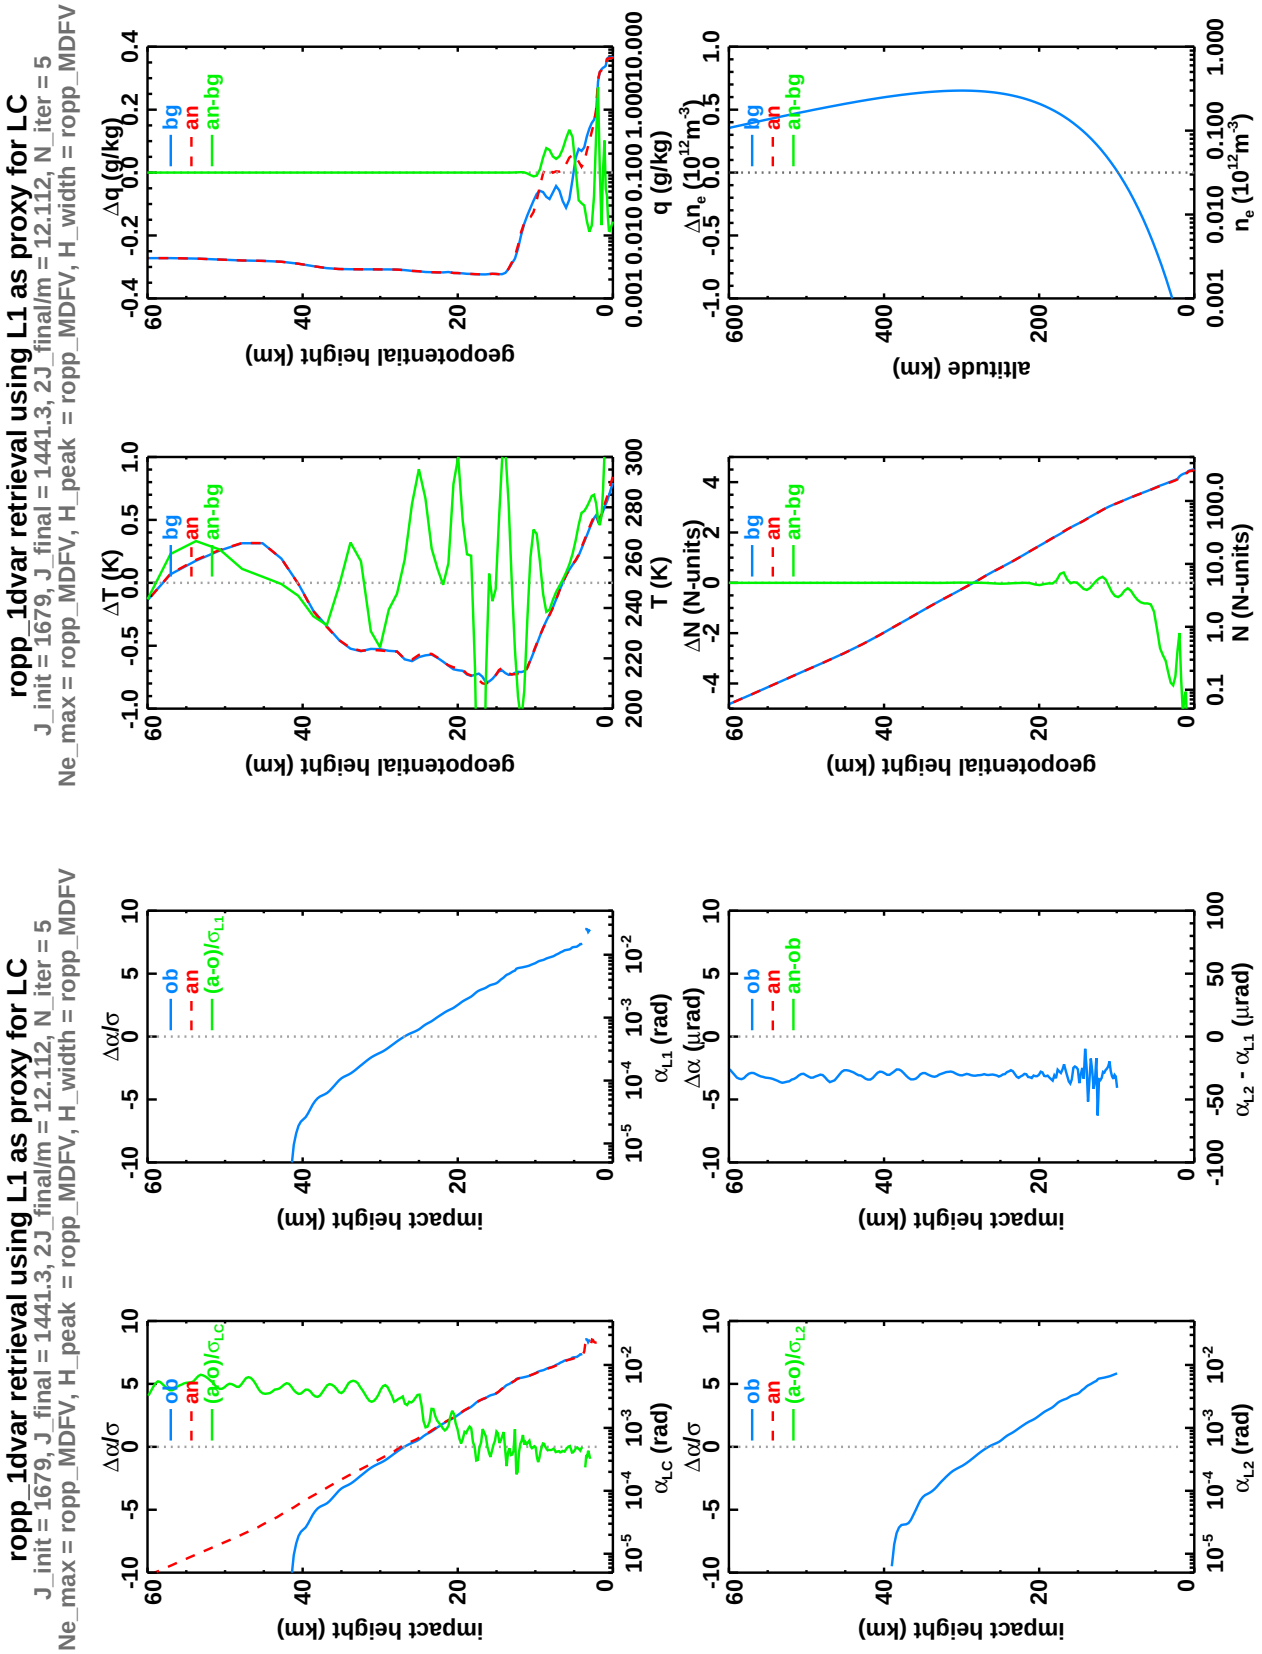


**Figure 3.9**: As for Fig 3.5, but using L1 as a proxy for LC.

## 3.2 Multiple profiles

We have seen that it is at least feasible to carry out 1D–Var retrievals using the L1 and L2 bending angles rather than the ionospherically corrected LC bending angles, which is the current operational practice. This procedure makes use of additional data in the lower atmosphere, mainly coming from L1. This has been shown, in two examples, to have some effect, particularly on humidity and surface pressure retrievals. The defining parameters of the model ionosphere can also be retrieved. The aim of this Section is to assess the generality and utility of these single profile indications, and to examine the integrity of dual (or single) frequency retrievals more closely.

To do this, all available 247 level GRAS profiles from five consecutive days in June 2012 (about 3500 profiles in total), and their colocated 91 level ECMWF background fields, were passed through the four experiments descibed in Sec 3.1, namely:

- NEUT: Use LC, no extrapolation of L2–L1.
- IONO: Use L1 and L2 directly, and a model ionosphere.
- IONO2: Use L1 directly, nullify L2, and use a model ionosphere.
- NEUT2: Use L1 as a proxy for LC, with no ionosphere in the forward model.

The statistics of these retrievals are briefly discussed in the following sections.

### 3.2.1 NEUT (LC)

The mean LC O–Bs and O–As are generally fairly small when normalised with respect to $\sigma(\alpha_{LC})$, as shown in Fig 3.10. The standard deviations of the O–As are a little smaller than those of the O–Bs, which shows that the retrieval is doing as expected. In state space, the A–Bs of temperature, humidity and pressure have a small mean, and standard deviations up to about 1.0 K, 0.2 g/kg and 1.3 hPa respectively. These increments reduce the mean surface refractivities by about 1 N-unit on average.

About 91% of the available profiles can be processed in this way. Most of the others fail to be usable because there are not enough L2 points to generate a usable LC.

### 3.2.2 IONO (L1 and L2)

The results of this experiment are plotted in Fig 3.11. The background L1 and L2 above 10 km have significant positive biases against observations. The 1D–Var retrieval procedure largely eliminates these, and also reduces the standard deviations. As a result, the 'observational' components of the cost function, $J_{\mathrm{OBS}}$, are reduced from their initial very large values to something similar to that seen in NEUT.

The standard deviations of the increments (A–B) of temperature, humidity and especially surface pressure are larger in this experiment than in NEUT, although the mean increments are about the same. These increments lead to an increase in the lower tropospheric refractivity, which is the opposite signal to that seen in NEUT. The components of the 'state' cost function, $J_{\mathrm{STATE}}$, show that the surface pressure is significantly perturbed by the retrieval.

About 82% of the available profiles can be processed in this way. It might have been hoped that this number would be higher than that of NEUT (91%), but many retrievals fail to converge at all. This needs further investigation.

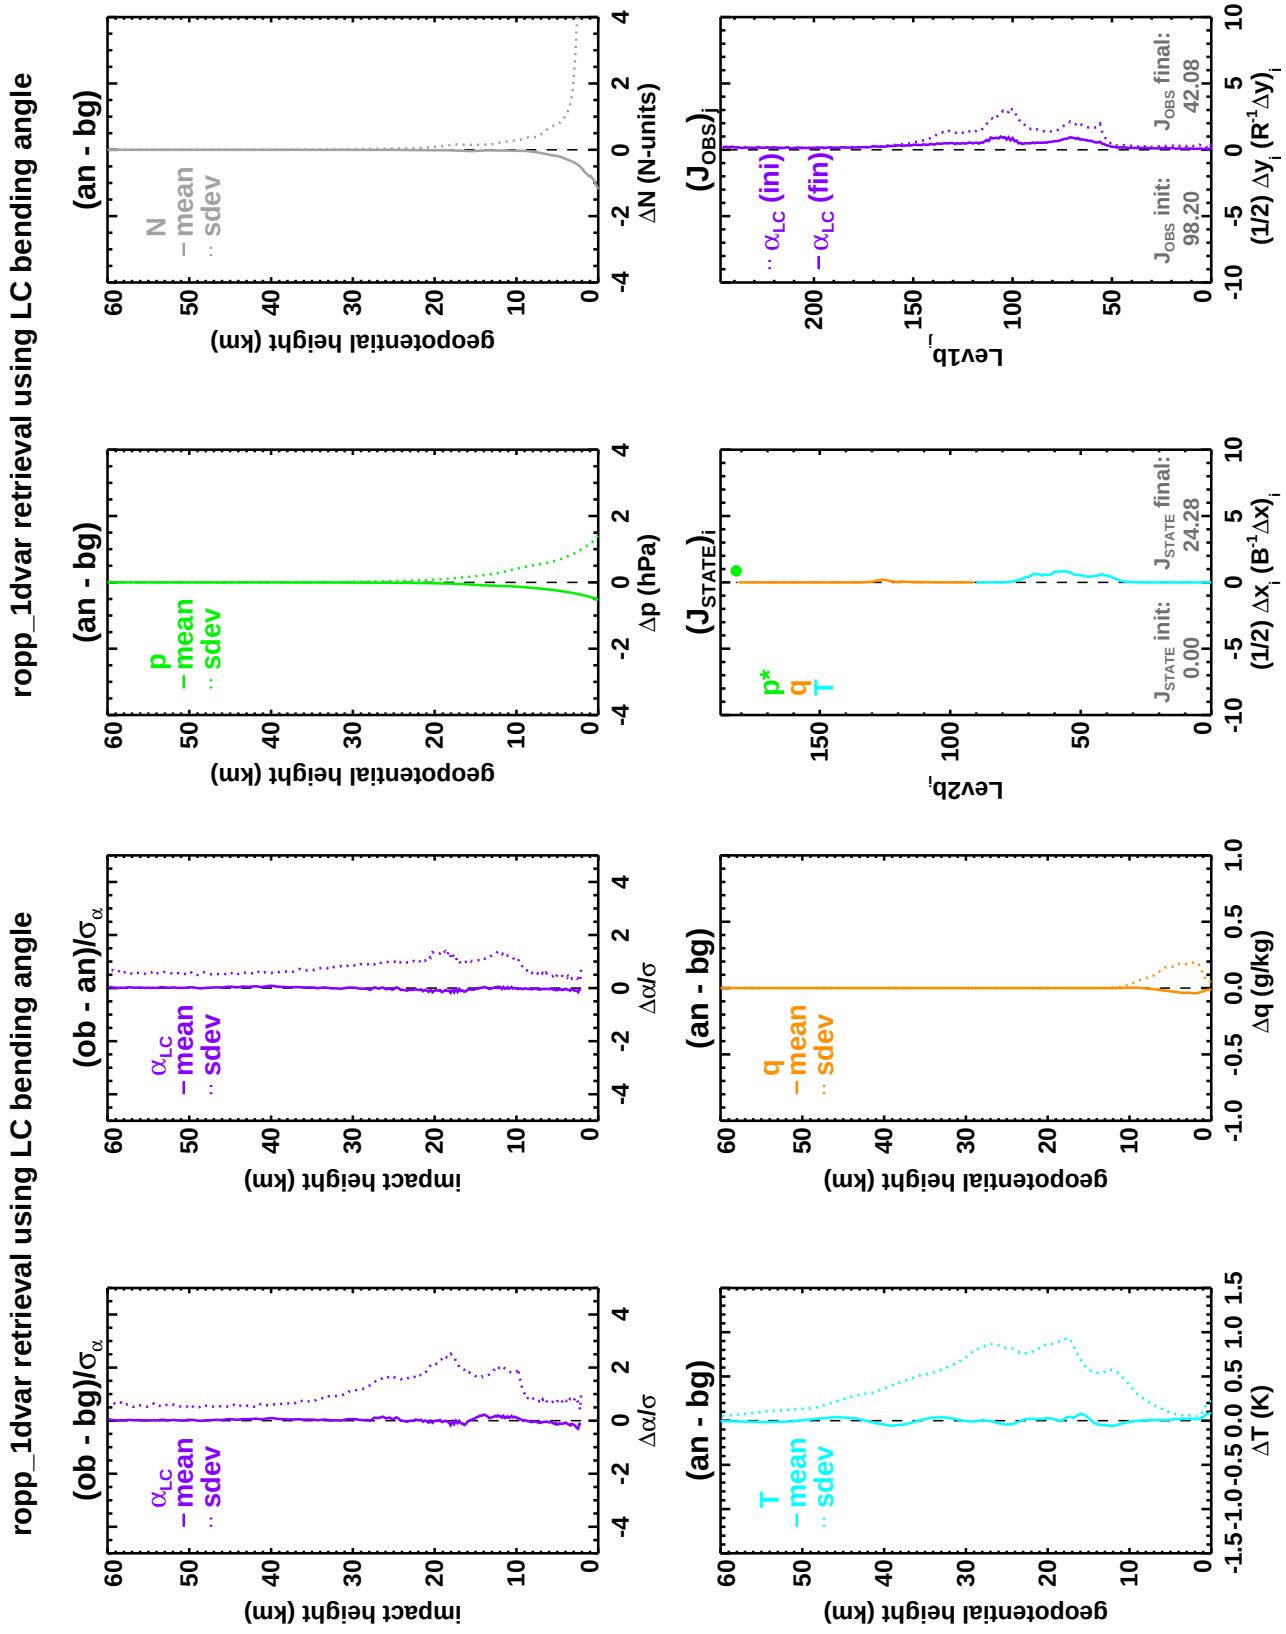


**Figure 3.10**: Statistics of 3500 GRAS retrievals using LC bending angles. Left hand panel: O–B, O–A, (O–B)/$\sigma_\alpha$ and (O–A)/$\sigma_\alpha$ for LC. Right hand panel: A–B of temperature and humidity, and level-by-level contributions to the state and observational components of the overall cost function.

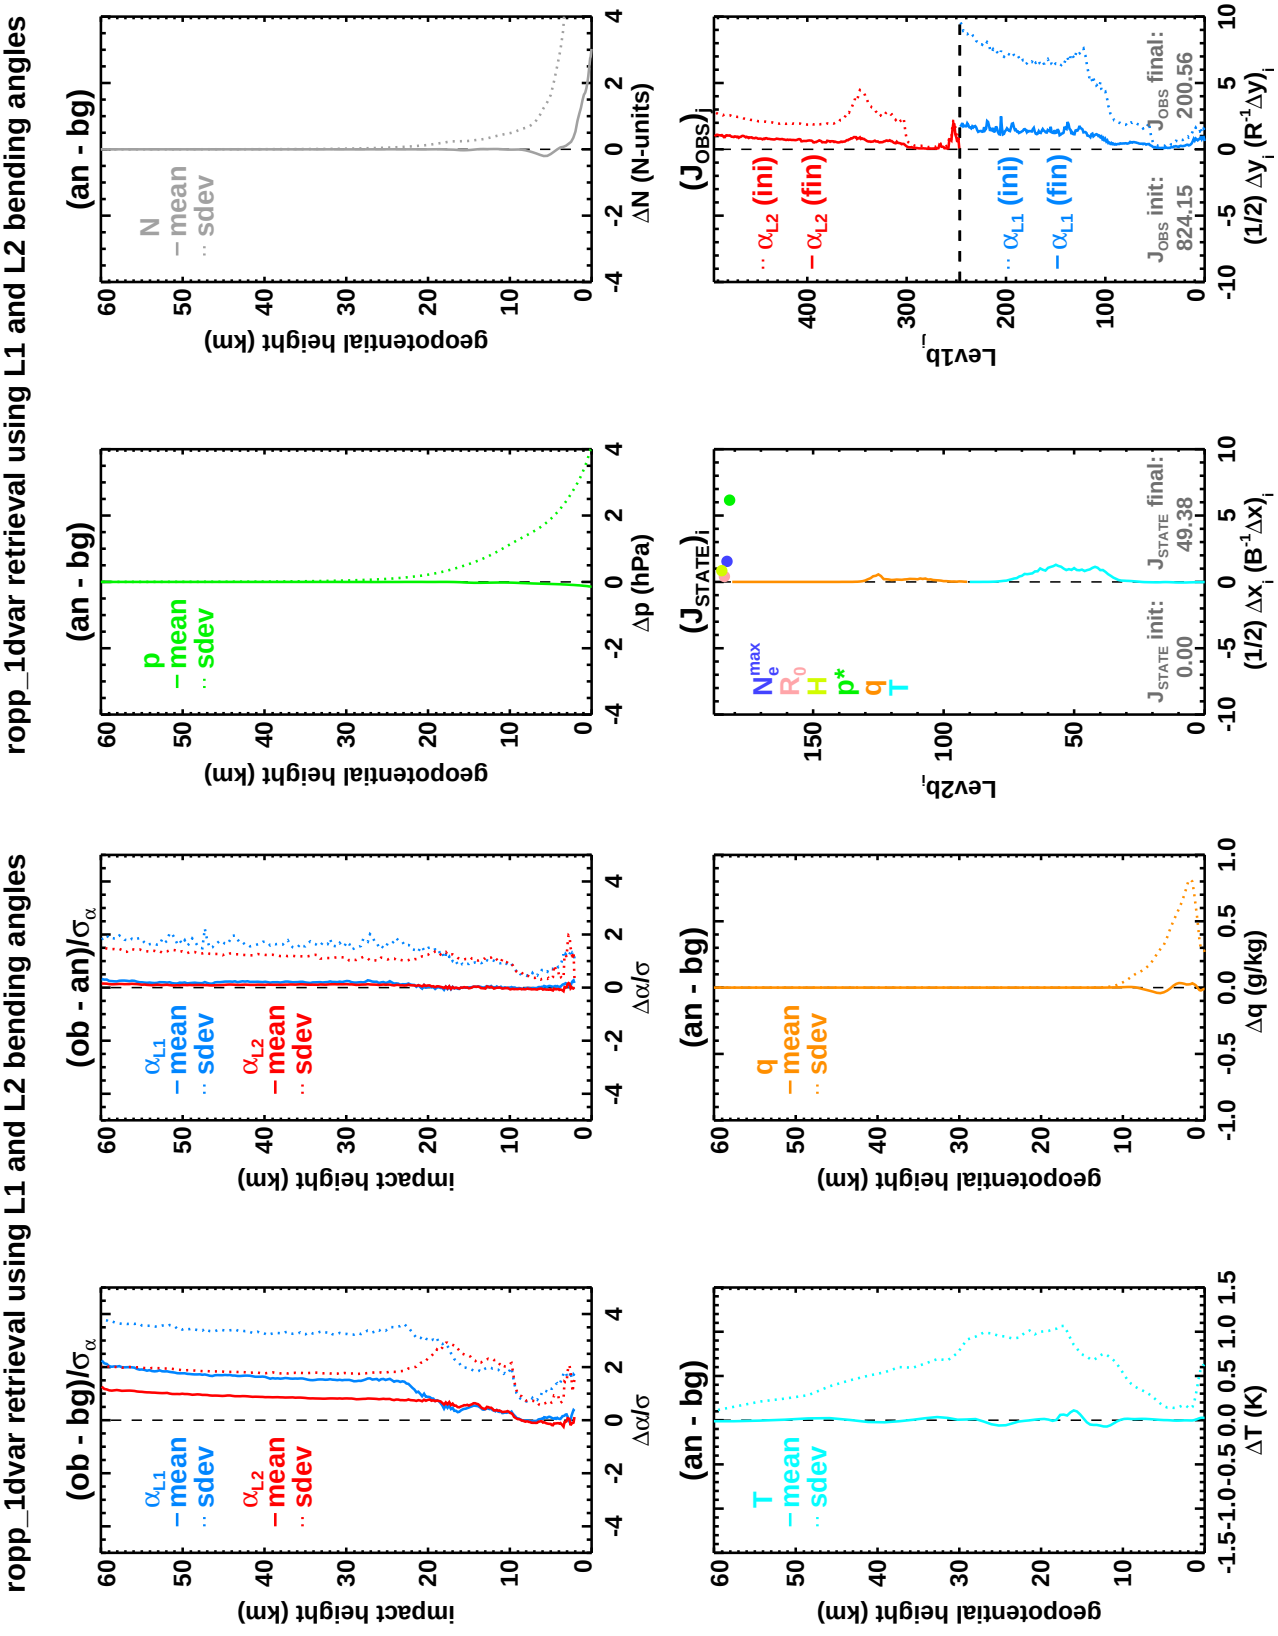


**Figure 3.11**: As for Fig 3.10, but using L1 and L2 bending angles, and direct modelling of the ionosphere in the forward model.

### 3.2.3 IONO2 (L1 only)

The results of this experiment are shown in Fig 3.12. In this case the profile-dependent adjustment of the background peak electron density cannot be applied (since there is no L2 from which to construct L1–L2), and this quantity therefore defaults to $3.0 \times 10^{11}$ m$^{-3}$. This leads to a significant bias in the background bending angles against observations above 15 km, which the 1D–Var retrieval does little to improve. The standard deviation is, however, reduced below 40 km, which means that the reduction in $J_{\text{OBS}}$ is greater than the increase in $J_{\text{STATE}}$.

The tropospheric increments are very similar to those of the IONO experiment.

The sensitivity of the retrieved ionospheric bending angles to the background peak electron density revealed by this analysis, and by the convergence/non-convergence of the retrieval of the anomalous profile discussed in Sec 3.1.2, strongly suggest that this question deserves further study.

Only 68% of the available profiles can be processed in this way. The rest fail to converge at all. Again, this needs addresssing.

### 3.2.4 NEUT2 (L1 a proxy for LC)

Finally, Fig 3.13 records the results of the NEUT2 experiment, in which L1 is simply used as a proxy for LC in a standard retrieval, without any accounting for ionospheric bending. Naturally, the O–Bs and O–As are very large in the ionosphere, as was seen in the single profile results, and not made any better by the retrieval, since the forward model contains nothing to alter the high level bending angles. Interestingly, large increments to temperature, humidity, pressure and refractivity occur at all heights, even where ionospheric effects might be thought to have less impact. This is presumably a sign of the minimiser struggling to find a solution that fits the unrealistic — as far as its forward model is concerned — observation. Further evidence of this is the very large (about 11 $\sigma$) surface pressure increments. Because of this, and its other large tropospheric impacts, this experiment remains the outlier.

About 96% of the available profiles can be processed in this way, 5% more than for the NEUT experiment, because the complete unavailability of L2 obviously does not prevent a profile being used in the this case. Note that these occultations may have formally converged, in the sense that the cost function or the state vector have stopped changing, but they have not necessarily converged well, in the sense of having a close least squares fit. For example, about 58% of the NEUT2 retrievals have a final scaled cost function $2J/m$ greater than 5. For NEUT, IONO and IONO2 the figures are 2%, 4% and 27% respectively, which further confirms the poor convergence of IONO2 suggested above.

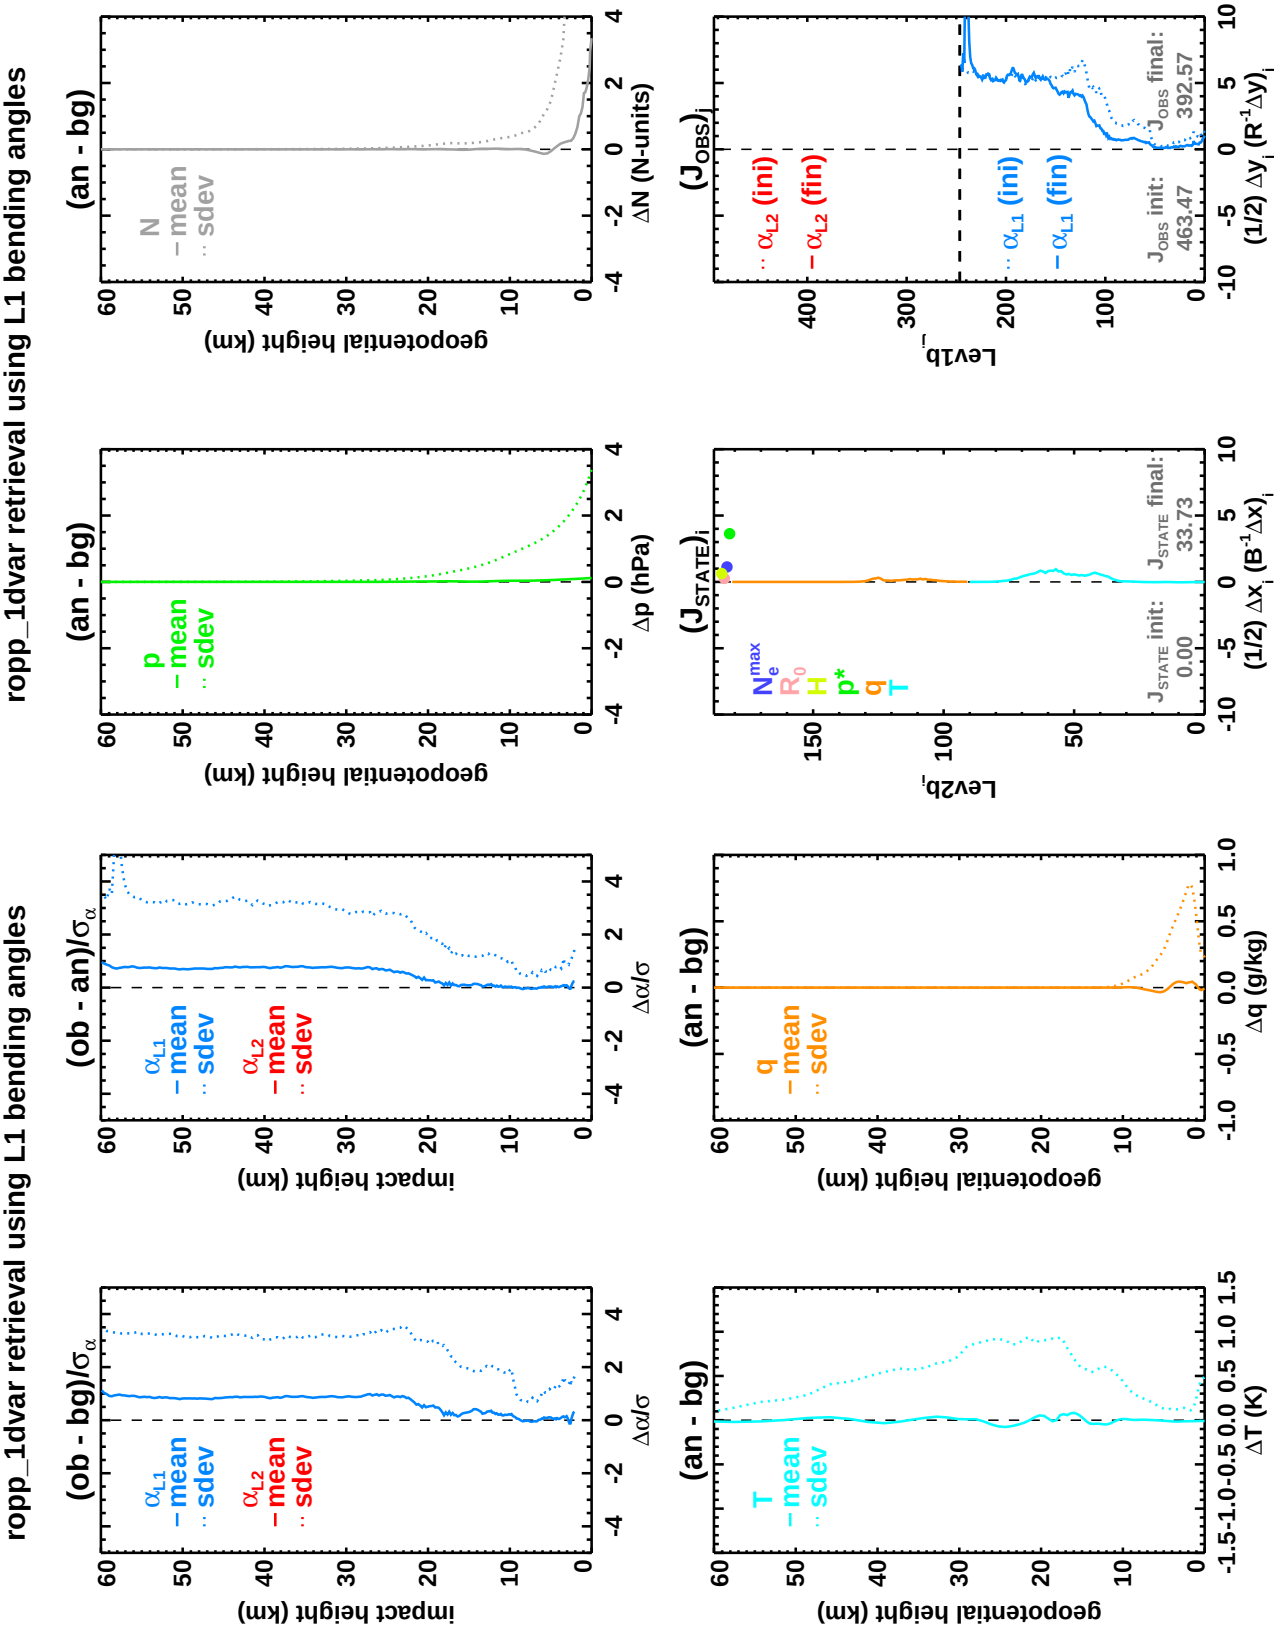


**Figure 3.12**: As for Fig 3.11, but using L1 bending angles only.

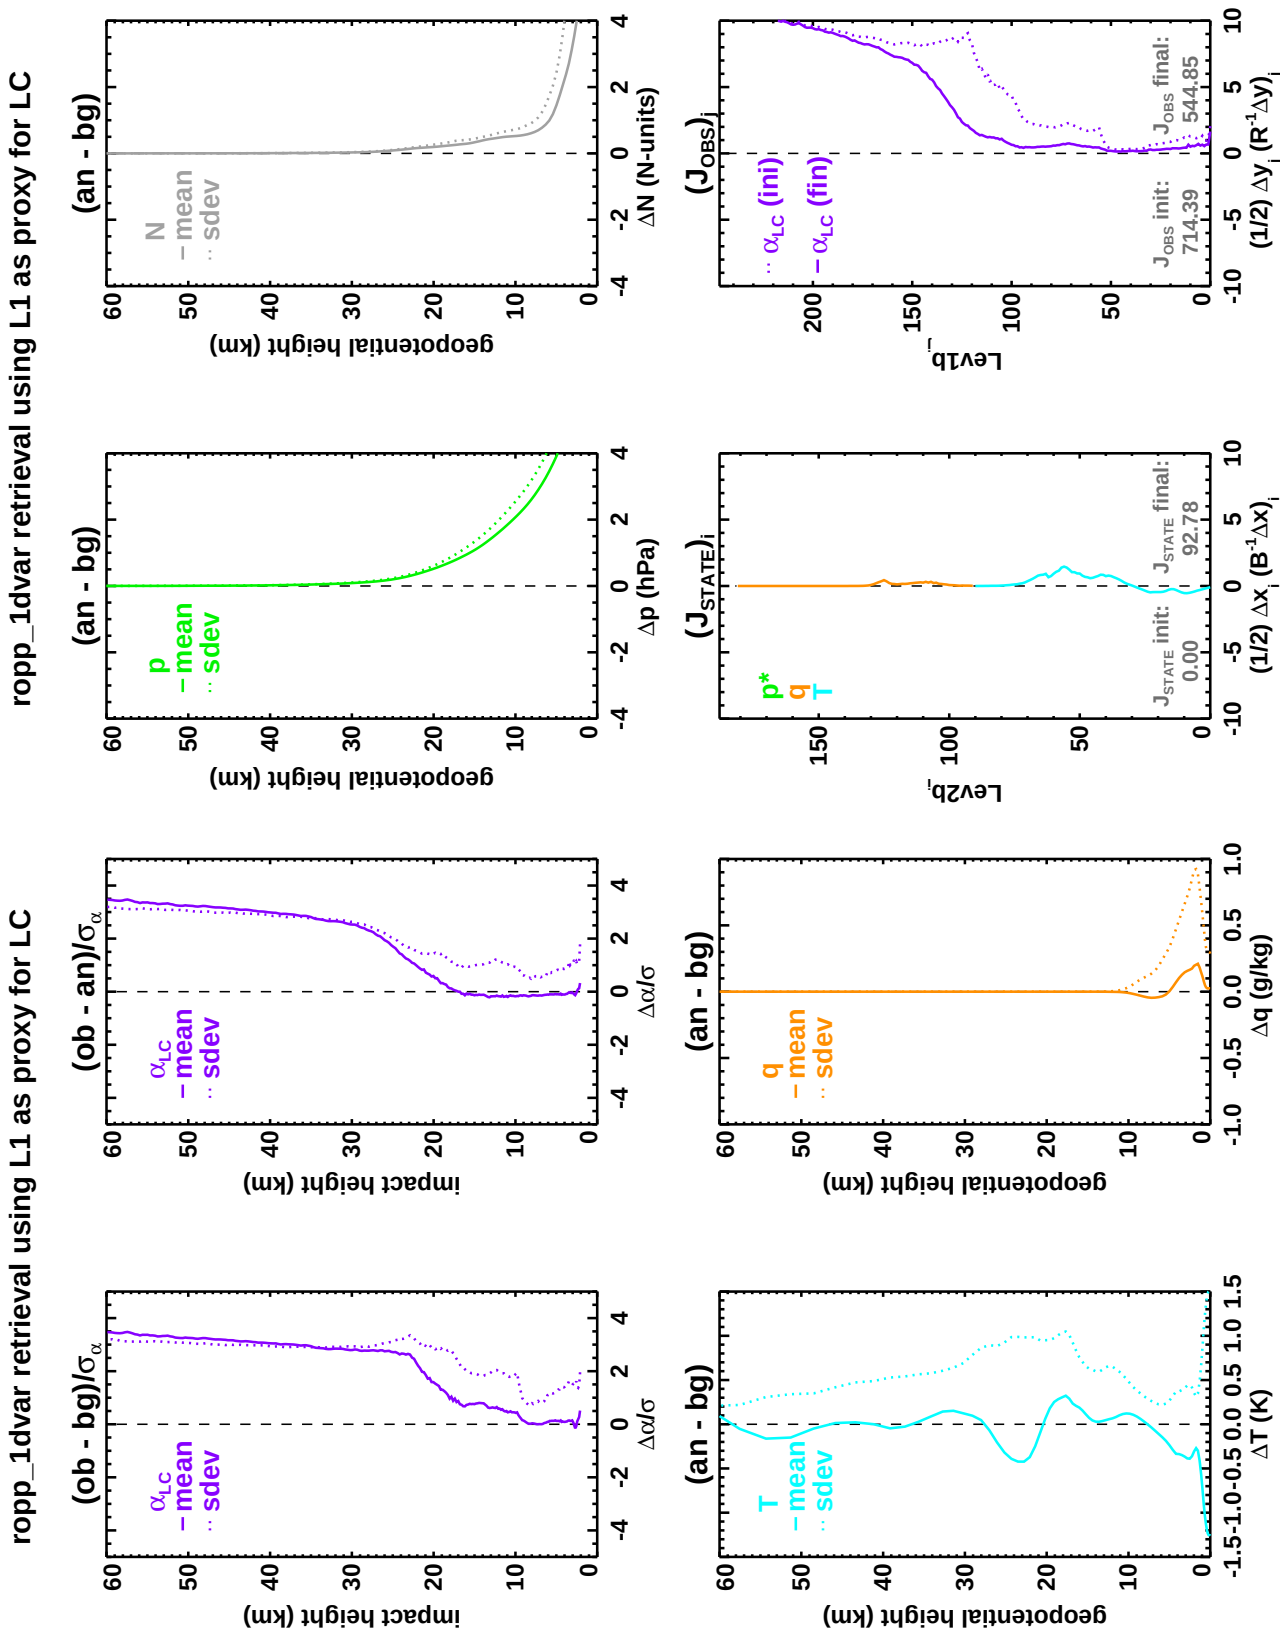


**Figure 3.13**: As for Fig 3.10, but using L1 as a proxy for LC.

### 3.3 Retrieved ionospheric parameters

It must be understood that the use of a simple ionospheric model is simply another way of accounting for the unknown and, for most NWP purposes, uninteresting properties of the ionosphere. The purpose of this report is to contrast this method with the usual one of linearly combining L1 and L2 to eliminate ionospheric effects. But, given that the three ionospheric parameters $n_e^{\max}, r_0$ and $H$ are adjusted and returned by the retrieval process, it would be remiss not to examine them.

The top left panel of Fig 3.14 shows the vertical total electron content (TEC) (equal to $\sqrt{2\pi e}\, n_e^{\max} H$ for a Chapman layer of peak density $n_e^{\max}$ and width $H$), derived from the 3500 profile L1 and L2 experiment IONO, gridded at 10° resolution. It peaks in the tropics at about 30 TECU (= $3.0 \times 10^{17}$ m$^{-2}$), and is larger in the summer hemisphere than the winter one. This is as expected (e.g. [1], [2]). The peak height $r_0$ and width $H$ are not on average much different from their background values of 300 km and 75 km respectively. The anomalous occultations (top right panel of Fig 3.14) occur more frequently in the Southern hemisphere than the Northern. The mean scaled cost function $2J/m$ at analysis is a reasonable 0.8, and the mean iteration count is about 5.

Fig 3.15 shows the analagous plots for the L1-only IONO2 experiment. The generally smaller TECs presumably arise from the different background peak electron density. $r_0$ and $H$ have small mean increments, as before. The higher final scaled cost function $2J/m$ and iteration count $N_{\text{iter}}$ show that the minimiser needs to do more work before it converges, and that this convergence is less complete. This too is likely to be a result of the poorer initial electron density estimate. Curiously, $2J/m$ is biggest in the Northern hemisphere, where there are fewer anomalous occultations but — or perhaps because — the TEC is largest. This suggests that the fitting is poorest where the ionospheric bending is largest, regardless of its sign. The hemispheric asymmetry is small, however.

For the record, the $2J/m$ and $N_{\text{iter}}$ maps for the two neutral bending angle experiments, NEUT and NEUT2, are shown in Fig 3.16. Interestingly, in the standard NEUT experiment, the final scaled cost function peaks in the tropics. Its mean value is a satisfying 1.1. The mean iteration count, which is also weakly peaked in the tropics, is about 4, which is a little less than in the IONO or IONO2 experiments. In the NEUT2 experiment, the large final scaled cost function and iteration count in the tropics and Northern hemisphere reflect the understandable difficulty in using L1 as a proxy for LC at times and places where the ionospheric component of the bending is large.

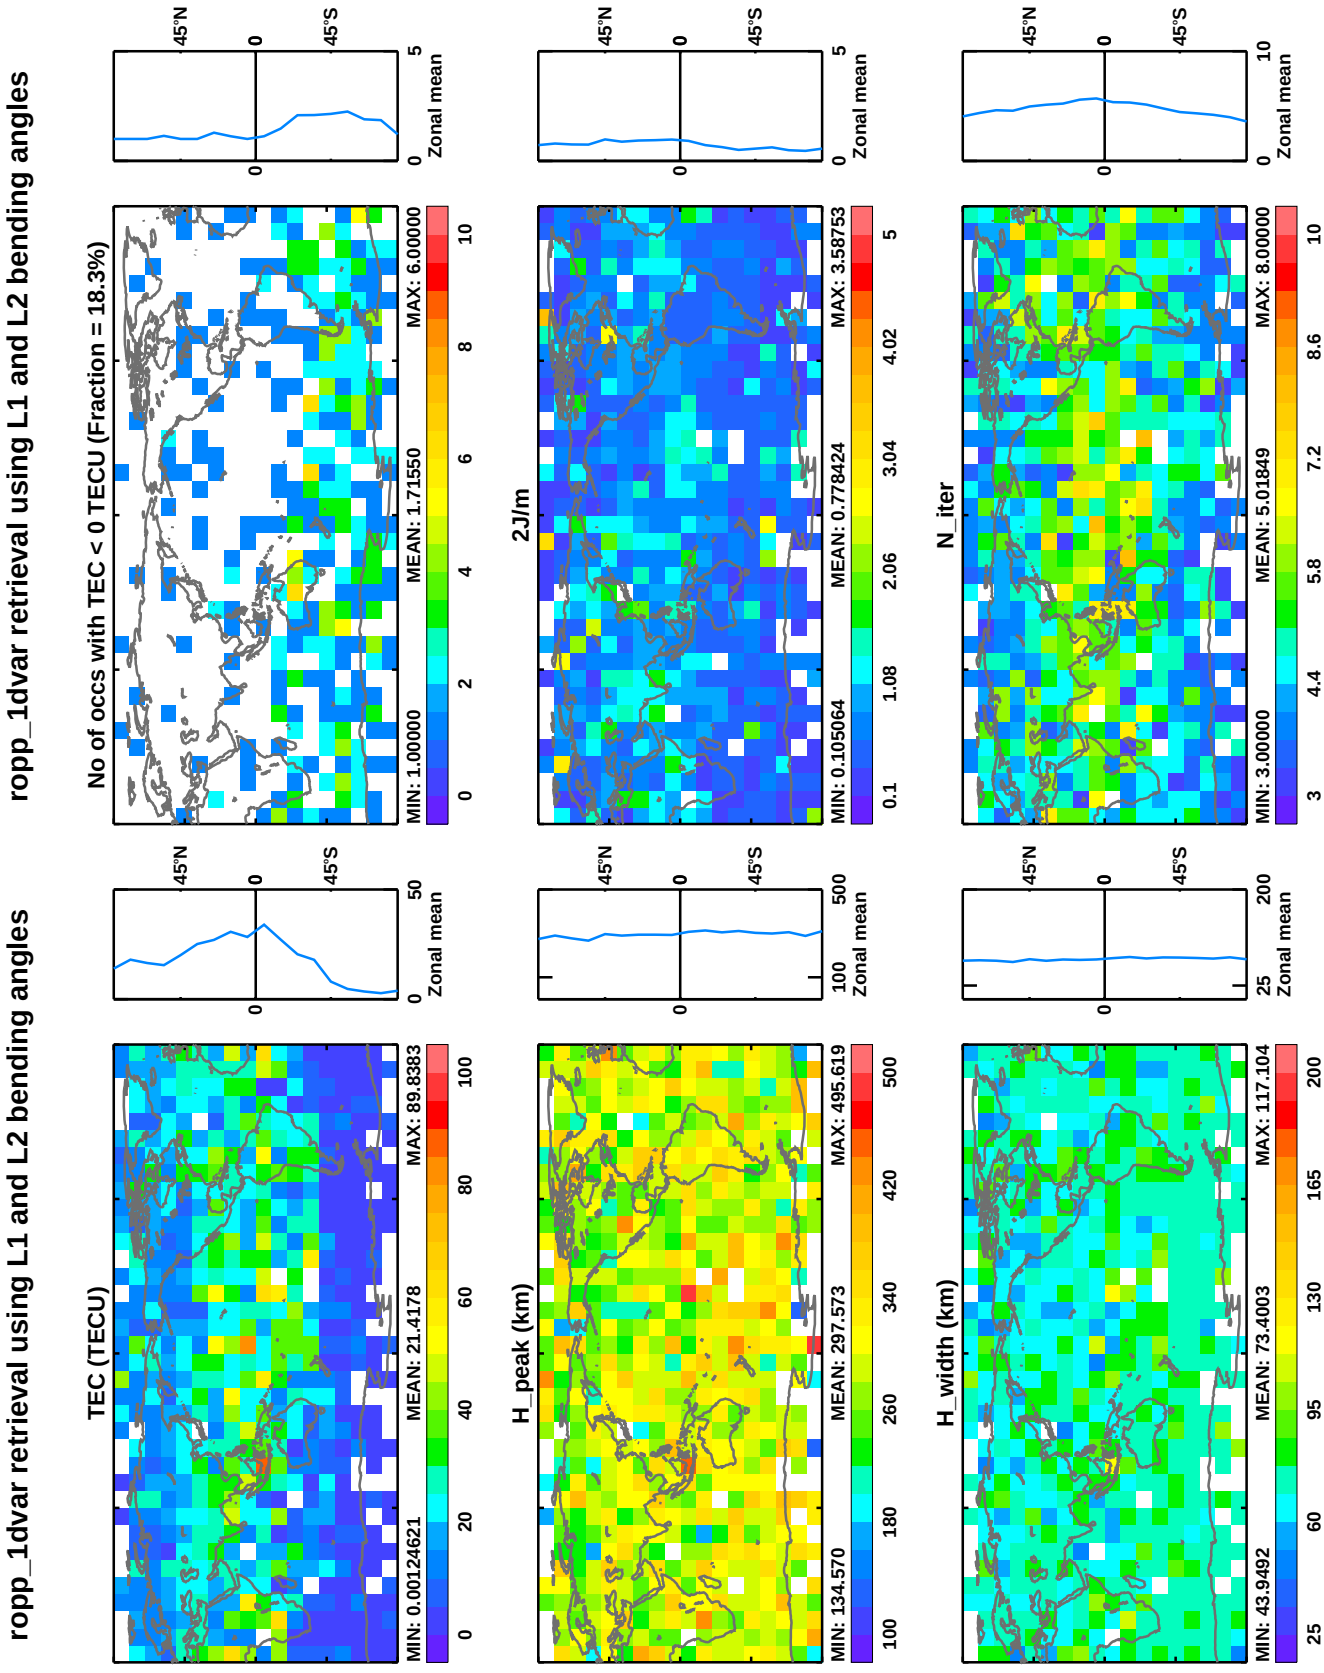


**Figure 3.14**: Gridded mean retrieved ionospheric parameters for the IONO (L1 and L2) experiment. Left hand column: total electron content, peak height and width. Right hand column: number of anomalous occultations, final scaled cost function and minimiser iteration count.

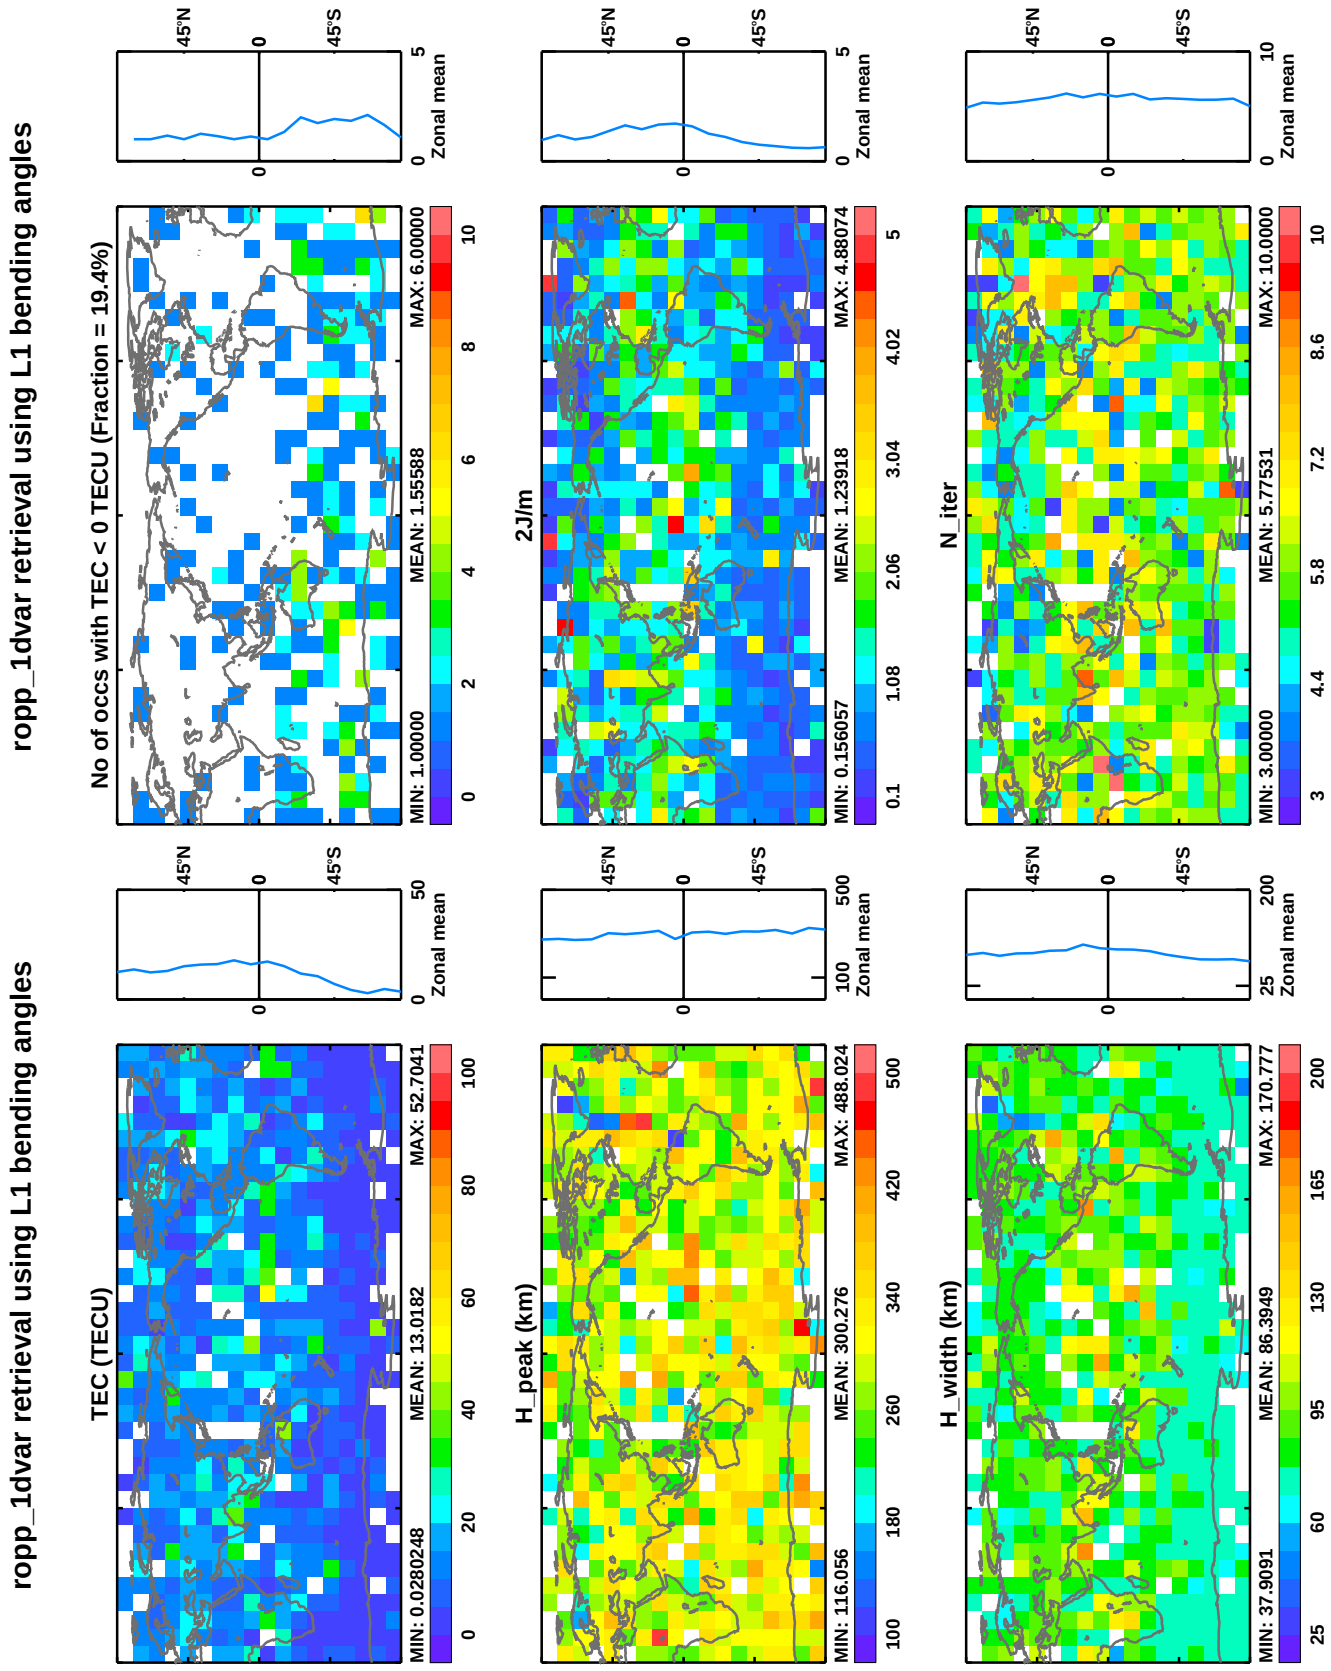


**Figure 3.15**: As for Fig 3.14, but for the IONO2 (L1-only) experiment.

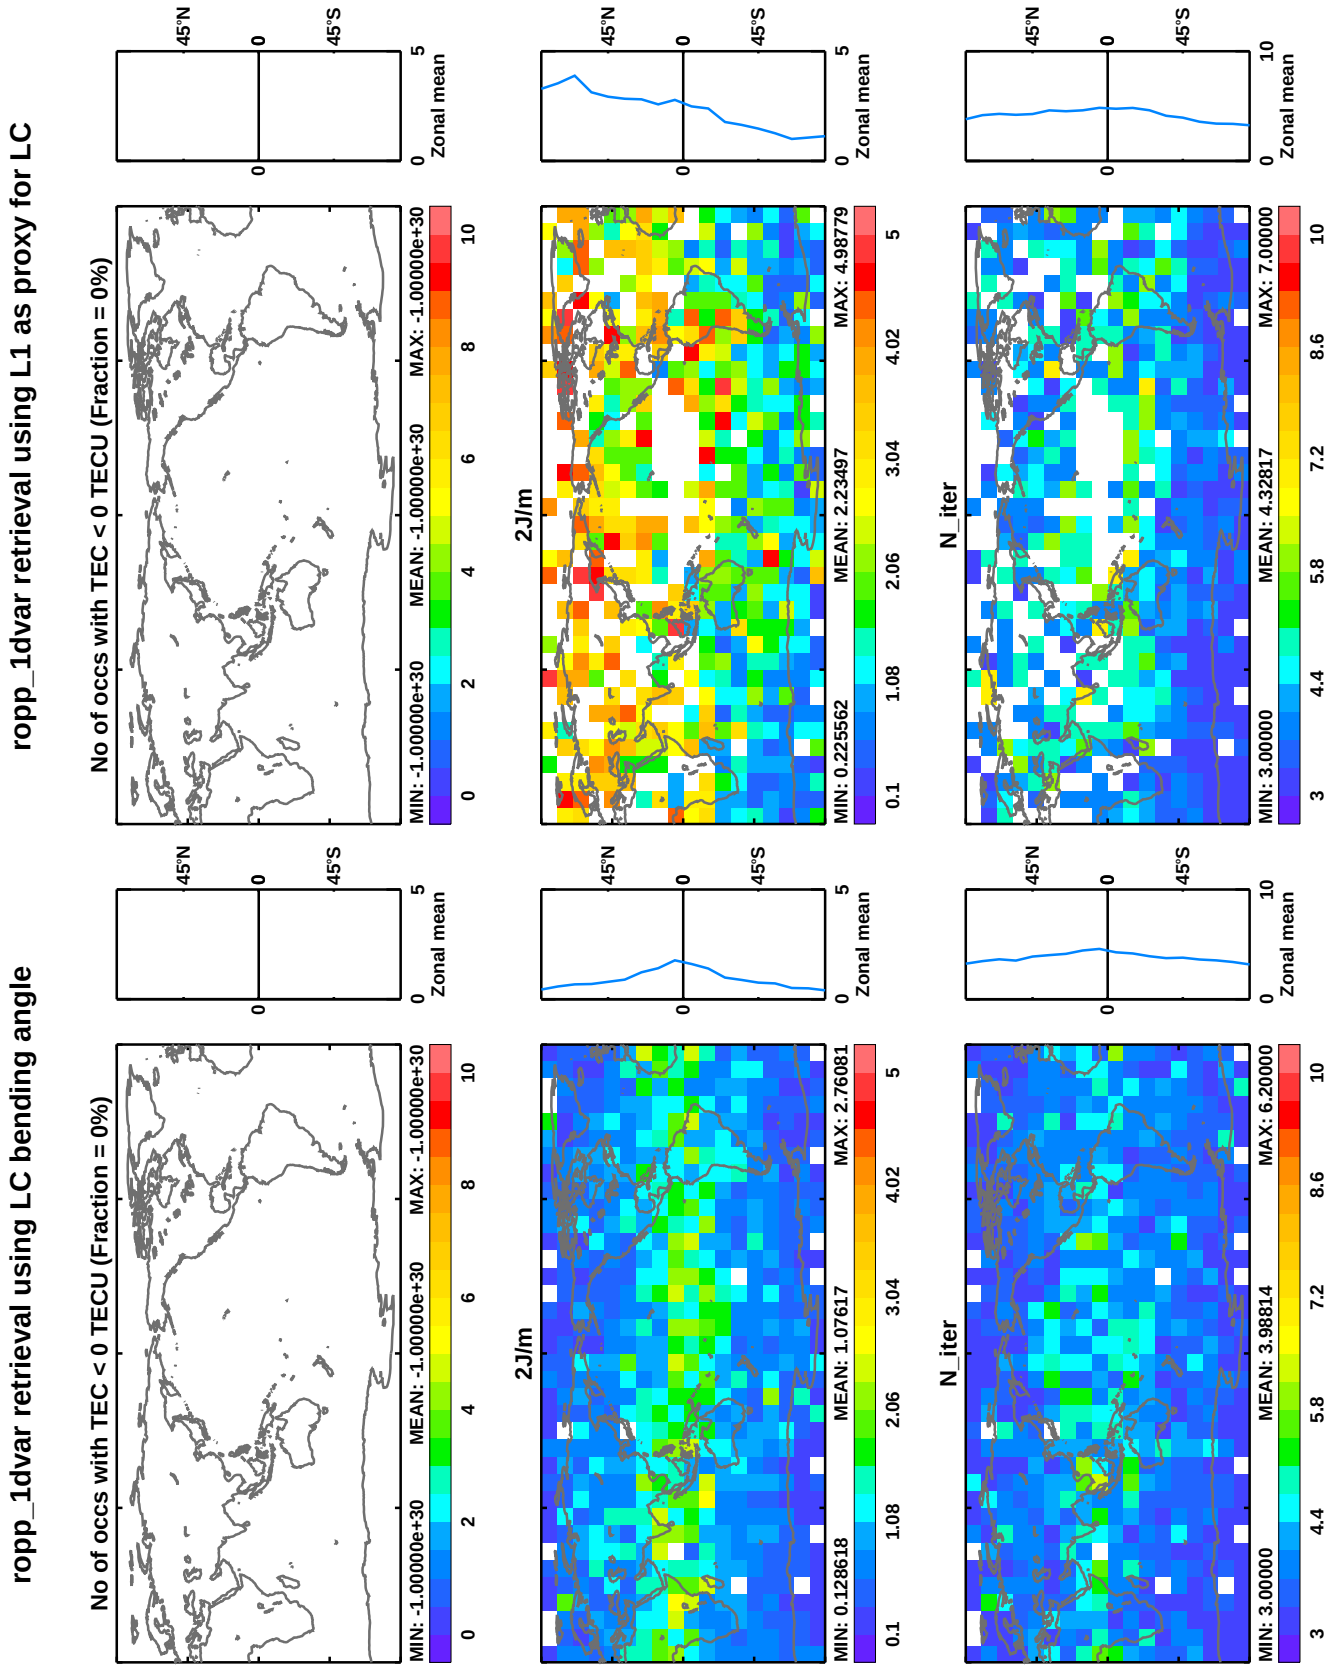


**Figure 3.16**: Gridded mean number of anomalous occultations, final scaled cost function and minimiser iteration count. Left hand column: the LC experiment NEUT. Right hand column: the L1-as-LC experiment NEUT2.

# 4 Summary, conclusions and future work

This report describes the results of a first investigation into the feasibility of carrying out 1D–Var retrievals using L1 and L2 bending angles, rather than the usual 'ionospherically corrected' linear combination LC. This option demands changes to all components of the retrieval system: the observation vector (the replacement of $\alpha_{LC}$ by $\alpha_{L1}$ and $\alpha_{L2}$); the observation errors (mainly, to increase them at heights where the L1 and L2 bending angles are much larger than those of LC); the state vector (which needs to be augmented by the ionospheric parameters $n_e^{\max}$, $r_0$ and $H$); and the background errors (to include error estimates of these quantities). Such changes have been implemented in the ROM SAF's Radio Occultation Processing Package, ROPP ([11]). This code was used to carry out four experiments, which were designed to assess the possible value of using L1 and L2 (or just L1) bending angles in 1D–Var retrievals.

The main results of these studies are as follows.

- It is possible in principle to forward model L1 and L2 bending angles by assuming a simple Chapman layer ionosphere.
- Doing so allows the use of (L1) bending angle data in (lower) regions where L2 is frequently absent. In these cases the LC bending angle is usually either taken to be absent in such regions, or calculated from an L2 that is derived by extrapolating L1–L2 from (higher) regions, where they are both present.
- The use of bending angle data in the troposphere leads to bigger impacts on tropospheric humidity, temperature, surface pressure and (therefore) refractivity.
- It is even possible to carry out single frequency retrievals, in which L2 is entirely absent. The results are similar to those of dual frequency retrievals, except for the ionospheric parameters and higher level bending angles. This is because the peak electron density is highly variable in space and time, and without information on L1–L2 to guide its initial value, we are forced to set this to a constant value, which is a poor approximation. This aspect of the method needs further investigation. Even so, using L1 with a simple ionospheric model is far better than simply ignoring the ionospheric component of the bending angles by adopting L1 as a proxy for LC.
- 'Anomalous' occulations, in which the L1 bending angle is greater than that of L2, are impossible to model by a spherically symmetric ionosphere without invoking negative electron densities. Heedlessly doing so, however, delivers reasonable fits to the observed L1 and L2 bending angles.

This work has suggested several areas that deserve further study. Principal among these are the following.

- The peak electron density, $n_e^{\max}$, has emerged as a key parameter in the calculation of the ionospheric component of the bending angles. It is therefore important to choose good values of the background, $b$, and the first guess in the variational retrieval, $x_0$. Currently the latter is crudely estimated, where possible, from the mean difference between the L1 and L2 bending angles above 30 km, and this value is also used as the background value in the 1D–Var scheme. The two variables are theoretically different, and, given that the background should really be

independent of the observations, it would be worthwhile redoing the experiments with such an observationally-independent background $n_e^{\max}$ — probably with a larger assumed error than $2 \times 10^{11}$ m$^{-3}$, which works satisfactorily when $b$ is set equal to $x_0$, but which would unduly constrain a variational retrieval algorithm presented with a background that was likely to be quite inaccurate. A simple climatology of $n_e^{\max}$, such as neQuick (e.g. [9]), might also provide a better background peak electron density, and this too could be investigated.

- The large increase in the variability of the surface pressure and humidity, resulting from the use of L1 and L2 in place of LC, should be investigated. It is known to be due to a general broadening of these distributions, rather than the distorting influence of a few outliers. Note, however, that about one third of the L2 bending angles fail to penetrate below 10 km, and therefore the same fraction of the ionospherically corrected LC bending angles will be absent from the standard NEUT experiment. Some of the increased tropospheric variability of the retrievals in the other experiments may therefore simply be due to the presence of more observations at lower levels.

- The other numerical details of the 1D–Var retrieval process, such as the lack of correlation between $\alpha_{L1}$ and $\alpha_{L2}$, and the assumed errors on the ionospheric parameters, could be usefully re-examined.

- The spectre of the anomalous occultations continues to haunt this field of work. Their abundance (about 13% of GRAS occultations) and curious distribution in space and time (see [2]) remain a puzzle, and should be investigated. We still do not know whether they are fundamentally flawed observations, which should be excluded from operational NWP assimilation systems, or whether whatever is causing their anomalous behaviour is being properly removed by the usual ionospheric correction method, leaving usable neutral bending angles behind. (Anomalous occultations currently pass Met Office quality control tests, and are therefore assimilated.) The methods of this report might be extended to investigate this question, by restricting attention to anomalous occultations, perhaps confined to certain regions/times/seasons. A comparison of the results of the LC and L1/L2 experiments might then throw some light on the question.

  In practice, however, anomalous occultations do not prevent 1D–Var retrievals from being undertaken, as long as the user is prepared to allow (and swiftly ignore) the implied negative electron densities.

- Studies of departures and increments can show the sensitivity to the different data and processing method described in this report, but they do not show whether they are better. This question could be answered by means of a trial of the new method, parallel to the current operational NWP data assimilation cycle, and this should be considered when some of the above difficulties have been overcome.

## Acknowledgements

Sincere thanks are extended to the reviewers of this report for their helpful comments and advice.

# Bibliography

[1] Culverwell, I. D. and Healy, S. B., Simulation of L1 and L2 bending angles with a model ionosphere, ROM SAF Report 17, http://www.romsaf.org/general-documents/rsr/rsr_17.pdf, 2015.

[2] Culverwell, I. D., Healy, S. B. and von Engeln, A., Investigation of anomalous operational GRAS occultations, Proc. 2nd International Conference on GPS Radio Occultation (ICGPSRO), Taiwan, May 2013, http://www.nspo.narl.org.tw/ICGPSRO2013/download/2nd_ICGPSRO_U3-3A-102-02_Healy_(for_Culverwell)_05-15-2013.pdf, 2013.

[3] Danzer, J., Healy, S. B. and Culverwell, I. D., A simulation study with a new residual ionospheric error model for GPS radio occultation climatologies Atmospheric Measurement Techniques 8, 3395-3404, 2015.

[4] EUMETSAT, Single Frequency Radio Occultation Retrievals, EUM/TSS/TEN/13/707742, Issue 2, 2013.

[5] Healy, S. B., and Culverwell, I. D., A modification to the ionospheric correction method used in GPS radio occultation, Atmospheric Measurement Techniques 8, 3385-3393, 2015.

[6] Kursinski, E. R., Hajj, G. A., Leroy, S. S. and Herman, B., The GPS Radio Occultation Technique, Terrestrial, Atmospheric and Oceanic Sciences, **11**, 1, pp53-114, 2000.

[7] Liao, M., Preliminary investigation with GNOS bending angle data: Monitoring, quality control and possible implications for ROPP, ROM SAF Visiting Scientist Report 32, SAF/ROM/DMI/REP/VS/32, 2017.

[8] Marquardt, M., Pers. Comm., 2017.

[9] Nava, B., Coisson, P., and Radicella, S., A new version of the neQuick ionosphere electron density model, J. Atmos. Solar-Terr. Phys., doi:10.1016/j.jatsp.2008.01.015, 2008.

[10] Rocken, C., R. Anthes, M. Exner, D. Hunt, S. Sokolovsky, R. Ware, M. Gorbunov, W. Schreiner, D. Feng, B. Herman, Y.–H. Kuo and X. Zou, Analysis and validation of GPS/MET data in the neutral atmosphere, J. Geophys. Res., 102, 29,849–29,866, 1997.

[11] ROPP User Guides, http://www.romsaf.org/ropp/index.php, 2017.

[12] Vorob'ev, V. V. and Krasil'nikova, T. G., Estimation of the accuracy of the atmospheric refractive index recovery from doppler shift measurements at frequencies used in the NAVSTAR system, USSR Phys. Atmos. Ocean, Engl. Transl., Vol. 29, 602–609, 1994.

**ROM SAF (and earlier GRAS SAF) Reports**

| | |
|---|---|
| SAF/GRAS/METO/REP/GSR/001 | Mono-dimensional thinning for GPS Radio Occulation |
| SAF/GRAS/METO/REP/GSR/002 | Geodesy calculations in ROPP |
| SAF/GRAS/METO/REP/GSR/003 | ROPP minimiser - minROPP |
| SAF/GRAS/METO/REP/GSR/004 | Error function calculation in ROPP |
| SAF/GRAS/METO/REP/GSR/005 | Refractivity calculations in ROPP |
| SAF/GRAS/METO/REP/GSR/006 | Levenberg-Marquardt minimisation in ROPP |
| SAF/GRAS/METO/REP/GSR/007 | Abel integral calculations in ROPP |
| SAF/GRAS/METO/REP/GSR/008 | ROPP thinner algorithm |
| SAF/GRAS/METO/REP/GSR/009 | Refractivity coefficients used in the assimilation of GPS radio occultation measurements |
| SAF/GRAS/METO/REP/GSR/010 | Latitudinal Binning and Area-Weighted Averaging of Irregularly Distributed Radio Occultation Data |
| SAF/GRAS/METO/REP/GSR/011 | ROPP 1dVar validation |
| SAF/GRAS/METO/REP/GSR/012 | Assimilation of Global Positioning System Radio Occultation Data in the ECMWF ERA-Interim Re-analysis |
| SAF/GRAS/METO/REP/GSR/013 | ROPP PP validation |
| | |
| SAF/ROM/METO/REP/RSR/014 | A review of the geodesy calculations in ROPP |
| SAF/ROM/METO/REP/RSR/015 | Improvements to the ROPP refractivity and bending angle operators |
| SAF/ROM/METO/REP/RSR/016 | Simplifying EGM96 undulation calculations in ROPP |
| SAF/ROM/METO/REP/RSR/017 | Simulation of L1 and L2 bending angles with a model ionosphere |
| SAF/ROM/METO/REP/RSR/018 | Single Frequency Radio Occultation Retrievals: Impact on Numerical Weather Prediction |
| SAF/ROM/METO/REP/RSR/019 | Implementation of the ROPP two-dimensional bending angle observation operator in an NWP system |
| SAF/ROM/METO/REP/RSR/020 | Interpolation artefact in ECMWF monthly standard deviation plots |
| SAF/ROM/METO/REP/RSR/021 | 5th ROM SAF User Workshop on Applications of GPS radio occultation measurements |
| SAF/ROM/METO/REP/RSR/022 | The use of the GPS radio occultation reflection flag for NWP applications |
| SAF/ROM/METO/REP/RSR/023 | Assessment of a potential reflection flag product |
| SAF/ROM/METO/REP/RSR/024 | The calculation of planetary boundary layer heights in ROPP |
| SAF/ROM/METO/REP/RSR/025 | Survey on user requirements for potential ionospheric products from EPS-SG radio occultation measurements |
| SAF/ROM/METO/REP/RSR/026 | Estimates of GNSS radio occultation bending angle and refractivity error statistics |
| SAF/ROM/METO/REP/RSR/027 | Recent forecast impact experiments with GPS radio occultation measurements |
| SAF/ROM/METO/REP/RSR/029 | Testing reprocessed GPS radio occultation datasets in a reanalysis system |

ROM SAF Reports are accessible via the ROM SAF website: http://www.romsaf.org