The referee did not point out any error in my paper; instead he evaluated it against his predictions as to how my proposed model would fare in the future, and against his misreading of several of my statements. If my paper is not published, then the prediction of the referee ("I do not think it will be possible to resolve these issues...") will fulfil itself automatically, without exposing my proposal to any kind of public scrutiny. This would be unfair.

Another general remark: what I propose in my paper is not a fully worked out single model of a GRB source, as the referee seems to assume, but a third step in a sequence of refinements of the model that I proposed in my Ref. [1]. To recall: it is one of 3 models presented there, the one that modeled the lowest-energy GRBs. That model correctly accounted for the existence and energy those GRBs and their afterglows, and nearly correctly for their angular diameters, but not for the duration of the bursts and afterglows or for their collimation. The second paper [2] proved that the collimation is present in the Szekeres spacetimes. The paper now submitted includes the collimation and improves on the implied angular sizes, but they are still nearly twice as large as the direction uncertainty of the arriving GRBs, so further improvements are needed. The durations are not yet dealt with, but I am now completing the next paper, in which they are explained via the cosmic drift effect (see my ref. [13] for a definition, and see below for more on this). Further improvements should result from refining the Big Bang (BB) profile so that the same range of GRB energies is obtained with a lower and narrower hump on the BB, leading to a still smaller angular diameter of the source. The current BB profile and the Szekeres deformation imposed on the L-T background resulted from a blind search for improvements within a 6-parameter class of profiles, see my Sec. XII (second paragraph). This class, labelled by 6 arbitrary constants, was still limited in its generality because it consisted of handpicked representatives of 2 arbitrary functions of one variable. A further limitation was that I assumed the energy function, E(r), to have the Friedmannian form $E = kr^2$, with constant k, throughout space, with a pre-assumed value of k, while in general E(r) is also an arbitrary function. The original profile that initiated the improvements was just guessed (it had to keep the light ray in the blueshift-generating region for sufficiently long time). All of these limitations are emphasized in my Sec. XII. It is highly unlikely that I could have spotted the best-of-all BB profile in this way. This is why I wish to expose my proposal to public scrutiny - the way of progressing to further improvements may be invented by someone else or by myself later.

Now I come to the specific statements of the referee. I apologize for the repetitions in my reply; they resulted from the complicated structure of the referee's arguments.

In his first sentence he says that my proposed mechanism of GRB generation is based on blueshifting the CMB. As I emphasized in Ref. [1] (see page 7 and Appendix B there), what is blueshifted are isolated rays emitted during the last scattering epoch, simultaneously with the CMB. The blueshifted rays must be radial (in L-T models) or axial (in axially symmetric Szekeres models), so they are not members of any specific spectrum - see below for more on this.

In the second sentence the referee says that my paper "includes new calculations to resolve some technical issues...", etc. The paper includes a definitive solution of one problem that Ref. [1] could not deal with: the collimation of GRBs into narrow beams. The collimation is necessarily present in the Szekeres model now used, but could not be explained in the L-T model used in Ref. [1]. This is not merely a technical issue but a substantial improvement. (This fact had already been demonstrated in Ref. [2], but on illustrative examples that were not related to cosmology.)

In his second paragraph the referee says that my paper appears to "grossly contradict a number of well-established observational facts". This is again not true. In Ref. [1] I listed 6 directly observed properties of the GRBs and I am striving to explain all of them by consecutive improvements of the model. I do not see at which point I contradicted any fact. I may have contradicted some model-dependent interpretations of some observational facts, but one should not evaluate one model using implications of another model as a criterion. It is in this paragraph that the referee dismisses my paper because of what he thinks would happen in the future.

In the referee's point (1), the following statements need to be corrected:

(1a) The QSS regions being large is not a necessary property of the class of models that I investigate. One of the aims of my present paper and of the planned future ones is to make them smaller.

(1b) Those regions are not placed "close to the last-scattering surface of the observer". The Szekeres region is a world-tube extending from the BB to the present time, and the last-scattering hypersurface (LSH) runs across it, just above the BB. (Besides, why "of the observer"? The LSH is an observer-independent object in spacetime.)

(1c) I have nowhere implied that there would be "large holes" in the CMB. Firstly, when a GRB arrives, it is superposed on the CMB sky. So, there is a "hole" in the CMB in those directions that are eclipsed by the GRB source, whether it lies at 5 Gyr from us (as is generally agreed by astronomers) or at 13 Gyr as in my model. Since the directions from which the GRBs come are determined only up to a disk of 1 degree diameter in the sky, this disk covers the "hole" that will have to be monitored during the GRB flash - we do not know what happens with the CMB rays that should come from those directions in that time interval. Secondly, the nearly-2 degree angular diameter of the "hole" implied by my model is not the ultimate implication of the whole class of models, but an intermediate result that needs to be improved, as I have said in several places of my papers (and also above). The way to achieve such an improvement is to modify the profile of the BB hump so that it yields the same GRB energy while being lower or narrower (modifying the Szekeres deformation superimposed on it should also help). The reason why I have not so far succeeded in making such an improvement is that each modification of the hump profile must be followed by numerical integration of a past-directed null geodesic running close to the hump (in order to verify the effect of the modification), and the numerical step must be small to keep the calculation sufficiently precise. Each numerical run thus takes much time. The limited numerical capabilities that I have at my disposal do not allow me to take full advantage even of the models I have already tested. Perhaps they are more efficient in generating blueshift than I could demonstrate, and this is one more reason why I wish to publish my current results. The referee's statement "I see no way that such large, extreme lenses could fail to seriously disrupt the CMB, even if the model were substantially modified." is premature, and refers to his prediction of the future. One cannot be sure of such predictions until they are verified, and I hope my paper will become an invitation to verification for other authors - maybe someone else will be able to do what I was not. The "hole" most probably can be made smaller, and then the "disruption" implied by the model will become comparable to the actual one that appears when a GRB is seen by a detector.

(2a) As I emphasized in Ref. [1] (page 7 and Appendix B there), my proposed mechanism blueshifts each axial ray separately. Only those rays are blueshifted that proceed along the preferred axis (in the Szekeres model) or are radial (in the L-T model). During the last scattering epoch, at a given non-central location at any instant there is only one such ray, so it has no chance to become part of the blackbody spectrum. The CMB consists of those rays that were not blueshifted. I know it is unrealistic to assume that every reader had read my previous papers on this topic, but I cannot repeat in every new paper all the things that I had already said.

(2b) That the durations of the GRBs and afterglows are not correctly reproduced using the L-T model I had already stated in Ref. [1]. The Szekeres model offers a way to deal with the durations correctly: if there is another Szekeres region between the BB hump and the observer, then the blueshifted rays running through it are deflected and the angle of deflection depends on time in consequence of the cosmic drift (see my Ref. [13]). Consequently, the maximally blueshifted rays will miss the observer after a while. For some time after that the observer will be receiving the less-blueshifted rays of the afterglow. I am now preparing a paper in which I work out a description of this process. The referee essentially demands that I solve all possible problems of my model in a single paper. This is both unrealistic (the paper would be very long, and it would take a very long time to prepare it) and unfair (does one require the same from all other papers that introduce new ideas?).

(3a) The objection of the referee concerning the pointing of the blueshift axis and covering the whole sky applies just as well to the actually observed GRBs. They are believed to be collimated into narrow beams, so they can be seen on Earth only if the beam is pointed at the Earth. After 27 years of BATSE observations, the number of detected GRBs exceeds 8000. With the resolution of 1 degree in diameter, about 41 000 observed GRBs will cover the whole sky (see my section XI). Assuming the rate of new detections will continue, this will happen in about 100 years from now, and then what? Newly detected GRBs will be coming from the same directions, from which other GRBs had already been registered. They all will be only a small subset of those that did occur, since a lot more were not aimed at the Earth. My model does not have this problem: the GRB sources (the humps on the BB) are still there, but they come into and out of view of the observer - see (2b) above. So, if the referee is not worried by the multitude of the observed GRBs, then this particular objection against my model does not apply.

(3b) The objection concerning "essentially covering the whole sky" applies only to the present form of my model. See (1c) above.

(3c) "why do we not see such regions at lower redshift?" - I suppose by "such regions" the referee means Szekeres regions. Of course we see them - they are voids and condensations in the distribution of galaxies. They do cause perturbations in the CMB temperature and were discussed in many papers (the pioneer and most prolific author of such papers was K. Bolejko, another name to search in the literature is R. Sussman). But they do not generate any blueshift for the reason explained in (3f) below.

(3d) "There are optical, IR, and UV backgrounds that could also be blueshifted into the gamma band." - Local blueshift is generated only below the extremum redshift surface (ERS), i.e. when the rays aimed at the observer are sufficiently close to the BB - see my section VII. Later, the rays acquire local redshifts that will, in most cases, obliterate the blueshift. The technical difficulty in my class of models is to ensure that the initially acquired blueshift is strong enough to survive the later redshifting. If the optical, IR, and UV rays mentioned by the referee are emitted later than the ERS, then they are not further blueshifted (but they may have resulted from earlier blueshifting).

(3e) "why should all the QSS regions have profiles that are so finely tuned to give such extreme blueshifts?" - Here the referee puts my reasoning head down. Inhomogeneous cosmological models imply that some rays emitted close to the BB will be blueshifted. Consequently, if these models are to be credible, we must find a place for blueshifted rays in the observed Universe. My hypothesis is that they are seen as the GRBs (possibly as only some of the GRBs). My Ref. [1] was a proof of existence of the mechanism that can blueshift the hydrogen emission radiation sufficiently strongly that it survives the later redshifting during almost the entire lifetime of the Universe and is presently seen in the gamma range of frequencies. It is incorrect to claim that the QSS regions are "so finely tuned" - see also below, in (3, the second paragraph). My model is just the first example, it remains to be investigated whether other models exist and how numerous they are. This model resulted from a series of improvements in a continuous family of models, and further modifications of it should be tested - but they will not be if my paper is blocked from public view.

(3f) "We might expect to see a population of objects with less extreme blueshifts in the X-ray or UV, for example.", and the next sentence. - This is a rewording of the objection in (3c), and I have already answered it: the less extreme blueshifts are present in the afterglows. The QSS regions lying closer to us do not generate blueshifts because the rays presently reaching the Earth went through them high above the ERS. Those regions are observed as voids or condensations.

(3, the second paragraph) My model is just a single member of a possibly large class that was not yet explored (see my second paragraph above). I have nowhere said or implied that it is unique; on the contrary - I have indicated in several places that it is an intermediate stage in a series of improvements. The accusation that the model is extremely fine-tuned is thus baseless - it must be verified what other models, and how many of them, can achieve the same. Besides, this accusation is a logical error - one could call "extremely fine-tuned" every theoretical model whose free parameters were fitted to experimental data. Using the referee's language, this argument of his is itself "extremely fine-tuned".

(4a) In view of what I said above, the first paragraph of section (4) of the report is incorrect. The model requires further refining, but some observational tests will become possible even in its present form in (perhaps near) future. For example, the model predicts a certain pattern of redshift across the observed image of a GRB source (see tables in my section IX), and these may be compared with observations as soon as the GRB detectors become able to resolve the images.

(4b, the paragraph after referee's (4)) Here the referee repeats the incorrect claim to which I have replied in (3, second paragraph) above. My model is a representative of a class whose extent is not yet known, and the values of its parameters were adapted to the lowest energy of the observed GRBs. Ref. [1] presents two other sub-classes of the same class, one of which is adapted to the highest-energy GRBs, and the other allows in principle to control the duration of the GRBs. By contrast, the FLRW models are severely limited in their generality even before the adaptation to observations begins.

(4c) The remark about "plausible mechanisms" does not invalidate my argument. It is one thing to assume an inhomogeneous model and show by calculations that it is "driven to a locally almost-FLRW solution", and another thing to assume that the Universe is already in the FLRW class having been previously driven to it, as most astronomers do. The first is legitimate science, the second is following a belief without verification.

(4d) "There is also the issue that the class of models usually used to describe the Universe on large scales is really perturbed FLRW, not FLRW in the strict sense." - The L-T and Szekeres models are classes of just such perturbations. They differ from those considered in the standard perturbation theory by being exact solutions of Einstein's equations. In consequence of the exactness, when using them we do not have to worry about going beyond the "linear regime": as long as we trust that general relativity is the correct theory of gravitation, the exact models can be extended arbitrarily far into the future. I am puzzled why most astronomers violently object against exact inhomogeneous models while accepting that FLRW require perturbations.

(4e) I will remove the parenthetical remark about inflationary propaganda if this is a necessary condition for accepting my paper for publication.

In conclusion, I ask the referee to revise his verdict. I may put some of the explanations given above into the text of the paper, if thereby it becomes clearer.

As a final remark: it is not my intention to force my proposed mechanism as the explanation of all the observed GRBs. From the point of view of relativity theory it suffices if only some of them are generated in this way. So, the observers would only need to verify whether some of the observed GRBs have features that are consistent with my model. If this is the case, then I will have identified a new subclass of GRBs. But, once more, this will never happen if I am not given a chance to present my paper to the public.