Teru Miyake **November 2010**
Stanford University

Underdetermination in Geophysics

Teru Miyake
Stanford University
PSA 2010, Montreal

Abstract: This paper examines the epistemological implications of a particular underdetermination problem from geophysics, with an emphasis on understanding how the scientists themselves tried to deal with the problem. The problem is from the highly influential work of the geophysicists Backus and Gilbert in the late 60's, who were trying to determine the internal structure of the Earth using seismic waves. I find that actual underdetermination problems can be vastly complex, with different sources of underdetermination having different epistemological implications. A better understanding of actual cases of underdetermination is needed before we can make epistemological conclusions based on underdetermination.

Word Count: 4,964 Words

1 An Underdetermination Problem from Geophysics

This paper will examine a specific underdetermination problem from geophysics. The particular underdetermination problem I am interested in is the problem of determining the internal structure of the Earth, given a certain set of observations at the surface of the Earth. I will be examining in depth the seminal work of the geophysicists Backus and Gilbert in the late 1960's to early 1970's. They identify the underdetermination problem for a particular geophysics problem, make the underdetermination precise by laying it out in an elegant mathematical framework, and then provide methods for dealing with this underdetermination by separating out various different sources of underdetermination.

This is an interesting case study for understanding the epistemological implications of underdetermination for four reasons. First, it is a case where underdetermination was recognized explicitly by scientists to be a significant problem, and methods were developed for dealing with underdetermination. Second, it is a problem in which various different sources of underdetermination are linked together in a complicated way, and we can see how scientists explicitly recognized these sources and tried to separate them out. Third, in this particular case, since we are talking about the internal structure of the Earth, it is less tempting to become an anti-realist, so that the problem of underdetermination can effectively be considered independently from the problem of scientific realism. Fourth, this particular underdetermination problem was highly influential—it led to the development of more general methods for addressing underdetermination (inverse problem theory), and it was a crucial step in the development of a hugely important model in geophysics.

In this paper, I cover very briefly the history of geophysics leading up to the work of Backus and Gilbert. I then describe in detail the underdetermination problem with which Backus and Gilbert were concerned, and the methods which Backus and Gilbert used to deal with the underdetermination. I then consider briefly the implications of the work of Backus and Gilbert for philosophical work on underdetermination.

2 Travel Time Inversion

The specific problem I will be interested in is the problem of finding the internal density distribution of the Earth, given observations at the surface. This is a problem on which we have made very good progress since the beginning of the 20th century, when geologists first began to use observations of seismic waves in order to gain information about the Earth's interior. The history of seismological research into the internal structure of the Earth in the 20th century can be divided into two periods, according to the kind of seismological observations that were available.

The earlier period is characterized by the use of observations of travel times of body waves through the Earth. When an earthquake occurs, two types of waves, P waves and S waves, travel through the Earth's interior and can be detected at seismographic receiving stations at various points on the Earth's surface. P waves are compression waves, involving a dilatational motion of the medium through which they pass, while S waves are rotational waves, involving a circular, non-dilatational motion of the medium. Assuming perfect elasticity and isotropy of the medium, wave equations can be derived and solutions can be found for both types of waves (Bullen, 111-113). According to these solutions, the speed of P waves depends upon the incompressibility, rigidity, and density of the medium, while the speed of S waves depends upon the rigidity and the density. Wherever there is a sharp discontinuity in the density, both types of waves will be reflected and refracted, analogously to the reflection and refraction of light waves. Since the density within the Earth also varies continuously, S waves and P waves will travel along curved paths, where the specific paths taken will depend upon the density, incompressibility, and rigidity of the medium through which they have traveled.

If we assume that the internal structure of the Earth is spherically symmetric, then the times for specific seismic waves[1] to travel specific distances (measured in degrees along the surface of the Earth) will be constant.The internal structure of the Earth was determined by compiling such travel times in the form of tables—important work was done here by Jeffreys and Bullen at Cambridge and New Zealand, and Gutenberg and Richter at Caltech.

[1] Seismic phases, to be more specific.

3 Normal Mode Inversion

In the 1960's, it became possible to bring a new type of seismological observation to bear on the problem of determining the internal structure of the Earth—observations of the normal modes of the Earth. The motion of an idealized mathematical string with tension τ and density ρ can be approximated by a wave equation, and with the addition of certain boundary conditions (such as fixing the ends of the string), only certain functions are possible solutions to this wave equation. These functions are called eigenfunctions, and the frequencies corresponding to the eigenfunctions are called the eigenfrequencies. Arbitrarily shaped waves traveling along the string can be represented as linear combinations of these eigenfunctions using the theory of Fourier analysis.

Although the mathematics is more complicated, essentially the same thing can be done for the vibrations of an idealized three-dimensional object. In particular, we can think of an idealized Earth as a spherically symmetric, isotropic, perfectly elastic, non-rotating body. The free vibrations of such a body will have eigenfunctions, often called 'normal modes', which in this case are complicated three-dimensional functions involving Legendre polynomials. It turns out there are two types of normal modes—torsional, or toroidal modes, which involve no radial movement, and spheroidal, or poloidal modes, which combine radial and transverse motions.

For an idealized Earth that is spherically symmetric, isotropic, perfectly elastic, and non-rotating, and given its rigidity, incompressibility, and density as functions of radius, the frequencies of the normal modes can be calculated. This problem is in principle solvable uniquely and exactly, although numerical integration may be required. The problem that geophysicists are interested in, however, is the inverse of this problem. The frequencies of the normal modes are observable, and we want to find out the internal structure of the Earth given these frequencies. The observation of the frequencies of the normal modes is by no means a trivial matter, but in order to make my point about underdetermination, I will simply assume that the geophysicist can observe the frequencies of the normal modes exactly. The problem geophysicists faced, then, was determining the internal structure of the Earth, given the frequencies of the normal modes exactly, and a geophysical theory that allows you to calculate the frequencies of the normal modes given a certain internal structure.

How can one go about solving such a problem? One natural way of thinking about solving this problem is as follows. You (1) construct a model of the Earth; you then

(2) calculate the normal mode frequencies you would expect to see for such an Earth; then you (3) compare these calculated frequencies with the observed frequencies; and (4) if they agree, you take that model to be confirmed. If they don't agree, then you change the model, and repeat steps (2) through (4).

Of course, we immediately run up against a problem that is familiar to philosophers of science. The procedure given above can be thought of as a kind of hypothetico-deductive inference. You hypothesize a certain model of the internal structure of the Earth, then you deduce its observable consequences, and you compare those consequences with actual observations. You take the hypothesis to be confirmed if the consequences of your hypothesis match up with the observations. The problem with hypothetico-deductive inference is well-known. Even if we find that the consequences of a hypothesis match observations, that does not mean the hypothesis is true, because there could be other hypotheses, known or unknown, that would match those observations just as well. Much has been written about hypothetico-deductive inference, and the issue of choosing between hypotheses in particular, in the philosophical literature. A strongly anti-realist position would be to say that we simply cannot say whether any of these hypotheses are true—all we can say is that they are empirically adequate. Others might use some kind of criteria (so-called supra-empirical virtues), like simplicity, to choose between hypotheses. In the following few sections, I will examine what geophysicists have actually done when faced with this situation.

4 Backus and Gilbert

The geophysicists George Backus and Freeman Gilbert wrote a seminal series of papers (1967, 1968, 1970) that explicitly considers the underdetermination of the internal structure of the Earth. Backus and Gilbert call data which are taken to correspond to features of the Earth as a whole 'gross Earth data'. These include the mass of the Earth, the moment of inertia, travel times of P-waves and S-waves, frequencies of the normal modes, quality factors of the normal modes, and so on. Since we can only observe a finite number of these gross Earth data, and the internal structure of the Earth presumably has a huge number (practically infinite) of degrees of freedom, it seems hopeless to try to determine the internal structure exactly. The natural question to ask, given this limitation, is an epistemological one:

> Human limitations are such that at any given epoch only a finite number of gross Earth data will have been measured. This paper is a discussion of the extent to

> which these finitely many gross data can be used to determine the Earth's internal structure. (Backus and Gilbert 1967, 247)

If we simply consider the number of observations we can possibly make, and the number of degrees of freedom of the internal structure of the Earth, it looks like the latter number is much greater than the former, so we will have a hopeless underdetermination. But just how hopeless is it? Might there be ways of dealing with this underdetermination?

The first step towards answering these questions is to make the underdetermination problem more precise. We must make explicit (1) what is being underdetermined, (2) what is doing the underdetermining, (3) what, exactly, the relation of underdetermination is. Backus and Gilbert set the problem up as follows:

> The problem we shall consider is the following: suppose a non-rotating, spherical, isotropic Earth of radius a has density $\rho(r)$, bulk modulus $\kappa(r)$, and shear modulus $\mu(r)$,[2] all functions only of r, the radial distance from the center. Suppose a finite number J of gross Earth data $\gamma_1, \gamma_2, \ldots, \gamma_J$ have been measured, and that these data depend only on the functions ρ, κ, μ. Given the observed values of $\gamma_1, \ldots, \gamma_J$, what can be said about the unknown functions ρ, κ, μ?[3] (Backus and Gilbert 1967, 248)

In terms of the three questions given above: (1) they take the density, bulk modulus, and shear modulus of the Earth as a function of radius to be underdetermined; and (2) they take

[2] Bulk modulus is a measure of incompressibility, and shear modulus is a measure of rigidity.
[3] Some of the notation has been changed for consistency with other parts of this paper.

these functions to be underdetermined by some number of gross Earth data, which is finite.

Backus and Gilbert then answer the third question (3) by coming up with a formal mathematical relation between the density, bulk modulus, and shear modulus on the one hand, and the gross Earth data on the other, using the theory of linear differential operators. According to the characterization of the problem given above, the Earth is assumed to be completely specifiable by the three functions $\rho(r)$, $\kappa(r)$, and $\mu(r)$. Thus, an Earth model can be represented by an ordered triple of real-valued functions defined on [0, 1] (normalizing for the radius of the Earth), $\mathbf{m} = (\rho, \kappa, \mu)$. Linear combinations of Earth models can be defined in a straightforward way: $a\mathbf{m_1} + b\mathbf{m_2} = (a\rho_1+b\rho_2, a\kappa_1+b\kappa_2, a\mu_1+b\mu_2)$. We can then think of Earth models as points in an infinite dimensional linear space M, the space of all possible Earth models. A natural inner product[4] can be defined on this space, and it can be completed to form a Hilbert space (the space $L_2[a,b]$).

In order to remove an ambiguity in the term ‘gross Earth data’, we distinguish between ‘gross Earth data’ and ‘gross Earth functionals’. Gross Earth functionals are simply real-valued functions g_1, g_2, … g_n on the space M of all possible Earth models.

[4] The inner product of $\mathbf{m_1}$ and $\mathbf{m_2}$ is defined to be the integral from $r = 0$ to 1 of $\rho_1\rho_2 + \kappa_1\kappa_2 + \mu_1\mu_2$. The choice of inner product is somewhat arbitrary; what is important at this point is that a Hilbert space can be constructed.

There are gross Earth functionals corresponding to the Earth's mass, its moment of inertia, its normal mode frequencies, and so on. The actually observed values of these gross Earth functionals, which we designate by γ_1, γ_2, … γ_n, and which are taken to be the values for the 'real' Earth, are then called 'gross Earth data'. Recall, for example, that the normal mode frequencies depend entirely on the internal structure of the Earth. So the frequency of each normal mode can be taken to be a function on the space M of possible Earth models. That is, for each normal mode *i*, there is a function g_i that associates, to each point in this space M, a real number representing the frequency of that particular normal mode. And each point in this space M of course represents one possible internal structure of the Earth, or one possible Earth model. The actually observed values of the frequency of the *i*th normal mode is given by γ_i.

We now have a precise way of stating the underdetermination. Given a certain set of observed gross Earth data γ_1, γ_2, … γ_n, can we pinpoint a single model $\mathbf{m}_E$ which we take to correspond to the 'real' Earth in the space M? Suppose, for the moment, that the observed gross Earth functionals are linear functions. This is not the case for the normal mode frequencies, and in fact, most gross Earth functionals are nonlinear functions. There are some linear gross Earth functionals, however, like the quality factors of the normal modes. I am doing the linear case first because the linear case must be understood in order to understand the nonlinear case. According to the Riesz representation theorem, every bounded linear functional $L(\mathbf{f})$ on a Hilbert space may be written as the inner product of $\mathbf{l}$ and $\mathbf{f}$, $\langle \mathbf{l}, \mathbf{f} \rangle$, where $\mathbf{l}$ is a point in the Hilbert space uniquely determined by the functional L (Parker 1994, 31-32).[5] Thus, every gross Earth functional $g_i(\mathbf{m})$ has associated with it a point $\boldsymbol{G_i}$ in the space M, such that

$$g_i = \langle \boldsymbol{G_i}, \mathbf{m} \rangle.$$

Backus and Gilbert call the point $\boldsymbol{G_i}$ associated with each gross Earth functional the 'data kernel' of that gross Earth functional (Parker (1994) calls this the 'representer'). Take a particular Earth model $\mathbf{m}$, which is a point in the space M. The value of a gross Earth functional g_i for that particular Earth model $\mathbf{m}$ is given by the inner product of $\mathbf{m}$ and the data kernel $\boldsymbol{G_i}$ of that gross Earth functional.

Now, suppose we are looking for a point $\mathbf{m}_{\boldsymbol{E}}$ in the space of possible models, which we take to represent the 'real' Earth. Suppose we only have a finite number of gross Earth data, say a thousand of them. We are in effect looking at the projection of $\mathbf{m}_{\boldsymbol{E}}$

[5] My explanation here follows that given in Parker 1994, since I find it much cleaner, but I have made the notation and terminology consistent with Backus and Gilbert.

onto the subspace A of M spanned by the data kernels $\boldsymbol{G_1}$, …, $\boldsymbol{G_{1000}}$ of the gross Earth data. But this space would only have a thousand dimensions, whereas the space M is infinite dimensional. The upshot is simply that there is an infinite number of models that will agree with the data just as well as $\mathbf{m}_E$, and hence are observationally indistinguishable from $\mathbf{m}_E$. In fact, the *space* of such observationally acceptable models is infinite dimensional. The underdetermination looks pretty bad, to say the least!

5 The Resolution Method

But maybe it's not so bad. This is actually something we already knew about, although we did not have a precise characterization of this underdetermination. For the internal structure of the Earth is something that we postulated at the beginning to be something with infinite degrees of freedom—the way we defined $\rho(r)$, $\kappa(r)$, and $\mu(r)$, these can be arbitrary functions of radius. And therefore we cannot hope to pin down a unique internal structure given a finite number, no matter how large, of gross Earth data. A more interesting question for the geophysicist is whether there are observationally equivalent models that are significantly different from a geophysical standpoint, as Backus and Gilbert point out:

> With only finitely many gross data we cannot expect to resolve details of arbitrarily small vertical scale; our vertical resolution is finite. This remark is sufficiently tautological as to be without geophysical interest. A geophysically more interesting question is whether there is any other source of non-uniqueness besides the finite resolving power inherent in a finite set of gross Earth data. We shall see in general there is. (Backus and Gilbert 1967, 251)

An infinite number of models will be consistent with the data—I will call such models 'observationally acceptable'. An infinite number of observationally acceptable models will differ from each other only in their fine-scale structure, but this is not important—if we are doing geophysics, we do not care about models that differ only on the scale of millimeters. The *real* question is, could there be observationally acceptable models that differ from each other in geophysically significant ways? And if there are, in what way do they differ? A further thought is that if all of the observationally acceptable models have certain geophysically significant features in common, then we can conclude that the 'real' Earth also has those features.

> If we are fortunate or shrewd in our choice of which gross data to measure, then all the different [observationally acceptable] Earth models may share some common properties. For example, they may all have a low-velocity zone in the upper mantle; or they may all become essentially the same when we take running averages of their ρ, κ, and μ over some fixed depth interval H. In the first example, we can definitely assert that the Earth has a low-velocity zone in the upper mantle.

> In the second example, we can claim to know ρ, κ, and μ as functions of radius *r*, except for unresolved details whose vertical length scale is H or less. (Backus and Gilbert 1967, 249)

This will allow us to take what looks like a hopeless underdetermination problem, and draw conclusions about the real world based on it. But how could we go about doing this? Suppose we happen to find an observationally acceptable Earth model, and it has certain features, such as a low-velocity zone in the upper mantle. Is this a feature that is particular to this one model, or is it a feature that all observationally acceptable Earth models share? One way of answering this question is simply to generate a large number of observationally acceptable Earth models and *see* if they all have this feature. But this procedure seems to be a rather inefficient way of going about answering this question. A further worry is that the simple fact that all the models you have generated *so far* have this feature, does not mean that all observationally acceptable models have this feature.

Now, another way of answering this question is to find out how much resolving power you have. Suppose your model has a feature near a point r_0. If that feature is smaller than the resolving power of the data, then you can conclude that this feature is an artifact of the model you happened to find. In (1968), Backus and Gilbert provide a method for finding the vertical resolving power of a given set of gross Earth data. The method they developed is called the Backus-Gilbert resolution method.

Suppose we are interested in how much we can know about the values of $\rho(r)$, $\kappa(r)$, and $\mu(r)$ for the 'real' Earth, given a set of gross Earth data $\gamma_1, \ldots, \gamma_n$. As I mentioned above, each gross Earth functional g_i can be written as the inner product of its data kernel $\boldsymbol{G_i}$ and a point in M representing an Earth model. Thus, the gross Earth data γ_i can be taken to be the lengths of the projections of the point representing the 'real' Earth $\mathbf{m}_E$ onto the data kernels $\boldsymbol{G_i}$. Now we ask how much vertical resolution we can expect to get at some point r_0. We will only get as much resolution as can be discriminated by the data kernels $\boldsymbol{G_i}$. To get some idea of the resolution, then, we try to construct, out of the data kernels $\boldsymbol{G_i}$, a function that gets as close as possible to a delta function at r_0. The criterion for delta-ness is somewhat arbitrary—Backus and Gilbert choose one that is convenient numerically. The aim of the procedure is, as I mentioned, to give the geophysicist some idea of what the vertical resolution around some point is, so the delta-ness criterion can be decided upon on a practical basis. If you cannot get a function that is close to a delta function around some point r_0, you can conclude that there are observationally acceptable models that have large-scale features that differ near that point.

6 The Nonlinear Case

I said above that most gross Earth functionals are actually nonlinear, including the frequencies of the normal modes. I postponed discussing the nonlinear case, since the linear case is simpler. For nonlinear gross Earth functionals, you must first find an observationally acceptable model $\mathbf{m_0}$. You then make the assumption that all other observationally acceptable models are sufficiently similar to this model that a linear approximation holds (we assume that the gross Earth functionals are Frechet differentiable), namely, that

$$f_i = g_i(\mathbf{m_1}) - g_i(\mathbf{m_0})$$

can be approximated by a linear function on the space M. You can then proceed with the resolution method exactly as in the linear case, using f_i instead of g_i. You are, in effect, exploring the part of the space of possible models M that is near the reference model $\mathbf{m_0}$. There is no way, however, of telling whether there are radically different models that are so far away that a linear approximation won't work. The upshot is that if the observational data consists of linear functionals, one can use the resolution method in order to tell whether there are observationally acceptable models that differ significantly from each other, or whether all observationally acceptable models differ only in fine-scale detail. If the observational data consists of nonlinear functionals, however, as is the case with normal mode frequencies, you have to make the assumption that there are no models that are so far away that a linear approximation won't work, in order to use the resolution method. And the only way of determining whether there are such models is to attempt to construct such models and see if they agree with observations. Given the appropriate computational power, one can systematically construct models and test them using Monte Carlo methods, but since the space of possible models is infinite dimensional, there are limitations to such methods, and drawing epistemological conclusions from them can be rather risky.

7 Philosophical Implications

The rest of this paper will briefly consider the work of Backus and Gilbert from the standpoint of philosophy. When faced with a situation where there are multiple models underdetermined by observation, you might simply deny that we have any right to choose what the 'true' model is like. This is particularly tempting for philosophers who have anti-realist tendencies. Of course, most philosophers are, I presume, rather reluctant to be anti-realists when it comes to talking about the internal structure of the Earth. We think there is a fact of the matter about what it's really like down there, and we want to find out as much as possible. A sophisticated understanding of the underdetermination problem gives us the tools to do something more than simply throwing our hands up and becoming anti-realists.

Backus and Gilbert provide us a way of sorting out the sources of underdetermination. Some of them can be dealt with, and some can't. The first step is to sort them out. Backus and Gilbert identify three sources of underdetermination. The first source of underdetermination is the finite resolving power of the data. The second source of underdetermination is the fact that radically different models can agree with observation. The third source of underdetermination is observational error. Now, the first source of underdetermination cannot be helped, since we are using a finite number of data to try to pinpoint a model that has infinite degrees of freedom. But this type of underdetermination is not that significant to the geophysicist, as long as the underdetermination is between models that only have differences on a small scale. The second type of underdetermination *is* significant, however, since if the second type of underdetermination obtains, there could be observationally acceptable models that differ significantly from a geophysical standpoint. Backus and Gilbert's resolution method is a way of telling the difference between the first type and the second type of underdetermination. Now, as mentioned above, the method is guaranteed to work for the linear case, but it will only work for the nonlinear case provided that certain assumptions hold. Backus and Gilbert (1970) give a method for carrying out the resolution method, taking into account observational error. They show that if there is observational error, there is a tradeoff between resolution and accuracy, and they provide a way of optimizing this tradeoff.

Backus and Gilbert do *not* provide an explicit method for dealing with a fourth source of underdetermination. All of the discussion of underdetermination so far presupposes that we have the right set of assumptions, and the right geophysical theory. But we might not have the right set of assumptions. It turns out, for example, that early models based upon normal mode inversion were making an incorrect assumption—they were assuming that velocity dispersion due to anelasticity of the medium would be negligible. We have also assumed that the medium is isotropic, but this is also, of course, an idealization. If the medium is not isotropic, then the constitutive equations of the medium, from which we derive the wave equations, will look rather different. In fact, we know that there cannot be a model corresponding to the 'real' Earth in the space M described above, since all models in that space are spherically symmetric, perfectly elastic, and isotropic, whereas the Earth is neither spherically symmetric, perfectly elastic, nor isotropic. What we have been taking to be the space of all possible Earth models has turned out not to be the space of all possible Earth models after all. How, then, can we be sure that the techniques we have been using are giving us information about the Earth at all? This is closer, I think, to the kind of underdetermination that is usually discussed by philosophers. I try to provide a detailed answer to this question elsewhere. For now, I will just say that the methods of Backus and Gilbert can at least help in sorting out the other types of underdetermination from this kind of underdetermination.

In conclusion, I want to make two points. First, I think that in general, philosophers have vastly underestimated the complexity—and the richness—of actual underdetermination problems. The emphasis in the philosophical literature on contrived or artificial examples of underdetermination makes it too easy for opponents of underdetermination to say that underdetermination is trivial. Actual underdetermination problems are complex—and philosophy of science can benefit from studying them in detail.

Second, I think that dealing with underdetermination involves many practical tradeoffs, and in order for us as philosophers to understand what's going on, we need to consider underdetermination in the context of ongoing research. The study of underdetermination in actual cases, then, will lead naturally to a deeper understanding of the methods and aims of scientists in ongoing and extended research programs.

REFERENCES

Backus, G. E. and Gilbert J. F. (1967) Numerical Applications of a Formalism for Geophysical Inverse Problems. *Geophysical Journal of the Royal Astronomical Society*, 13: 247-276.

Backus, G. and Gilbert, F. (1968) The Resolving Power of Gross Earth Data. *Geophysical Journal of the Royal Astronomical Society*, 16: 169-205.

Backus, G. and Gilbert, F. (1970) Uniqueness in the Inversion of Inaccurate Gross Earth Data. *Philosophical Transactions of the Royal Society of London, Series A, Mathematical and Physical Sciences*, Vol. 266, No. 1173: 123-192.

Bullen, K. E. (1975) *The Earth's Density*. London: Halsted Press.

Lanczos, C. (1997) *Linear Differential Operators*. Mineola, NY: Dover.

Parker, R. L. (1994) *Geophysical Inverse Theory*. Princeton University Press.

Riesz, F. and Nagy, B. (1990) *Functional Analysis*. Mineola, NY: Dover.

Tarantola, A. (2005) *Inverse Problem Theory and Methods for Model Parameter Estimation*. Philadelphia: SIAM.