

©David Lucy - Don Ortner, 2001 - cite as:

Lucy, D. and Ortner, D. 2001 Immune Response Modelling - unpublished informal Report.

Immune Response Modelling - Informal Report.

David Lucy¹ and Don Ortner

1 Introduction

The model elaborated upon here was devised by Don², and was published in a Cambridge volume by Grauer and Macadam. It was Don himself who suggested I take a look at it to see what I could make of it.

The model itself is based on the empirical observation that males seen in the palaeopathological record seem to have a much higher degree of skeletal manifestation of disease. The crucial link realized by Ortner is that there is some underlying male/female sex differences which can lead to this. Ortner's primary analytical construct is the immune system. I think for the purposes of this paper what we can imagine the immune system to be is a generalized propensity to resist the changes in bone formation and destruction processes seen in some of the more chronic diseases. I think at this stage seeing the immune system in all its complexity would not be as helpful as a simplified notion of it. Ortner analysis the differences seen in morbidity between males and females in a number of diseases seen in the modern world, then relates these to male/female differences in immune response recorded in the literature.

Ortner's model is revolutionary in some respects, by citing a quantity such as the simplified idea of immune response seen here, and by mixing it with a little more mathematical playing about, some very useful quantities can be calculated. Not only that, but some of what the model does relates to the *osteological paradox* in a more testable way than has been seen before in much of the literature. The *osteological paradox* is an important theory in many respects, principally because it gives some notion³ that the skeletal population doesn't relate directly to the population of diseased humans, but also that it is an explicit palaeopathological theory independent of the parental medical theories.

The reason why I discuss the *osteological paradox* in this document is because this development of Ortner's model implies the existence of the *paradox*. This is why this model is important. It unifies several aspects of current osteological theory into a cohesive whole, whilst allowing a more intimate conceptual understanding of the causative phenomena behind palaeopathological observations.

¹Department of Archaeological Sciences, University of Bradford, Bradford, West Yorkshire, BD7 1DP until 1st March 2001; Department of Mathematics and Statistics, The University of Edinburgh, Edinburgh, as from the 1st March 2001.

²Don Ortner of the Smithsonian Institute

³But no precise values to.

2 Model

2.1 Original model

The original model proposed a main structure of immune response, all other quantities being seen in relation to this major axis. To explain the sex differences seen in skeletal samples⁴ Ortner then proposes that each sex⁵ has associated with it a distribution of immune response to any disease to which it is exposed which is Gaussian, $N(\bar{x}_1, \delta_1)$ and $N(\bar{x}_2, \delta_2)$, where \bar{x}_i is the mean and δ_i variance. The selection of Gaussian functions to describe the frequency distribution is expected and well in order. Ortner then proposes two points R_1 and R_2 on the major axis which define cut off points. Individuals whose immune responses lie below R_1 are interpreted as being killed off by the disease immediately. Those whose immune responses lie above R_2 are said to recover immediately. Both the groups consequently display no signs of the disease, hence the interaction with the *osteological paradox* as people examining the skeletons from these individuals would not be able to tell them apart. Only those with immune response values between R_1 and R_2 would display any skeletal manifestation of the disease.

2.2 Problems with the original model

By careful juggling of where \bar{x}_1 , \bar{x}_2 , δ_1 and δ_2 are on the immune response axis in relation to each other, and where \mathbf{R} is in relation to the other parameters some estimate can be made of the relative numbers of individuals from either sex with skeletal involvement by integrating \mathbf{N}_1 and \mathbf{N}_2 between \mathbf{R} . However, that really is all it can do, and the step function represented by \mathbf{R} has little aesthetic appeal, and no grounding in any physically transcribable mechanism.

I propose a minor modification to this model which will allow a few more quantities to be calculated such as the expected distribution of severity of skeletal affectation whilst retaining the ability to deduce the proportions of those dying from any disease. I also expect to make the elements in the model more meaningful in terms of common conceptions.

3 Modified model simple case

Take a frequency distribution of immune response for a single, but large sample. As before Gaussian is a good choice defined as: $N(\bar{x}, \delta)$. Let us now imagine that for every value of immune response there is an associated probability of death per unit time, and another associated probability of immunity per unit time.

3.1 Probit functions

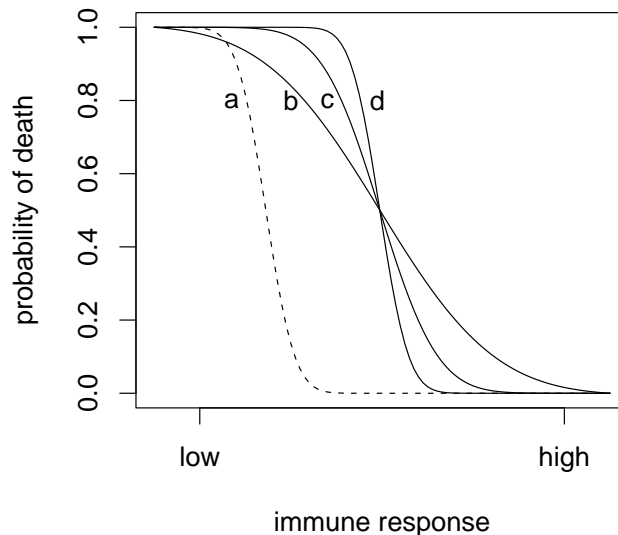
The functions which describe the probabilities of death and immunity are curious. First if one imagines the probability distribution function (pdf) for death then we expect it to be high for those at the worse end of the immune response scale, and low at the higher end. We

⁴Using the term sample here in the statistical sense.

⁵Male and female only please - things just get too mind-boggling if we account for other sexes.

could imagine it to be some step function, but intuitively expect the probability of death to vary inversely in some continuous manner with the immune response. Hence a function belonging to the sigmoidal⁶ family might be expected. For this contribution I have selected a probit⁷ function to represent the pdf of death per time unit as it is easy to control (see Figure 1) for where the $P(d) = 0.50$ point occurs, where P is probability and d is the event of death. Additionally the *shape*, or how spread out the function is, is simmerly easy to control precisely (again see Figure 1).

Figure 1: A small family of probit functions: *a* (short lines) has it's $P(d) = 0.50$ at the lower end of the immune response range otherwise has the same shape parameter as *b* which has it's $P(d) = 0.50$ similarly in the response range. *c* and *d* are likewise centered, but use different shape parameters to *spread* them out a bit. Notice all the distributions featured here are vertically scaled to $[0, 1]$ implying that an individual from the lower range of immune response will stand a very high probability of dying, whereas those further up the scale will have a decreased probability



The function which describes the probability of total immunity for a given immune response is very similar in that it too is a probit, but this time turned around so that the function starts low at the lower end of the immune response scale, and gets higher as we go up the scale. There are probably a couple of points which should be made:

1. The probability functions describing both immunity and death for immune response are for a given disease or condition have be notated as $P(d|\mathbf{x})$ but are more correctly notated as $P(d|\mathbf{x}, c)$ where \mathbf{x} is the immune response and c is the given disease.
2. The two functions $P(d|\mathbf{x}, c)$ and $P(i|\mathbf{x}, c)$ can be varied for c leaving the frequency distribution stationary. This may or may not reflect the real case, it could be that the frequency distributions of immune response may be disease defandant which would be less satisfactory, but doesn't really compromise the model. Experiments later in this report may give some clues as to their behavior.
3. $P(d|\mathbf{x}, c)$ and $P(i|\mathbf{x}, c)$ don't have to be probits, they could just as easily be any sigmoidal function which gave the correct type of form. Unlike the Gaussians used

⁶S shaped.

⁷In fact I used $\text{cumsum}(N(\bar{x}, \delta))$ rescaled to $[1, 0]$, my colleague, Dr. R.G.Aykroyd (Leeds Mathematics and Statistics) thinks this is a probit function.

for the frequency function there is no expectation of any particular shape or family of functions derived from theory.

4. Very important point this: $P(d|\mathbf{x}, c) + P(i|\mathbf{x}, c) \leq 1$, as the death and immunity are exclusive and exhaustive they two probability functions cannot be summed between themselves to be greater than unity.
5. Here I have scaled $P(d|\mathbf{x}, c)$ and $P(i|\mathbf{x}, c)$ to be $[0, 1]$. I'm not absolutely convinced that they have to be as large as 1, so long as both are > 0 and they meet condition 4 above the model should be fine.

3.2 What we can do with these

Let us build in an overall property of immunity in that once any individual is immune then they suffer from no effects, so even were they affected, they survive and the skeleton gets back to normal⁸. This means that we are really only interested in the bit of the sample which succumbed to death, and that we see no manifestation of disease on those which didn't die from the disease⁹. In which case the proportion of the sample dying for any given level of immune response will be in proportion to the probability of them dying and the probability of them having total immunity, multiplied by the frequency of that immune level in the population sample.

3.2.1 Small example to help clarify

The way in which we can calculate the proportion of individuals in the population sample dying from the disease considered by the model is best illustrated by a small example.

Figure 2 gives a frequency distribution of immune responses for a large sample c and the probability of death function a , and the probability of immunity function b . If we consider just a small slice of the total scale of immune response, say that lying on the vertical line A , we notice that the intersection with A of c is at 0.304. Given that the total integrated area under c is 37.32 this means that $(0.304/37.32) \times 100 = 0.814$ percent of the population. We wish to know what proportion of this 0.814% will die, we know from Figure 2 that associated with an immune response at point A there is a probability of immunity per unit time of 0.715 and a probability of death per unit time of 0.063. The proportion of those dying will be as a proportion to these two probabilities - that is $0.063 \times (0.063 + 0.715) = 0.049$, or about 5% of the proportion represented by vertical A .

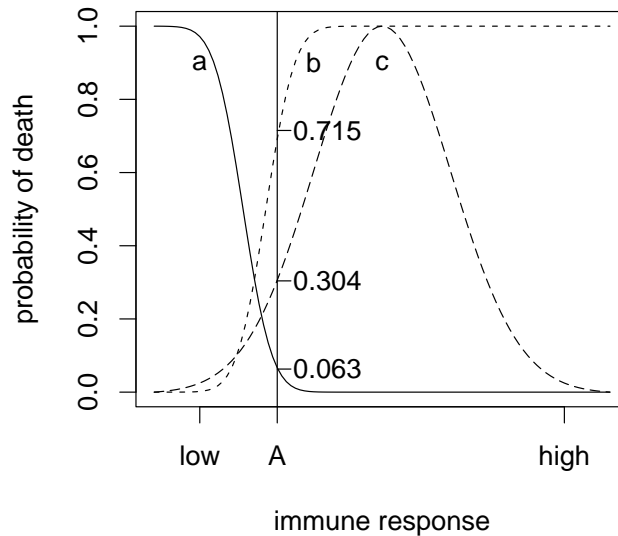
3.2.2 Whole calculation

The calculation demonstrated in Section 2 should be applied across the range of immune response to get a frequency distribution for the *dead* section of the population. This has been done in Figure 3.

⁸Don't know how true this is.

⁹It could be the case that people who had the disease, and had some skeletal manifestation actually died of something else before they completely remodeled all trace of the first disease, in which case they would be picked up on. I think at the moment we should assume those that don't die from the disease recover completely.

Figure 2: Example calculation of death population: c (dashed line) is a frequency distribution for the immune response of the sample, a is the function for the probability of death per unit time, b is the probability of immunity for per unit time for the sample. Vertical line A is the point along the immune response scale we will consider.



4 Linking to a distribution of symptoms

Here is to some extent the clever bit, but possibly the most contentious part of this model.

We have worked out the proportion of those dying from any given disease, we now wish to work out what sort of distribution of symptoms might be expected from the disease. To do this it is necessary to borrow from biological theory and say that bone remodelling activity under the influence of some disease or condition continues at a constant rate. Bone remodelling activity cannot continue after the death of the organism, so the amount of bone remodelled is going to be proportional to the amount of time lived with active remodelling. As stated in Section 3.2 all those who do not die from the condition eventually recover and display no symptoms when they eventually turn up in the skeletal record. Therefore the focus shall again be on the *dead* section of the population.

Referring back to Figure 2 let us take the section of the *dead* population on the vertical line A . We know that eventually $0.063 \times (0.063 + 0.715)^{10}$ of 0.304 will die from disease or condition c . However, associated with this $\approx 5\%$ we have a probability *per unit time* of death, this means we can work out an expected mean life for those individuals who die with this particular immune response value.


4.1 Inverse exponential functions

Inverse exponential functions are one of the families of functions all archaeologists and biologists should be familiar with. They crop up in ^{14}C dating,¹¹ as the Gompertz function and in so many other applications. I'm surprised they haven't managed to be mentioned earlier

¹⁰0.063 is $P(d|x, c)$ and 0.715 is $P(i|x, c)$ - see sections 3.1 and 3.2.1.

¹¹In fact in all radiometric dating.

Figure 3: Example calculation of death population continued: dashed lines a, b, c are as described in Figure 2. The solid line is the distribution of immune response for the proportion of the population which died, and has been calculated by following the procedure set out in Section 3.2.1 but for all values of the immune response axis. The scale of the distribution d has been magnified by a factor of seven for display purposes, and is a frequency distribution. The integrated area under d is 3.2% of the area of c , meaning that 3.2% of the population would be expected to die given the parameters of a and b .



in this report.


In way of explanation they are inextricably linked with the attritional population. Let us say we have a population, we can't add to the population, however individual members of that population happen to die. Every member of the population has an equal probability of dying per unit time. What we want to know is *how many members of the population are there still left alive as we go through time?*

If we set the probability of dying as above to 0.063, and we start with 1000 members in the population then at $t = 0$ we have 1000 members still left. After one time unit, as each member of the population has a 0.063 probability of dying, and living and dying are exhaustive and exclusive¹² then each member of the population has a probability of $1 - 0.063 = 0.937$ of still being alive. This means that after one time unit there is $0.937 \times$ the initial population size, which in this case is $0.937 \times 1000 = 937$. After two time units I expect to have 0.937 of the population surviving again, obviously not $0.937 \times$ the original population, but $0.937 \times$ the population after the first time unit: that is $0.937 \times 937 = 878$. After three time units there will be $0.937 \times 878 = 822$, after four $0.937 \times 822 = 771$ after five $0.937 \times 771 = 722$ and so on until eventually we have very few of the population members still left. This process is illustrated in Figure 4 using 100 time units for two inverse exponentials viewed as smooth functions.

If we now examine what we have done to get to the population level after five time units. First we had 1000 individuals, then we had 1000×0.937 after one time unit, then we had $1000 \times 0.937 \times 0.937$ after two time units. Then after three time units we had $1000 \times 0.937 \times$

¹²You are either alive or dead - you can't be neither or both.

Figure 4: Two inverse exponential functions: the function B has a probability of death per unit time of 0.03, A is 0.063 as in the example followed from Figure 2. The horizontal (dashed) line is when the population has fallen to 377 individuals, the population level corresponding to the mean life for both functions. The vertical lines a and b are the times corresponding to the mean lives, 16 time units for a , and 33 time units for b .



0.937×0.937 individuals. After four we have got $1000 \times 0.937 \times 0.937 \times 0.937 \times 0.937$. Essentially we for each successive time unit we merely multiply what we had before by 0.937. So after five we have $1000 \times 0.937 \times 0.937 \times 0.937 \times 0.937 \times 0.937$ members left in the population. If all the 0.937's were grouped together we would have $0.937 \times 0.937 \times 0.937 \times 0.937 \times 0.937$ which can be written as 0.937^5 . We can see that the 5 corresponds to the time elapsed, so for any time the population level will be 1000×0.937^t where t is time elapsed. For some peculiar reason¹³ the $N_0 \times (1 - \lambda)^t \approx N_0 e^{-\lambda t}$ form of the equation¹⁴ (where N_0 is the population size we started with in this case 1000, λ the probability of death per unit time, and t elapsed time). This is where the similarities between this part of the model and radiocarbon dating models are more obvious.

The significant part here is λ which is the probability of death 1 per unit time ($P(d|\mathbf{x}, \mathbf{c})$). Were the *dead* section of the population to be examined and a list made up of how long all the individuals survived the lengths of time in the list could be summed and then divided by the number of individuals, this would be the mean time before death¹⁵. Actually we don't have to sit there and members of the population dying as for this sort of function the mean life is equal to $1/\lambda$. The mean lives have been marked on Figure 4.

4.2 Mean lives from the example

As argued in Section 4 the severity of manifestation is proportional to the amount of time the individual lived with the condition or disease. So we can say that the mean expected life is also proportional to the severity of the condition seen on the skeleton. So far the

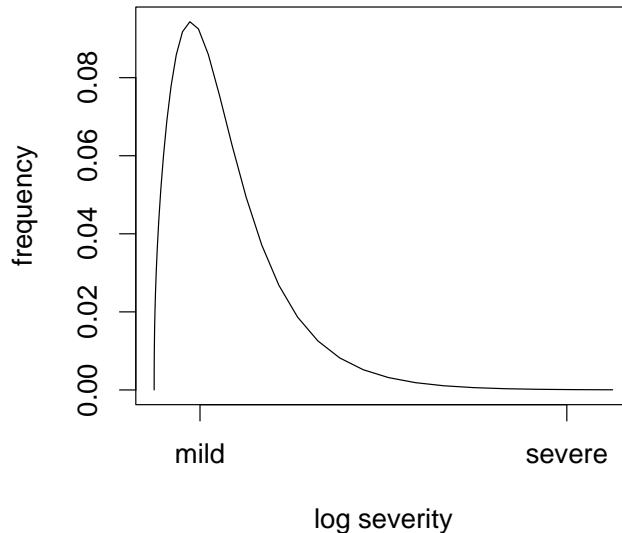
¹³I cannot remember the derivation at the moment and all my books are packed away.

¹⁴The only significant differences are in the first half life.

¹⁵Very similar to *mean time before failure* ideas for electrical goods and automotive components.

graphs have had two axis, immune response and frequency or probability. We need a third axis which is just a construction from $P(d|\mathbf{x}, \mathbf{c})$ each point on curve d in Figure 3 having an associated mean expected life.

Figure 5: Distribution of severity: this is only the first part of the distribution, as can be imagined for the higher immune response part of the population $P(d|\mathbf{x}, \mathbf{c})$ (λ) approaches zero so $1/\lambda$ will approach infinity making the x axis very long and the behavior of the majority *dead* section of the population difficult to see. Here the severity (time) axis has been logarithmically transformed for additional clarity.



There are a couple of relevant points which should be mentioned about the shape of the distribution seen in Figure 5 and particularly it's highly extended x axis (we only see the first forty time units displayed in Figure 5):

1. It can be expected that some form of truncation applies to the x axis as a skeletal element subject to any condition which leads to abnormal bone formation can only get so bad. For example: were the condition one which left the bone surface with some sort of lesions then the process can only continue until the entire surface is covered. Any further formation of lesions will be lesions on top of lesions, and will therefore just be a lesion.
2. Another thing which may *squash up* the x axis is that those individuals from the population who live longest, and are consequently worst affected, are also those whose immune responses are highest. It may be suggested that contrary to the assertion in Section 4 that the rate of change once an individual has the condition is constant, that the rate of change is somehow inversely related to the immune response scale. The *log* of severity I have used in this report may not be so far wide of the mark, however I shall not pursue these points any further than this and leave the really hard work to others.

5 Modeling different disease types


The first question is how do the changing the disease parameters affect what we see in the modeled skeletal record? The sort of thing I should like to do in this section is just model

two types of disease and see whether the model is capable of generating the sorts of numbers we know from clinical studies and the archaeological record. In Sections 5.1 and 5.2 I have tried to simulate the effects of two diseases on a single population. In particular we are interested in the distribution of severity in each case, so the general methodology has been to posit a *baseline* disease, the probability functions for death and immunity being the same, and then varied the parameters for the other disease.

5.1 Chronic low level infection

From what I gather most disease examined by palaeopathologists is of the long standing chronic variety. To look at these types of diseases I have tried to generate two suitable sets of probit functions to see whether the model can give two separate chronic diseases which kill a small percentage of the population, but give two different sets of severity.

Figure 6: Model of a generic chronic disease: Top graph - the dashed lines are the probability of death, and probability of immunity functions for a mild chronic disease. The solid lines are the same functions for a more severe chronic disease, note that neither function is scaled $[0, 1]$ as in Point 5 in Section 3.1. The dotted line is the distribution of immune response upon which both diseases are operating. The two dot-dash lines are the frequency distributions for the deaths from each disease. In this instance the proportion of individuals dying from disease one is 4.71% and from disease two is 4.51%. There are 0.96 as many deaths from disease two as disease one. The modal severity for individuals dying from disease one is 28.22, from disease two is 3.24. Probability of death function - disease one: 50% point 5 spread 15 - disease two: 50% point 20 spread 5. Probability of immunity function - disease one: 50% point 50 spread 15 - disease two: 50% point 30 spread 5.



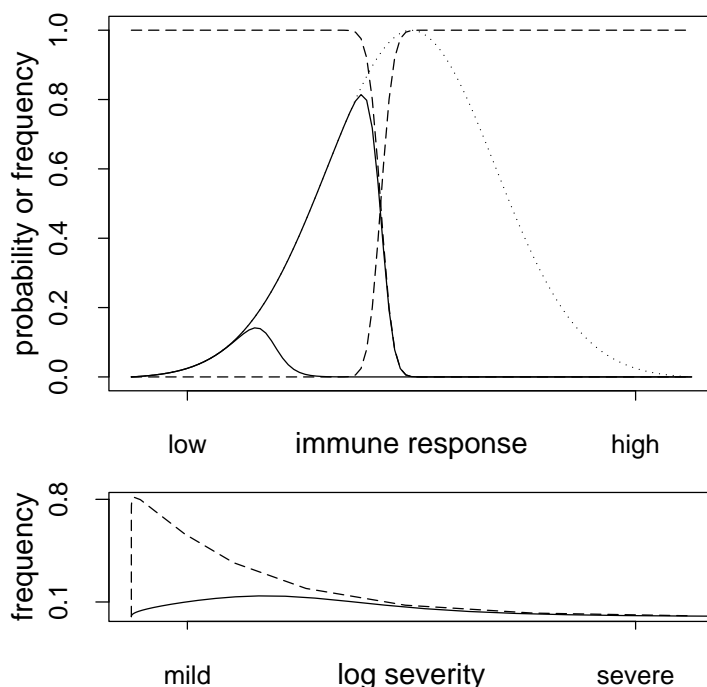
From Figure 6¹⁶ we can see that indeed two low level chronic disease functions can be generated which each kill about 4.5% of the population, the severity of the more severe disease being very much greater than the *baseline* disease. We do this by using a very low probability of death function which is not scaled $[0, 1]$ as suggested in Point 5 in Section 3.1.

¹⁶This is where a colour printer would be most useful.

5.2 Acute high level disease

A disease group palaeopathology cannot address directly¹⁷ is that which occurs quickly, causes death very quickly, and may have little effect on the skeleton. The model is illustrated in Figure 7.

Figure 7: Model of a generic acute disease: Here in the top graph the probability of death, and the probability of immunity functions for the baseline chronic disease have been omitted for clarity. The dashed lines represent the probability of death and immunity functions for the illustrated acute disease. The immune response distribution is the faint dotted line. The death functions for both diseases are shown, the larger being the one for the acute disease. The bottom graph shows the severity of symptoms, the solid line is the baseline chronic disease function, the dashed the acute disease function. In this instance the proportion of individuals dying from disease one is 4.51% and from disease two is 35.69%. There are 7.91 as many deaths from disease two as disease one. The modal severity for individuals dying from disease one is 3.24, from disease two is 1.03. Probability of death function - disease one: 50% point 20 spread 5 - disease two: 50% point 45 spread 1.8. Probability of immunity function - disease one: 50% point 30 spread 5 - disease two: 50% point 45 spread 1.8.



The disease modeled in Figure 7 would represent something like bubonic plague or smallpox. Individuals dying very quickly, approximately one third of the population succumbing. The probability functions need to be adjusted so that they are almost step functions, and that the sum of the two functions is nearly one across the whole range of x . The proportion of the population having been said to die is simply a function of where the step appears relative to the population immune response distribution. Basically all this does is kill off every individual with a immune response below a certain cutoff level. The bottom graph shows the severity of the disease for the subset of the population which died from the acute disease, and shows this in comparison to the baseline chronic disease. Unfortunately, although the behavior predicts that the majority of individuals dying from the acute disease will live for such a short time they will manifest no symptoms, the number of individuals who survive for

¹⁷Although burial studies can reveal some details - for example *plague pits*.

a long time is still greater than that from the baseline disease, thus were this model correct we should still expect to see the acute disease giving more symptoms than the *baseline* chronic disease.

5.3 Different diseases: summary

In both the examples above (Sections 5.2 and 5.1) we have compared the effects of several sets of parameters for the two probit functions on the distribution of severity for a single population. In the case of chronic disease we have been able to simulate the sort of distributions one imagines from different low level diseases. For the more acute diseases the modeling has produced results which are believable, except for the problem that the more acute case should, according to the model, leave more on the skeleton than the *baseline* chronic disease. It is early days with the model, and possibly some parameters for the probit functions associated with the more acute diseases will be found which don't have this undesirable property.

6 Different populations

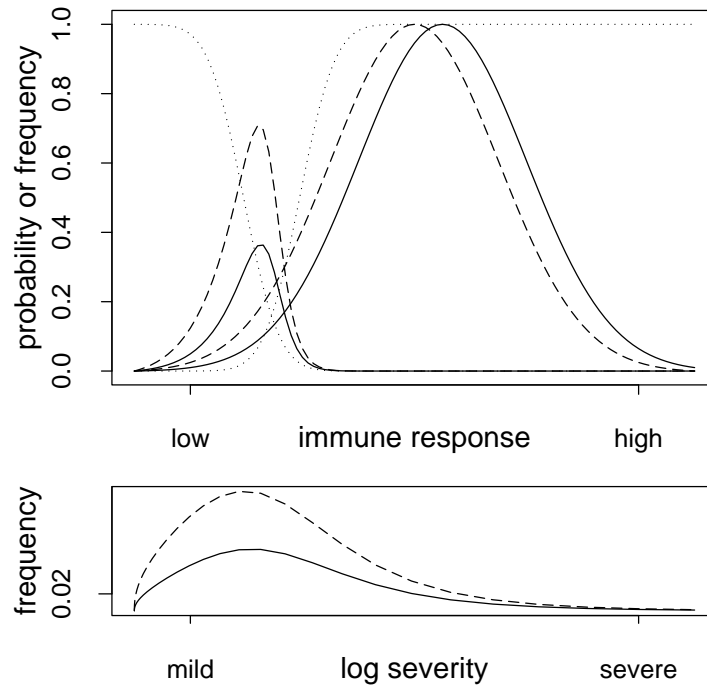
In this section we shall look at the effects of a given disease upon different populations of immune response. Different populations in this instance could be sex differences, or could relate to different geographic populations which may also have different response functions to each other. The two disease types will again be chronic and acute, the probability function parameters being taken from the *baseline* disease in Section 5.1 and the acute disease in Section 5.2.

6.1 Chronic disease

The disease represented in Figure 8 has the same parameters for probability of death and probability of immunity as the *baseline* chronic disease, this time applied to two different populations of immune response. On the top graph the population immune response distributions are given by the large Gaussian curves, the disease parameters are the probit functions. The smaller *death* distributions are scaled up by a factor of five for clarity. The bottom graph gives the distributions for severity of the disease on the two populations. It is interesting to note that the *male* population, despite having a greater proportion of it's members dying in fact shows that most of the members of the *dead* portion of this population will have less severe symptoms than the *female* population. Although there will be more *males* than *females* at any point along the severity axis. This could be simply because the model is not yet accounting for the immune process in a sufficiently informed manner, or could be real. On the whole the model is giving sensible results from what I imagine to be the real situation, and the separation in *male* and *female* immune response functions is fairly small to get credible results - I certainly would imagine that the immune response distributions would not be too different¹⁸.

¹⁸Although I may be wrong here, we only need to look at geographic populations for this in circumstances such as the introduction of smallpox to the South Americas where the Europeans had a small immunity to it, but the natives had none whatsoever.

Figure 8: Model of a generic chronic disease on two populations. All disease functions are taken from the *baseline* disease in Section 5 and are represented by the dotted line. The population parameters are the same except the two populations have means which are $1/3^{rd}Z$ (both population variances are the same). The solid line in both graphs could represent a female response function with a higher mean than the male (dashed). Here 4.51% of males die from the disease, and 2.22% of females giving a male to female ratio of approximately 2 to 1. The mode of the severity for males is about 3.2 whereas for females it is 4.1



6.2 Acute disease

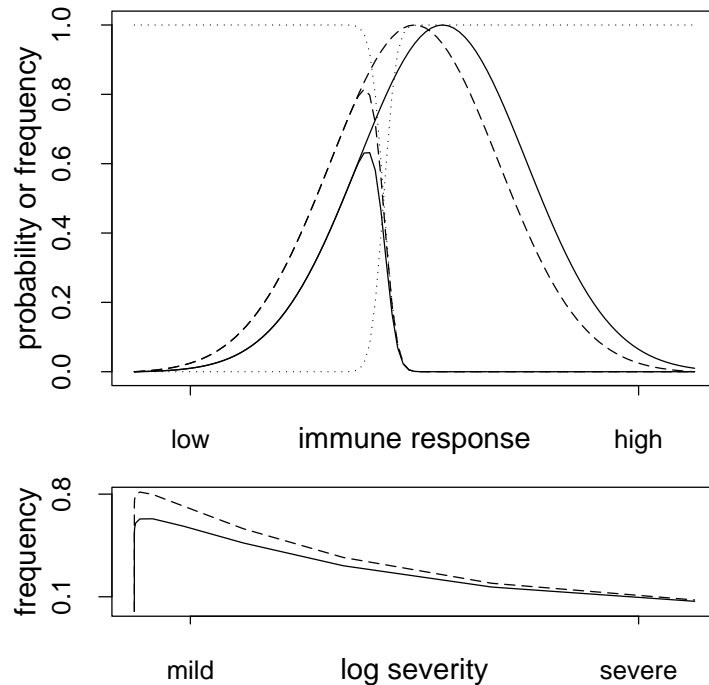
The model featured in Figure 9 has its disease parameters taken from the acute disease featured in 5.2 the probit functions are represented by the dotted lines. As in Figure 8 the two immune response populations are the Gaussian distributions, but in this case the *death* distributions are quite strangely shaped. The *death* distributions are displayed without being scaled up. As might be expected the ratio of *males* to *females* is lower this time, although one may anticipate that for a high level acute disease both populations might be affected more or less equally. More in line with this expectation are the distribution of severities which are virtually identical.

6.3 Different populations: summary

The model works reasonably well on the different immune response populations problem. Again as with the *two diseases* problem the main weakness seems to be when we're trying to model acute high level diseases¹⁹. The differences between the modeled chronic and acute diseases in the way in which the *male* and *female* ratios change is a prediction of the model which seems to make some sense, as is the near identical severity distributions. A surprising feature is in the chronic case in which there are twice as many *males* dying, but the distribution of symptoms tends towards a lower severity.

¹⁹By *high level* I mean causing a large portion of the population to die from them.

Figure 9: Model of a generic acute disease on two populations. All disease functions are taken from the acute disease in Section 5.2 and are represented by the dotted line. The population parameters are the same except the two populations have means which are $1/3^{rd}Z$ (both population variances are the same). The solid line in both graphs could represent a female response function with a higher mean than the male (dashed). Here 35.6% of males die from the disease, and 24.3% of females giving a male to female ratio of approximately 1.4 to 1. The mode of the severity for males is about 1.02 whereas for females it is 1.08 which are virtually equal



7 Final points

In conclusion there are a number of points which need to be addressed about the success and failures so far of the model, and as guides to the next steps.

1. I think the least successful part of the model is where it is applied to simulating acute diseases and deriving the distribution of symptoms. The problem is that the model predicts that there should be more symptoms for high level acute conditions than for chronic conditions. We know this to be wrong. Ways around this might be:
 - (a) employ stepped functions for the probits for the probability functions
 - (b) some explanation in the form of the progress of the condition.

Neither of which is truly satisfactory as they introduce some elements of *special pleading* for a model which has pretensions to universality.

2. More successful has been the application in prediction of chronic disease. I have left the assumption that the progress of any disease is linear. This is fairly obviously not the case whether it is because of some reason such as Point 1 in Section 4.2, or Point 2 in the same section. But I think this will be a crucial area of progress for the model.
3. It may be the case that the immune response distributions may be disease specific. This may be expected, but would not be elegant in terms of mathematical modeling.

4. The major problem will be the fitting of parameters to the various functions and distributions which comprise the model. I haven't thought this through thoroughly, but parameters may be available from the epidemiological literature. Most of them can be fairly approximate in form, the problem would be knowing where the form lies in respect of the immune response distribution. I have a notion that some of their characteristics may be deducible from the vectors of transmission.
5. I have omitted low level acute diseases - it may be a good idea to examine the expected behaviour of these as well.

On the positive side:

1. The model gives credible results in terms of the distribution of severity, and the proportion of individuals who die from a particular condition.
2. In the *two population* problem the population immune response distributions don't have to be very far apart, which is what one might expect.
3. Because of the time element introduced for the severity distribution the model would lend itself to some more dynamic modeling. We have only considered a stationary cohort here, and although it would give asymptotic predictions for stationary populations, the iterative dynamic approach would be useful for the highly non-stationary populations suspected in ancient times.
4. The model is capable of providing predictions which are unexpected consequences of the various parameters. These may be real, or if false, may be used to develop the model.
5. The model can also be used to generate²⁰ a distribution for the part of the population which caught the disease, but didn't die from it.
6. I have generated the results in this report empirically, that is by crunching vectors of real numbers which represent smooth functions. Whilst this approach is great for rapid model development and seeing approximately what is going on, it is also inelegant. I suspect that the distribution functions which are generated, that is the *deaths* distributions and the severity distributions are simple functions from known families. For example the severity distributions look distinctly like γ functions. If this is the case then if this model is to really be credible, these functions, and how to calculate their parameters from the model parameters must be understood.

More generally:

When discussing this particular model with my colleagues from the world of mathematics we were unable to believe that this sort of thing hadn't been done before. We came to the conclusion that a hunt through the epidemiological literature (or chat to an epidemiologist) would reveal a wholly suitable model with all the problems which this model is experiencing already ironed out.

Six months later up here in Edinburgh, and having asked as many people as to whether anything like this has been done before it would appear that it is a bit of a first, and could do with a Ph.D. or something like it to remove most of the major difficulties.

²⁰I have not done the calculations here for the non-dying population segment.