# Compound (or exponential) growth and decline: a conceptually minimal approach

Michael Wood (MichaelWoodSLG@gmail.com)
August 2015

*Mathematics is there to make difficult things easy.* Eugenia Cheng (Cakes, custard and category theory)

Introduction ........ 1
A conceptually minimal approach to compound or exponential growth ........ 2
Summary ........ 6
Some examples ........ 6
Conclusions and general lessons ........ 7

## Introduction

Excel spreadsheet for calculating compound growth: http://woodm.myweb.port.ac.uk/SL/compoundgrowth.xlsx

Simple Learning home page: http://woodm.myweb.port.ac.uk/SL/simplelearning.htm

The notation used for mathematical formulae is discussed here: http://woodm.myweb.port.ac.uk/SL/mathformulae.pdf.

Many things grow or decline at a constant, or roughly constant, rate. The obvious examples are investments, debts on which interest must be paid, and biological or human populations. Slightly less obviously, the example which triggered my interest in this topic is the analysis of the reliability of manufactured components: if you start off with 1000 light bulbs, how will the number that are still working decline through time, assuming that this rate of decline is roughly constant?

For example, suppose you borrow £100 at a rate of interest of 10% a week. After a week you will owe £110, which isn't too bad. But if you don't manage to pay it back for two years (104 weeks), the interest on the original £100 (known as *simple* interest) will be £1040 and you will owe a total of £2040. However, in practice, you will have to pay interest on the interest as well - this is known as *compound* interest - which increases the total amount owed far more than you might expect. With compound interest the amount owed after two years will be £2 017 619 - over two million pounds, which is almost two thousand times what you would pay with simple interest! This is known as *exponential growth*, which intuition often under-estimates. (See Example 1 below.)

Normally compound or exponential growth is analysed using some fairly advanced mathematics which makes the concepts inaccessible to many. My aim here is to create a new set of concepts which I hope will be accessible to someone without the mathematical background to understand the conventional approach. The next section covers the new approach in detail, and the section after it gives a summary which, I hope will make some sense to those who have skipped the next section.

## A conceptually minimal approach to compound or exponential growth

The standard mathematical way of analysing these problems involves the use of the exponential function , $e^x$, and natural logarithms. Understanding the rationale behind this requires a background in calculus and the theory of logarithms. (The Wikipedia pages on the Exponential function, Logarithms and Natural logarithms will give you an idea of this background, although this goes beyond what is necessary.) My problem was teaching this aspect of reliability theory to students who did not have this background. What I did is to follow the approach adopted by almost all teachers of mathematical techniques to non-mathematicians: you teach them how to use the formula without mentioning where the formula comes from or why it works. If you don't really understand the next bit, don't worry: this is really the point!

In my reliability classes I tried to teach the students that if *r*% of bulbs fail in any hour, then the proportion of light bulbs remaining after *t* hours will be

$$e^{-r\,t}$$

So, for example, if *r*% = 1% and *t* = 100 hours, then the proportion of bulbs still working is

$$e^{-0..01x100} = e^{-1} = 36.8\%$$

Going the other way, if we know that the proportion of bulbs still working is *p* after *t* hours, then the proportion of bulbs failing per hour will be

$$-\log_e(p)/t$$

and there is a similar formula for working out *t* if we know *r* and *p*.

In practice most of my teaching efforts were devoted to helping the students use the $e^x$ and $\log_e$ buttons on their calculators. This involved a moderate amount of effort, so we all felt that some useful education had occurred.

However, I don't really think this is so because all they could do was use the buttons and get the answer: they didn't know what the functions did or what the answers represented, and they had little appreciation of the rationale behind the formulae, what assumptions are made and whether they are reasonable. And because it meant little, once they closed their books, they probably did not remember much.

So I decided to try and simplify the underlying concepts so that they would make as much sense as possible to someone with a certain amount of arithmetical common sense but no knowledge of the underlying theory of calculus and logarithms, or any desire to acquire such knowledge. What follows is the result of many false starts and more effort than I care to admit. (The formulae below all use spreadsheet (Excel) notation so they can be pasted directly into a cell on a spreadsheet. Also, please remember that a percentage is simply a way of expressing a proportion: 0.1 and 10% mean the same thing and Excel will interpret them both in the same way - if you right click a cell in Excel containing a number, and then click Format, you will get a variety of options for changing the way numbers are presented.)

I'll use interest rates on loans as the initial example to illustrate the concepts. Suppose you need to borrow £100 for a week, and are offered a loan at an interest rate of 10% for the week. You will pay back £110 at the end of the week. If nothing else is available this perhaps doesn't sound too bad. The difficulty, of course, is what happens if you don't manage to pay it back at the end of the week. Let's imagine that you have the loan for 20 weeks. Also let's suppose to begin with that you just pay simple interest - in other words you pay interest on the original £100 loan but not on the interest that you have to pay after one week, two weeks, and so on. Over 20 weeks the interest will be 200% (=20*10% using spreadsheet notation for formulae) of the loan, or £200.

However, the true situation is much worse than this if, as is almost certainly the case, the interest is "compounded" - you pay interest on the interest and this can mount up surprisingly quickly if you don't manage to pay it back in the first week. The interest on the first week is £10, and for the second week it is 10% of the total owed - £110 - which comes to £11. At the end of the second week the total owed is £121 so the interest is £12.10 and the total owed is £133.10. And so on.

Each week 10% is added to the amount owed so the amount owed at the end of the week is 1.1 (1 plus 10%) times what it was at the beginning of the week. After two weeks the amount owed will be multiplied by another 1.1 so the overall multiplier will be =1.1*1.1 =1.1^2 (^2 means "to the power of 2" or squared), and after 20 weeks it's =1.1^20. The interest as a proportion of the amount borrowed will be this minus 100% (representing the amount borrowed at the start), which is

=1.1^20-1          (remembering that 1 is the same as 100%)

which is 572.75%. This is a lot more than the 200% growth with simple interest, all due to the interest being compounded.

Many banks add on interest every day, which raises the question of what will happen if the interest is added more frequently. Obviously the more frequently the interest is added the more interest on the interest you will pay so the higher the final debt. But how much more?

Let's suppose the *n*umber of time *p*eriods is *np*. If the interest is compounded every week np=20, but if we work it out in the general case we can then experiment to see what happens if np is larger. The proportional growth in each of these time periods is 200%/np. Using the same method as we used before for the 20 week case, the proportional growth (or the overall interest rate)

pg =(1+2/np)^np - 1       (remembering that 200% is the same as 2)                    Equation 1

If we replace np by 20 this becomes

=(1+0.1)^20-1

which is the same as before, as we would expect.

The advantage of the general formula with np is that we can try different numbers for np. For example:

| *np* | *Proportional growth* |
|---:|---:|
| 1 | 2.000 (200%) |
| 20 | 5.727 (572.7%) |
| 1,000 | 6.374 (637.4%) |
| 1,000,000 | 6.389 (638.9%) |
| 1,000,000,000 | 6.389 (638.9%) |

As expected, the more time periods there are the greater the proportional growth is, but the increase slows down so that it will never be more than 6.389. In the final row there are a billion time periods so the time intervals will be infinitesimally small. This is, in effect, *continuous growth* - we get to know what will happen if interest is added on continuously based on the amount owing (or the size of the population) at each point in time. If we tried the analysis with an even bigger number (easy to do with a spreadsheet) the proportional growth would still be 638.9% to one decimal place. This is referred to by mathematicians as the limit as the size of each time period tends to zero.

(Another, more intuitive, example of a limit is 1 + ½ + ¼ + 1/8 + 1/16 .... The first two terms comes to 1½, the first three to 1¾ and so on. The sum obviously gets as close as we want to 2, but never quite reaches it. The limit, as the number of terms becomes larger and larger, is 2.)

So, to recap, if the proportional growth with no compounding (which we'll abbreviate to *pgwithnocomp* ) is 200%, the proportional growth with continuous compounding (*pgwithconticomp*) is 638.9% assuming that the 200% growth rate applies at each point in time. We could do a similar analysis for any other value of *pgwithnocomp*. If you prefer shorter names for these quantities I'd suggest *pgnc* and *pgcc*.

Excel has many built in functions: for example there is one for finding a square root. To see how this works, type =SQRT(9) into a cell of an Excel worksheet. The number "returned" in the cell will be 3, the square root of 9. The function SQRT(number) returns the square root of whatever number is.

Unsurprisingly, there is no built in Excel function to work out *pgwithconticomp*. However, there are two functions EXP (the exponential function) and LN (the natural logarithm) which we can use to work out *pgwithconticomp* and to get back to *pgwithnocomp*:

PGWITHCONTICOMP(pgwithnocomp) = EXP(pgwithnocomp)-1 Equation 2

PGWITHNOCOMP(pgwithconticomp) = LN(pgwithconticomp + 1) Equation 3

So, for example, if *pgwithnocomp*=200%, Equation 2 says that *pgwithconticomp* =EXP(200%)-1, and if you key this into a cell in Excel, the answer is 638.9%. Which is just the same answer as above.

Equations 2 and 3 can be proved using the methods of mathematics; alternatively the fact that they give the same answer as the argument above should, I hope, convince you.

The table and graph below shows the relationship between pgwithnocomp and pgwithconticomp. Negative values are included because the same formulae still work - the third example below illustrates how this works. I've also put a spreadsheet on the web (http://woodm.myweb.port.ac.uk/SL/compoundgrowth.xlsx) which allows you to convert *pgwithnocomp* to *pgwithconticomp* and also allows you to see what happens if compounding is done in a particular number of steps (Equation 1).

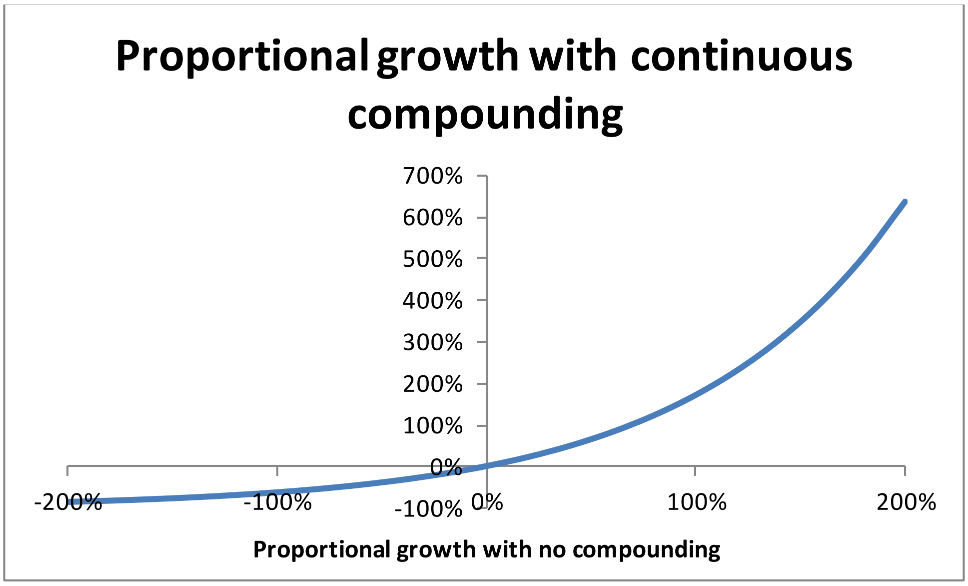


| Proportional growth with no compounding (pgwithnocomp or pgnc) | Proportional growth with continuous compounding (pgwithconticomp or pgcc) |
|---:|---:|
| 8.00 | 2,979.96 |
| 7.50 | 1,807.04 |
| 7.00 | 1,095.63 |
| 6.50 | 664.14 |
| 6.00 | 402.43 |
| 5.50 | 243.69 |
| 5.00 | 147.41 |
| 4.50 | 89.02 |
| 4.00 | 53.60 |
| 3.50 | 32.12 |
| 3.00 | 19.09 |
| 2.50 | 11.18 |
| 2.00 | 6.39 |
| 1.50 | 3.48 |
| 100% | 172% |
| 50% | 65% |
| 0% | 0% |
| -50% | -39% |
| -100% | -63% |
| -150% | -78% |
| -200% | -86% |

As you will see, the difference between the two columns gets far greater for larger values of pgwithno comp - if this simple interest is 8 times (800%), the corresponding pgwithconticomp (continuously compounded interest) is almost three thousand times. This is so-called exponential growth.

## Summary

I'll illustrate the concepts with the example of a 10% per week interest rate over a period of 104 weeks as outlined in the introductory section.

The two core concepts are:

1. Proportional growth with no compounding (pgwithnocomp or pgnc): 1040% in our example.
2. Proportional growth with continuous compounding (pgwithconticomp or pgcc): 3285862.57% (which is over three million percent).

In the introduction above I explained what would happen if the interest is compounded every week. If the interest is compounded every day, or every hour, the total amount of interest would obviously be greater. Pgwithconticomp is the proportional growth if the interest is compounded in infinitesimally small time intervals. This describes the exponential growth of things like populations.

There is a simple relationship between these two concepts. For any value of pgwithnocomp we can work out the corresponding pgwithconticomp. This relationship is described by Equations 2 and 3 above, and shown in the graph and table above, and by the spreadsheet at http://woodm.myweb.port.ac.uk/SL/compoundgrowth.xlsx.

The examples below should make this clearer.

## Some examples

*All the answers below can be worked out from Equations 1, 2 or 3. These are incorporated in the Excel spreadsheet at http://woodm.myweb.port.ac.uk/SL/compoundgrowth.xlsx which can be used to work out the answers.*

1. If you don't manage to pay the 10% per week loan back for two years, or 104 weeks, pgnc (the simple interest) is 104*10% or 1040%. If you put this in the pgnc green cell in the spreadsheet, and 104 for the number of steps in the yellow cell, the pgsc figure in the blue cell is 2 017 519.45%, corresponding to the compound interest. The pgcc figure - if the interest were compounded continuously is more than three million percent. The fact that both of these figures are far larger than you might expect has got many people into trouble.

2. Financial products like loans and mortgages are complicated by the charges made, and by the frequency with which interest is added. There is a short explanation at http://www.theguardian.com/money/2007/oct/25/debt.savings which gives this example:

> " For example, an account offering a rate of 6.25% paid annually may look more attractive than an account paying 6.12% with monthly interest payments, however the AER on the

monthly account is 6.29%, as opposed to an AER of 6.25% on the account with annual interest payments."

AER is the annual effective rate which, without charges, is the proportional growth over a year. Using the spreadsheet http://woodm.myweb.port.ac.uk/SL/compoundgrowth.xlsx you should be able to work out the 6.29% AER. 6.12% is the interest that would be paid if there was no compounding so put this in the pgnc green cell, and put the number of steps - 12 - in the yellow cell. The answer should then be in one of the blue cells.

3. Let's try the formula with a reliability problem . Suppose that 1% of light bulbs fail every hour. What proportion will fail over 200 hours? Pgwithnocomp, the equivalent of simple interest, is *minus* 200% (=200*(-1%))because the number still working will be declining, not growing. If you put this in the pgnc green cell, the value of pgcc will be -86.47%. In other words the number working will have declined by 86%, so the number still working will be 14%.

Notice that pgwithnocomp is minus 200% which seems to suggest that the number of bulbs left working is negative - which does not make too much sense. This is a hypothetical figure we need to get the answer: it has no easy interpretation.

4. Finally let's try a population problem. According to Wikipedia the 2013 list by the CIA World Factbook gives the UK's growth rate as 0.54%. Imagine it stays the same for 10 years, Then

pgwithnocomp = 10*0.54% = 5.4% and

PGWITHCONTICOMP(5.4%) = 5.55% (or use the spreadsheet)

which is much the same. Now let's try 1000 years:

pgwithnocomp = 1000*0.54% = 540%

PGWITHCONTICOMP = 22,041%

indicating that the population will have multiplied by a factor of 220 - making it a little over 14 billion assuming a current population of 65 million. This is about double the present (2015) population of the entire world! With small increases, or short time periods, compounding makes little difference. With large increases, or long time periods, compounding makes a massive difference, far bigger than most people's intuitions suggest.

## Conclusions and general lessons

All I've done is rewritten some very standard mathematics using slightly different concepts and notation. Why bother? What's gained? My assertion is that the new version is simpler, but this obviously needs a bit of explanation. Let's look at the senses in which the proposed framework above is simpler and more useful than the conventional one involving exponential functions and natural logarithms.

1. Firstly, the conceptual background necessary for a thorough understanding is much reduced. The argument above presupposes an understanding of basic arithmetic including proportions and percentages. Let's call this *arithmetical common sense* (ACS). In addition the argument above

introduces the reader to ideas of infinitesimals and continuous growth, and the functional notation used in Excel. A similarly thorough understanding of the conventional framework presupposes all this (although conventional mathematical notation might be used instead of Excel functional notation), and also aspects of the theory of logarithms and calculus - which is actually an enormous area of study. This means that the conceptual background required for the conventional approach is far more extensive.

2. My use of spreadsheet (Excel) notation for formulae - which means they can be pasted straight into a spreadsheet. This tactic is discussed at http://woodm.myweb.port.ac.uk/SL/mathformulae.pdf).

3. I chose the names of the key quantities - pgwithnocomp and pgwithconticomp - with some care. These need to be short enough to fit into a formula, and to remind users what they mean. The term "exponential" has passed into everyday language (although in a rather vague way), but logarithm gives no sense of being the inverse function for getting back to where we started from. I have a suspicion that, unlike terms in ordinary natural language, technical terms are notoriously resistant to change. People tend to stick to the names used by, or derived from, the original inventors (e.g. the "central dogma of molecular biology" was called a "dogma" by its originator, Francis Crick, and the name seems to have stuck despite Crick's admission, according to Wikipedia, that "dogma" is an inappropriate term). My suggestion for the names of the two concepts are just suggestions and can doubtless be improved. But I do think names are important.

4. One of the problems with the function $e^x$ is that it gives you the proportional change in the *total*, whereas $x$ is the proportional *growth*. So, if for example, $x = 0$ (no growth), $e^x = 1$ indicating that the total (debt, population, etc) is unchanged. This lack of symmetry is a confusing factor I wanted to avoid, so I decided that both of my concepts would refer to growth.

5. The reliability formula, $e^{-r\,t}$, has two input quantities and a negative sign, whereas PGWITHCONTICOMP(pgwithnocomp) just has one input quantity and no negative sign. In this sense, my version is simpler. The snag, of course, is that users need to work out themselves that (using the example above) 1% is a rate of decline so must be negative in the formula, and that for 200 hours the pgnocomp will be 200 times this. However, making sensible use of PGCONTICOMP requires ACS (arithmetical common sense), and for anyone with ACS, this calculation would be trivial.

6. My two equations (2 and 3 above) are in this sense minimalist: they just tell you the bit that's difficult to work out for yourself. Encouraging users to use ACS for the easy bits seems likely to ensure that the ideas are better understood.

7. This brings me to my final, very practical, advantage. We have just one formula covering finance, populations, reliability, and many other things. This is obviously simpler and more economical than having a different formula for each application.

8. It's also intriguing to mull over whether the intuitive interpretation of the new functions would bring any value to more complicated formulae starring the exponential function like the standard normal distribution function. (We can always replace $e^x$ by 1+PGWITHCONTICOMP(x) to obtain an equivalent formula.) I'm not sure.

Mathematicians would probably also object that these proposals mean depriving people of the elegance of the conventional mathematical functions. This ignores the fact that such elegance tends to only be perceived by experts, but it does highlight an important issue. Aesthetic appeal is important. I find the symmetry between the two concepts pgwithnocomp and pgwithconticomp, and the neatness of the link between them as illustrated by the graph above, very appealing, and more elegant than the formulation in terms of the conventional functions. I am less keen on the names of the functions, which somehow lack gravitas.

Finally, you may be wondering if I tried these new functions in my reliability class, and how well they worked. I didn't, of course. My excuse is that I didn't work out the ideas here until after I had finished teaching the class, but there is a bit more to it. Students expect to be taught the standard fare, and they need the standard fare to understand the subject as it's presented in textbooks, software, regulations and so on. This inertia is a very powerful force making change difficult. And, my suggestions are untested: there are doubtless many improvements that could be made. We need some experimental epistemologists to carry out some controlled trials.