The Ocean Physics behind Net Zero
Professor Myles Allen
31st January 2023

Have you ever wondered why the deep ocean is cold? Most people think the answer is obvious: the sun doesn't penetrate down there, so there is nothing to heat it up. But that could only explain why the deep ocean isn't warming up. Not why it is cold in the first place.

And it is cold, really cold. Arctic, to be precise. This is a vertical cross section through the Atlantic Ocean from south to north – it's expanded by a factor of 1000 in the vertical so you can see what is going on. The sea-mounts aren't nearly this spiky.

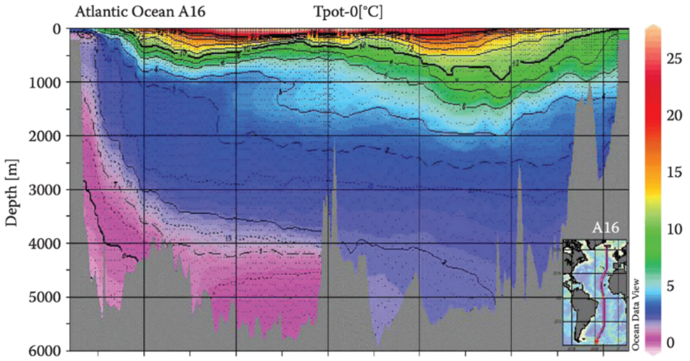

When we see these kinds of sections through the Earth's oceans and atmosphere, it is easy to forget the planet is covered by just a thin skin of fluid. Which makes it even more remarkable that we see temperature contrasts of over 20 degrees over a few hundred metres in the vertical, and essentially uniform temperatures over thousands of kilometres in the horizontal.

It is even more striking when we look at a south-to-north section through the Pacific.

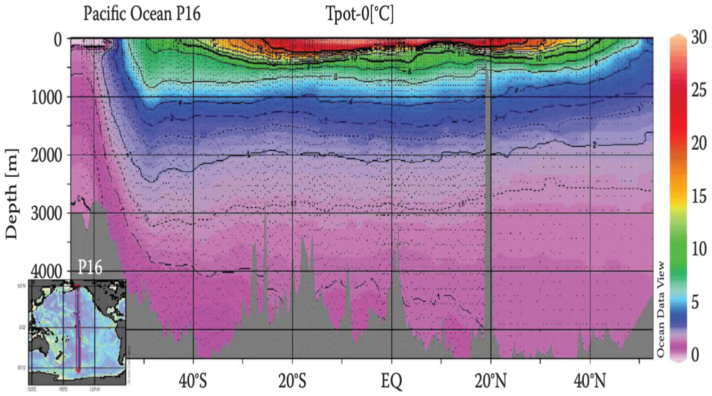

In the northern tropics we see a temperature contrast of 20 degrees over less than one hundred metres of the water column. A scuba diver could start at the surface in near-bath-water temperatures and get hypothermia within amateur diving depths.

This situation is perfectly stable in the short term, because warm water is less dense than cold water, inhibiting any mixing in the vertical. We call this “stable stratification.” But that explains why this temperature gradient can persist for a while: it doesn’t really explain why it is there in the first place.

Of course it’s cold, you might think: sunlight can’t penetrate that deep to warm it up. But, again, that only explains why it isn’t warming up, not why it is so cold to start with.

Water conducts heat, as anyone who has been soaked to the skin knows. But even though water is one of the best liquid heat conductors, conduction is a relatively slow way of transporting energy. If you heat up one end of a 10cm steel nail, it would still take a couple of minutes for the temperature to even out all along the nail (although if the point was really hot, you might have to drop it before then). If you heat the top of a 10cm deep beaker of water, ignoring conduction through the walls of the mug, and made quite sure not to stir it, it would take several hours for the temperature to even out through the whole depth of the beaker (if that sounds implausible, leave out a Styrofoam cup of water on a frosty night and see how long it takes to freeze all the way down).

The length of time it takes for temperatures to even out over a distance by conduction depends on the square of the distance – that should make sense to you, because the longer the distance, the smaller the average temperature gradient driving energy from the hot end to the cold end. So if it takes a few hours for temperatures to even out by conduction through water over 10 centimeters, it could take 10,000 years for temperatures to even out over 500 metres.

Ahah, you might think, 20,000 years ago we were in the middle of an ice age. Are those cold temperatures evidence of the oceans still recovering from the last Ice Age? And doesn’t that mean sea level rise is entirely natural. Be careful, next thing you’ll be sending me an angry email proving that climate scientists are all wrong, global warming is entirely natural, to follow up that one about Angstrom and saturated carbon dioxide absorption bands. Simple and plausible back-of-the-envelope arguments are great when they actually capture the essential physics driving a phenomenon in the real world, but can go badly wrong when they don’t. Lord Kelvin used a similar argument based on temperature gradients down through the Earth’s crust to prove that the Earth was only about 6,000 years old, and hence the Bible was right and Darwin was wrong. He convinced a lot of people. He just didn’t know about natural radioactivity in rocks.

In any case, even in the last ice age, tropical ocean temperatures were only a few degrees colder than they are today, so that can’t possibly explain why we are seeing these Arctic temperatures only just below the surface in the Tropical Pacific. The oceans have been there for billions of years – more than enough time for conduction to even temperatures out. No, there has to be another explanation.

And there is – which turns out to be directly relevant to our Net Zero story.

In the last lecture, we talked about energy flows through the atmosphere, how increasing the carbon dioxide blanket around the Earth reduces the rate at which energy escapes to space, causing an imbalance between incoming energy from the sun and energy flowing back out into space. We talked about how the climate system has to warm up to restore the balance between incoming and outgoing energy.

And we talked about how many things change as the world warms, affecting the efficiency with which the Earth gets rid of energy per degree of surface warming, the apparently all-important “sensitivity parameter”, and so how hard it is to use climate models to predict the equilibrium warming due to a doubling of carbon dioxide concentrations – or, indeed, the long-term equilibrium response to any constant atmospheric composition.

But we are a very long way from climate equilibrium. The climate system will not reach equilibrium again in our lifetimes, or the lifetimes of our great-great-grandchildren. So, we need to understand how the climate system behaves when it is out of equilibrium. And that is where the oceans come in.

To come back to the puzzle of why the deep oceans are cold. Over most of the oceans, warmer, less dense, water, heated by the sun, lies over colder, denser, water (kept cold for reasons to be revealed), so there is very little mixing between them. But temperature isn’t the only factor affecting the density of sea water.

You’ve all heard about reading newspapers in the Dead Sea. Increasing salt-content, or salinity, makes sea water denser.

And salinity becomes particularly important when water is cold, because as water temperatures approach freezing, the density increases less rapidly as it cools. It even eventually reverses as freezing actually occurs, which is why ice floats. Water is very unusual this way. Most liquids sink when they freeze. But it is a very good thing, otherwise the oceans would have frozen solid billions of years ago and life on Earth would have had no chance at all.

So, in only a few a cold regions of the world, in the North Atlantic and near Antarctica, where surface temperatures are approaching freezing, it only takes some evaporation of wind-blown spray, or freezing into sea-ice, to make surface waters more saline, and they may then become denser than the water below, sinking down to the ocean depths. Crucially, this can only happen when surface waters are near freezing to start with.

And once water has sunk down to the ocean depths, then your intuition is right, there is absolutely nothing there to warm it up again (we can ignore the tiny trickle of geothermal heat). No sun. No equivalent of cloud formation and rainfall in the atmosphere. A body of water moving round the oceans “remembers” the temperature and salinity it had when it last encountered the surface, even as it drifts around for hundreds of years. Which creates what we call the Great Ocean Conveyor, or the thermohaline circulation. Water sinks down into the ocean depths in just a few of isolated “deep water formation” regions, and then circulates around for decades to centuries before eventually emerging back to the surface potentially on the other side of the world.

Oceans are coloured by surface density, so water only sinks down to the ocean depths in these regions where it is densest at the surface. And then it travels southwards, not mixing with warmer, less-dense waters above, all the way around the world. My one grumble with this amazing animation is that it implies water pops up again somewhere else. In fact, there are no “shallow water formation” regions to match those “deep water formation” regions: the process is more like an aquarium pump in reverse, drawing surface water off and pumping it down to the deep ocean and then letting water rise very gradually everywhere.

So, what has this got to do with Net Zero?

In the previous lecture, with the help of our Gresham Climate Model, we introduced the bathtub analogy for energy flow through the climate system.

If there is any imbalance between incoming and outgoing energy, resulting, for example, from an increase in greenhouse gas concentrations, that energy has to go somewhere. Over the most recent decade, the direct impact of human activities changing the composition of the atmosphere was to reduce outgoing energy to space, or increase the net downward flow of energy into the climate system, by 2.5 Watts per square metre. Now the Earth has already warmed up by around a degree, which is driving energy back out into space at around 1.75 Watts per square metre. But that leaves 0.75 Watts per square metre, which is an awful lot of energy when you add it up over the whole Earth's surface. Energy is conserved, so it has to go somewhere. Where has it all gone?

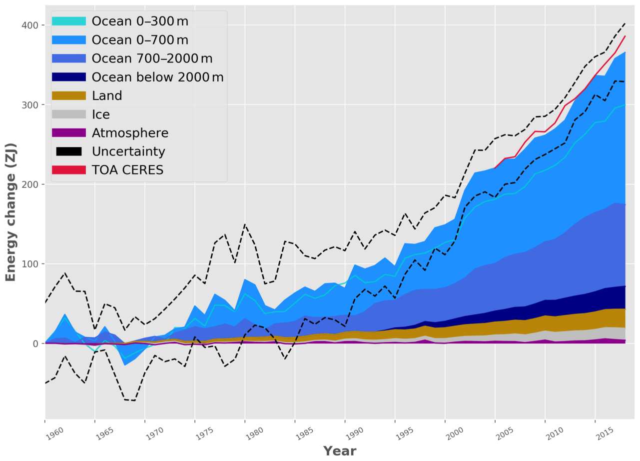

von Schuckman et al, 2020

This is a plot of energy accumulation in the climate system, since 1960, in ZettaJoules. A ZettaJoule is a lot of energy. A billion trillion Joules, to be exact. The additional energy accumulating in our climate system as a result of past greenhouse gas emissions dwarfs the energy we actually produce from burning fossil fuels. And don't forget there is more than twice as much again being radiated away because of past warming.

People sometimes ask me whether the actual heat produced in power stations or driving cars or whatever matters for climate. It can have a local impact, but globally, the overwhelming impact of energy use on climate is through their greenhouse gas emissions. This matters because, for example, a nuclear power plant produces just as much direct heat as a fossil fuel power station: but it doesn't produce greenhouse gases, at least not from operations.

The coloured bands show where this heat energy is going. The vast bulk of it accumulates in the oceans, some on land, some is taken up in melting ice, only a very small fraction in the atmosphere, because the heat capacity of air is so much smaller than the heat capacity of water. But most of it ends up in the near-surface ocean: half in the top 700 metres, even though these comprise only one-sixth of the oceans' volume. We've only been able to monitor this accumulation of heat in the oceans quite recently, thanks to initiatives like the Argo float programme, so estimates of global ocean heat content trends back in the 1980s are very uncertain. But the trend since 2000 is rock-solid, and matches the energy imbalance we can detect from satellites very well, which is this red dotted line, so we're able to "close the energy budget". And it's clear that energy is accumulating in the climate system at a rate of 0.75 Watts per square metre, or 12 ZettaJoules per year, which is over 20x world primary energy production.

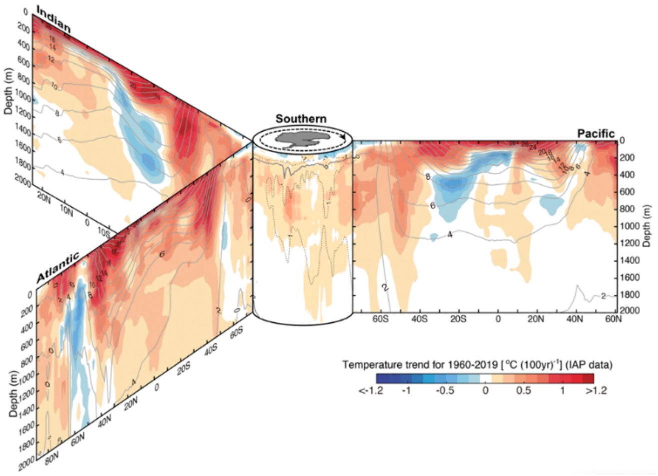

Cheng et al, 2020

Here we see how this energy is penetrating down into the deep oceans. Trends are very uneven – almost no penetration in the tropical Pacific and Indian Oceans (and remember that is more than half the world's oceans by area), more in the Atlantic and Southern Ocean, where heat from the surface is being transported down by deep water formation. I've told you deep water is cold, but what we are seeing here is that it is becoming slightly less cold (these units are hundredths of a degree per year), transporting energy down into the ocean depths.

And to understand what this means for surface warming we turn back to our Gresham Climate Model. We introduced this model in the last lecture. When I turn it on, this gentle trickle of fluid represents the natural flow of heat energy through the climate system in from the sun and back out into space. The natural flow has been scaled down so the whole model is manageable, but changes in energy flow are to scale. So when I turn up the speed of the pump, like this, that represents an extra net flow of energy in at the top of the atmosphere. Because it's a net flow in, that might arise from an increase in the power output of the sun, a decrease in the reflectivity of clouds (by China and India cleaning up their atmospheres, for example), or, of course, an increase in greenhouse gas concentrations.

The fluid level in this small tube adjusts pretty quickly to a change in flow because it is a narrow tube – representing the relatively small heat capacity of the atmosphere, land and near-surface ocean. So, on slow timescales, and ignoring the impact of the deep ocean, this model is close to equilibrium, with the increase in height of the fluid in this tube (or Earth's surface temperature) driving fluid (or energy) out faster "to space" and maintaining the balance between incoming and outgoing fluid (or energy).

We can express with an equation, which is:

$$F = k \times h$$

- F is the additional rate of fluid flowing in, above "background" flow.
- h is the increased water depth above the initial equilibrium level.
- k is the openness of the outlet pipe.

The additional flow is simply proportional to the additional height of the fluid in this tube.

And exactly the same equation applies to the climate system:

$$F = \lambda \times T$$

- F is net additional energy flow in due to increased greenhouse gases.
- T is global average surface temperature increase above pre-industrial.
- λ is efficiency with which Earth gets rid of excess energy to space.

Going back to our plastic tubes (Demonstration), now we're going to uncork the deep ocean, to allow fluid to pass between these two tubes, or energy to be transported from the surface to deep ocean. This tube on the right here is much fatter than the "atmosphere and near-surface ocean" tube, because the deep ocean has a much greater heat capacity. It takes more fluid to raise the fluid level in the ocean tube, just as it takes more heat energy to raise the temperature of the deep ocean.

Imagine, to start with, the deep ocean is infinitely large, so it never warms up at all. This corresponds to me leaving this tap open so the level of fluid in this ocean tube cannot rise. What happens to the response of the fluid level in the "atmosphere and surface" tube to an increase in the flow rate from the pump? Let's give it a try.

What you see is that the level increases, but not nearly as much as before, because it doesn't need to rise as much because there are two pipes allowing it to escape rather than one.

This is what this looks like in our equation:

$$F = (k + k_2) \times h$$

- F is the additional rate of fluid flowing in, above "background" flow.
- h is the increased water depth above the initial equilibrium level.
- k is the openness of the outlet pipe flowing out to "space".
- k_2 is the openness of the outlet pipe flowing into the "ocean pool".

And this is what it looks like for the climate system:

$$F = (\lambda + \gamma) \times T$$

- F is net additional energy flow in due to increased greenhouse gases.
- T is global average surface temperature increase above pre-industrial.
- λ is efficiency with which Earth gets rid of excess energy to space.
- γ is efficiency of excess surface energy transport to the deep ocean.

The initial warming that we see in the decade or two following an increase in greenhouse gas concentrations is only about half to two-thirds of the long-term equilibrium warming if we were to stabilise concentrations forever at their new level. That fraction is hotly contested, of course, precisely because, while we can see this initial warming, we don't know what the long-term equilibrium warming will be.

But this doesn't stop us understanding what happens on much shorter timescales, a decade or two for the climate system, or just a few seconds for this Gresham Climate Model:

$$\Delta F = (k + k_2) \times \Delta h$$

- ΔF is the *change* in rate of fluid flowing in over a time-interval.
- Δh is the *change* in water depth, after initial rapid adjustment.
- k and k_2 are constant.

We've adjusted the viscosity of the fluid to make the timescales match at one minute in this room is about 120 years in the climate system, so you don't get too bored watching it.

This proportionality applies to *changes* in fluid flow as well as the total excess flow over the "pre-industrial" natural level. This big Delta symbols denote a change in fluid flow and change in fluid level after a few seconds of adjustment.

And exactly the same equation applies to changes in global temperature in response to changes in greenhouse gas concentrations on a timescale of one or two decades, when the surface and atmosphere have had a chance to adjust, but the deep ocean has yet to respond.

Let's just rearrange this equation: if fluid level is proportional to fluid flow, then fluid flow is also proportional to fluid level – I'm just replacing the fraction, one over k plus k2, with a single kappa symbol for convenience.

$$\Delta h = \kappa \times \Delta F$$

- Δh is the *change* in water depth, after initial rapid adjustment.
- ΔF is the *change* in rate of fluid flowing in over a time-interval.
- κ is a constant, $1/(k + k_2)$.

Now let's make it more interesting, and close off this tap so the deep ocean tube is no longer, in effect, infinitely large. We are now allowing the deep ocean to warm up in response to this additional energy penetrating down from the surface. But it is still quite large, so it doesn't change very fast.

Let's try some experiments to see how the fluid level in the Atmosphere and Surface tube responds to changes in fluid flow. If we suddenly increase the flow, the level jumps up just as it did when the tap was open, because the deep ocean hasn't had time to respond at all. And if we increase the flow by twice as much, it jumps up twice as much. So one component of the change in depth over a short interval is proportional to the change in flow, just as was the case when the tap was open.

But there is now something else going on as well, which is most obvious if we look at the behaviour after we have stopped changing the flow. Now the input flow is constant, corresponding to constant concentrations of greenhouse gases in the atmosphere, but the level of fluid in the atmosphere-and-surface tube is still rising. Eventually, it will come back into equilibrium again, when the level of fluid in both tubes is the same and fluid only escapes through the outlet to space.

That takes a very long time, because this second tube is quite fat and because the closer we get to equilibrium the slower the flow of the fluid through this pipe. That's analogous to deep ocean temperatures catching up with the surface: once the world has come back into equilibrium, the deep ocean will still be colder than the surface, but if all ocean temperatures have all warmed by more-or-less the same amount, there will no longer any net export of heat out of the surface layers to the deep.

On shorter timescales, analogous to decades to a century in the climate system, the level of fluid in both pipes rises at a more-or-less constant rate with constant input flow. And if we do the experiment again, doubling the increase in the input flow, then after the initial adjustment, it settles down to increase again at more-or-less twice the rate as before.

So, over a time-interval much shorter than the deep ocean adjustment time, which is centuries in the climate system, but longer than the atmosphere-and-surface-ocean adjustment time, which is only around a decade, the change in fluid level in the atmosphere-and-surface tube depends on two things: one term proportional to the change in input flow over that time-interval, and another proportional to the average flow over that time-interval multiplied by the length of the time-interval.

By now, you'll be relieved to see some more maths. Here you are:

$$\Delta h = \kappa \times \Delta F + \kappa_2 \times \bar{F} \times \Delta t$$

- Δh is the *change* in water depth, after initial rapid adjustment.
- ΔF is the *change* in rate of fluid flowing in over a time-interval.
- $\bar{F}$ is the *average* rate of fluid flowing in over that time-interval.
- Δt is the length of the time-interval.
- κ and κ_2 are constants.

This is just saying what I've just said in the form of an equation. The change in fluid level, Delta h, depends on the sum of one term that is proportional to the change in fluid flow rate, Delta F, over that time-interval, and another proportional to the average flow rate, F-bar, multiplied by the length of the time-interval. Bear with me, because this is an incredibly important equation, and it's half the story in understanding the need for net zero, so hang in there.

Here is exactly the same equation again, just replacing kappa-2 with kappa rho, where rho is a fractional rate of change, like an inflation rate.

$$\Delta h = \kappa \times (\Delta F + \rho \times \bar{F} \times \Delta t)$$

- Δh is the *change* in water depth, after initial rapid adjustment.
- ΔF is the *change* in rate of fluid flowing in over a time-interval.
- $\bar{F}$ is the *average* rate of fluid flowing in over that time-interval.
- Δt is the length of the time-interval.
- κ is a constant "Transient Level Response to Flow".
- ρ is a constant "Rate of Adjustment to Constant Flow".

If we multiply it by 100, rho has familiar units of percent per year. What does rho mean? It's the fractional rate at which the level of fluid in the atmosphere-and-surface tube continues to rise after the fluid flow is stabilized at some new faster rate following a relatively rapid increase from the initial equilibrium level, before the level of fluid in the second tube has had time to adjust very much. It's a fractional rate, which means the bigger the initial jump in the input flow, the higher the level in this atmosphere-and-surface tube immediately afterwards, and the faster, in absolute terms, the level continued to rise after the flow is stabilized. But as a fraction, it still goes up at the same rate, in percent per second.

Now comes the punchline. Here is exactly the same equation again, but now applied to the climate system.

$$\Delta T = \kappa_F \times (\Delta F + \rho_F \times \bar{F} \times \Delta t)$$

- ΔT is the *change* in global average surface temperature.
- ΔF is the *change* in energy flow in due to *change* in greenhouse gases.
- $\bar{F}$ is the *average* energy flow in due to *level* of greenhouse gases.
- Δt is the length of the time-interval.
- κ_F is a constant "Transient Climate Response to Forcing".

ρ_F is a constant "Rate of Adjustment to Constant Forcing".

The change in global temperature over a time-interval ranging from 20 years to a century or more is simply proportional to two things: first, and generally the dominant term in these days of rapidly increasing greenhouse gas concentrations, a term proportional to the change in energy imbalance due to changing atmospheric composition over that time-interval. And second, a term proportional to the total energy imbalance due to the total change in atmospheric composition since pre-industrial. This second term is dwarfed by the first one at the moment, but will become important if we ever actually manage to slow down the rate of greenhouse gas increase.

We call kappa-F, the sensitivity of the climate system to fast changes in energy imbalance, the “Transient Climate Response to Forcing”, or TCRF – changes in energy imbalance that result directly from external drivers like rising greenhouse gases are called “radiative forcing”, hence the name. The TCRF has units of degrees Celcius per Watt per square metre, or (by analogy to the Equilibrium Climate Sensitivity), degrees per CO2-doubling, given a CO2-doubling causes approximately 3.7 Watts per square metre of radiative forcing. Plausible values of the TCRF are around 1.5 to 2 degrees Celcius per CO2-doubling, so around half to two-thirds of the Equilibrium Climate Sensitivity (although, because we don’t know what the ECS is, this fraction is kind of moot).

By the way, enthusiasts may have heard of the Transient Climate Response, which is the warming we get at the time of CO2 doubling after 70 years of CO2 levels increasing at 1% per year. This is closely related to the TCRF, but about 10-15% larger, because after 70 years the deep ocean has already started to adjust. So the TCR is actually rather a messy combination of the long-term and short-term response: the TCRF is much cleaner.

This symbol rho-F represents the slow fractional adjustment to a constant radiative forcing. It is called, appropriately enough, the Rate of Adjustment to Constant Forcing, or RACF. It has units (multiplied by 100) of percent per year, and typically has values of around 0.3 percent per year, 3% per decade, or halving over 200 years. That gives you some indication of the slow adjustment timescale, and also what to expect if we were to stabilize atmospheric concentrations tomorrow. We are now at one and a quarter degrees above pre-industrial and warming at about one quarter of a degree per decade. If we were to stabilize atmospheric concentrations of everything by the time temperatures reach 1.5 degrees (which is not the same as net zero emissions), temperatures would continue to rise by about 0.3% of 1.5 degrees per year, or about 20% of the current rate of warming.

Remember how the Rio Convention committed the parties to “stabilization of atmospheric concentrations of greenhouse gases”. Well, now we know what that would achieve: if we were somehow to stabilize atmospheric concentrations by the time temperatures reach 1.5 degrees, it would slow the rate of warming by a factor of 5 from what we are experiencing today. That would be a great step, but it isn’t enough.

So, how well does this work? Well, for the global average temperature response to an energy imbalance due to changing atmospheric composition, pretty well: the dots show simulations of the response to a scenario in which we keep emissions more-or-less constant over this century (which, depressingly, is what current policies actually imply, if we don’t believe all the aspirational goals of net zero by mid-century, most of which don’t seem to be backed up by actual policies at the moment).

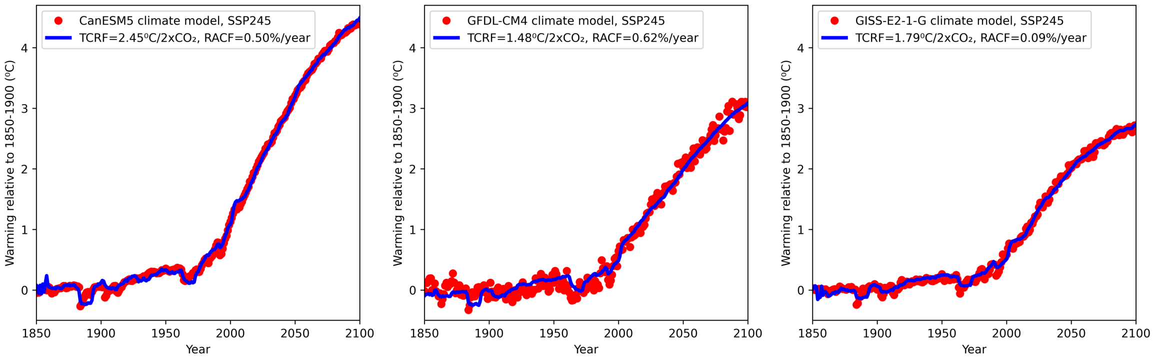

These simulations were done by some of the most advanced climate models in the world, participating in the Sixth Coupled Model Intercomparison Project for the latest assessment of the IPCC. Each represents many millions of processor-hours on the worlds most advanced supercomputers. The lines show what we get simply by taking the time-series of energy imbalance and applying our simple formula, with the values of TCRF and RACF appropriate to that model. As you can see, it works pretty well.

Just in case any politician is watching online and might be tempted to shut down the Met Office's climate program and replace it with an assemblage of plastic pipes, I must stress that these complex climate models give us much more information than just global temperature. They tell us about changing weather patterns, sea level rise, extreme weather risk and so on. But for understanding the response to a global energy imbalance on decade to century timescales, you don't need a complex climate model. You just need a simple equation.

Full-complexity climate models take millions of processor-hours per simulation

$$\Delta T = \kappa_F \times (\Delta F + \rho_F \times \bar{F} \times \Delta t)$$

This is important, because these climate models are so complex, there are plenty of ways they can go wrong. For example, if we look at the Met Office model here, you'll see how it barely warms at all over the 20th century and then takes off like a rocket, warming at almost 0.4 degrees per decade in the 2020s. That doesn't tell us what the immediate future will hold, still less cause us to revise our views on warming to date: these models are neither crystal balls nor substitutes for climate observations. They are useful tools for understanding observations and their implications for the future.

And one question we can use our new-found understanding of the climate system to answer is what is actually needed to "halt" the warming.

$$\Delta T = \kappa_F \times (\Delta F + \rho_F \times \bar{F} \times \Delta t)$$

- We want $\Delta T = 0$, meaning no further surface temperature change.
- So either $\bar{F} = 0$, meaning atmospheric composition is back to pre-industrial, or…

$$\frac{\Delta F / \Delta t}{\bar{F}} = -\rho_F$$

- So global energy imbalance needs to *decline* at a fractional rate equal to the RACF, or about 0.3% per year.
- Not constant, but declining, concentrations of greenhouse gases.

This means Delta-T has to be zero: no more change in global temperature. Of course, you might want Delta-T to be negative, to cool things off, but that is a policy decision. We are just trying to answer the geophysical question here of what would it take to halt the warming. It's an important question, because the Paris Agreement is very clear we are aiming to limit warming to well below 2 degrees, so whatever happens to 1.5 degrees, we will definitely need to halt the warming sometime this century. So, what would it take, in terms of energy imbalance? A constant energy imbalance, Delta-F equals zero, corresponding to "Rio-stabilization", is not enough, unless we first reduce all greenhouse gas concentrations back to pre-industrial levels, making F-bar zero. But if we rearrange this equation again, we see that if we can get the energy imbalance to decline by about 0.3% per year, or 3% per decade, then these two terms balance each other and we get no further warming.

We can see this work in the Gresham Climate Model. If we set the pump speed to increase relatively rapidly, but then decline by the equivalent of 0.3% per year, so halving over 50 seconds, the declining pump speed, corresponding to declining energy imbalance due to increased greenhouse gases, balances the adjustment of the fluid in the deep-ocean tube and we get no further increase in the level of the fluid in the atmosphere-and-surface tube. So, to halt global warming, we need overall concentrations of greenhouse gases, aggregated in terms of the total equivalent concentration of carbon dioxide, to decline by about 0.3% per year. As you can see, that is not a climate in equilibrium: our climate will not be in equilibrium again for many, many centuries, and probably not even then if we include the adjustment of the ice sheets. But it is enough to halt the rise in global average surface temperature, which is what drives most climate impacts.

So, there we have it. A simple equation that turns out to be extremely useful relating changes in global energy imbalance due to atmospheric composition to global temperatures.

- Changes in temperature over decade to century timescales:

$$\Delta T = \kappa_F \times (\Delta F + \rho_F \times \bar{F} \times \Delta t)$$

- Warming over a multi-decade time-interval is proportional to:
    1. the *change* in global energy imbalance due to any *change* in greenhouse gas levels over that time-interval plus
    2. a contribution from the gradual adjustment to the current *level* of energy imbalance relative to pre-industrial.

The enthusiasts among you will no doubt be wondering, “hang on, that can’t be the whole story, because the climate can’t respond instantaneously to a change in energy imbalance, and nor does it continue to warm indefinitely in response to a constant energy imbalance.” This is true: this equation only refers to decade to century timescales. On shorter timescales, we need to understand the fast response, and internal climate variability generally dominates anyway. On multi-century timescales, we need to worry about lots of other things as well, like melting ice-caps and changing patterns of vegetation.

If you want a more comprehensive model, here it is – you’ll find, if you assume the near-surface heat capacity is much less than the deep ocean’s, which it is, then to a good approximation this reduces to the simple equation on decade to century timescales. There is only one term here that we haven’t seen before, which is this additional energy lost to space due to the climate being out of equilibrium. This is important, because if we use today’s energy budget to predict the equilibrium climate sensitivity, and don’t take into account the changing pattern of warming as the climate warms, we potentially get the wrong answer. This is another reason why the equilibrium climate sensitivity is so uncertain.

Climate models predict quite pronounced changes in the pattern of warming as the climate system re-equilibrates – but, of course, we have no way of testing these with direct observations, although observations of past climates can help. So, we have once again failed to pin down the equilibrium climate sensitivity, but it turns out it doesn’t matter! To work out how much warming to expect over the coming decades, we mostly just need to know the Transient Climate Response to Forcing, which we can observe pretty much directly from changing atmospheric composition and the current rate of warming. To know what it will take to halt the warming, we need to know the Rate of Adjustment to Constant Forcing. Here, we depend on models more, but we do know it has to be around 0.3% per year. We also have to recognise that we can model until we are blue in the face, but we won’t really know for sure what it will take to halt the warming until we significantly reduce the rate of increase in greenhouse gas concentrations: which means big cuts in emissions. And that will be the subject of our next lecture.

References and Further Reading

Archer David, Global Warming: Understanding the Forecast, John Wiley & Sons; 1st edition (18 Oct. 2006)

Myles R. Allen, Pierre Friedlingstein, Cécile A.J. Girardin, Stuart Jenkins, Yadvinder Malhi, Eli Mitchell-Larson, Glen P. Peters, Lavanya Rajamani

Annual Review of Environment and Resources 2022 47:1, 849-887