CONVERTING CUMULATIVE CARBON BUDGETS TO ANNUAL EMISSION REDUCTIONS

The central thread wending its way through this book is the necessity of adhering to a carbon budget, the maximum cumulative carbon emissions that preserve a better-than-even chance of keeping global warming below a specified limit. The book adopts a 2°C target for global average temperature above its preindustrial (mid-nineteenth-century) level; 1.5° would be better if we could manage it, but it is not likely we can. This corresponds to a budget of 1,070 gigatonnes (Gt) of carbon dioxide as of 2019, and it applies to the world as a whole.

Alligators argues that the first step any country should take is to announce the share of this global budget it wishes to claim for its own future emissions, followed by putting into place policies designed to keep emissions below this limit. A key element is translating the cumulative budget into a sequence of annual reductions. This is a fantastically complex problem, since the amount of each year's proposed emissions logically depends on the whole future pathway extending to a planning horizon of many decades (at least), and that in turn depends on a range of expectations, especially for new technologies that are just hopes today but may become realities at some point in the future.

For obvious reasons, I am reluctant to propose specific emission paths for individual countries or the world as a whole. Were the direction advocated in this book to be adopted, commissions or expert panels would be created to do that job. No doubt they would be influenced by political as well as economic and environmental considerations. There is little point to predicting in advance what choices they will make or how likely they will be to meet the budget constraint.

Nevertheless, when the book discusses the cost of decarbonization and the political economy dimension of the problem, it's helpful to have a working approximation of the stringency of emission cuts we face, so some method for translating budgets into annual targets is unavoidable. One approach would be to appeal to the integrated assessment models (IAMs) that generate scenarios based on simplified representations of economic and environmental interrelationships, filled out by a host of assumptions about future production and consumption options. There are two problems with using the IAMs, however.

First, as argued in Chapter 5, while the patterns that emerge over many models and runs have much to tell us about the challenges facing climate policy, little credence can be placed in any one scenario. Second, their reliance on a multitude of assumptions, separately and in combination with each other, renders them too opaque for our purposes. In fact, you will look in vain for any single, comprehensive list of all the assumptions that a single model run incorporates; it would be too long and too complexly related to the run's output. This entanglement of predictive and modeling assumptions lurking beneath the hood underlies the long discussion in Chapter 5 of the role of carbon removal in the mitigation scenario literature: just one set of assumptions is placed under the microscope, and its effect varies enormously depending on what model employs it and what other assumptions it interacts with.

Since so much of the book's argument rests on the general speed and scale of the required emission cuts, I have resorted to a simple strategy whose purpose is not to "solve" the pathway problem but only to benchmark it. By this, I mean that it provides a rough indication of how rapidly decarbonization needs to proceed, and it provides a single metric for comparing the stringency of different mitigation scenarios. In fact, it is a type of solution, but only under a narrow set of conditions.

The benchmarking formula for converting cumulative budgets into annual reductions appears first toward the end of Chapter 3, in the form of a rough estimate of what the world's remaining carbon budget might be in the year 2024, the number of years it would take for 2023's emissions to exhaust it, and the fixed rate of emission reduction that would keep this years-to-exhaustion number constant. In the coming pages we will look at the assumptions underlying this approach and how the annual reduction formula is derived. We will then use this formula to shed light on two additional issues: the cost of delaying action and the thorny problem of allocating a global carbon budget to individual countries.

Underlying assumptions: Above all, the benchmarking method rests on the assumption that the mandatory carbon budget is known at the outset of the program. This sounds straightforward but is actually rather far-reaching. In situations likely to arise in real life, those proposing an emission reduction path may begin with an agreed-upon budget, but they may also take into account the possibility the budget will be altered in the future. If they think it is much more likely the budget will be tightened than relaxed, it would make sense to make larger near-term cuts, so the population down the road is not faced with overwhelming decarbonization demands – and, of course, vice versa. So the initial assumption needs to be qualified: the budget is known at the outset, and either there are no expectations at all about future adjustments to it or the expectations are balanced in the sense that they don't call for any anticipatory policy modifications.

Just as important, and certainly far more controversial, is the assumption that carbon removal at a scale permitting net zero or net negative emissions is not an option climate policy can rely on. Ruling out net negative emissions is obviously a precondition for setting and adhering to a carbon budget; otherwise, it would be possible to exceed the budget (overshoot) and then claw back to the 2° (or other) warming target by massive carbon withdrawals from the atmosphere later on. In such a scenario the budget ceases to be a constraint. But net zero emissions would also undermine the logic of the formula, and it is important to understand why.

It is extremely unlikely that human beings will eliminate carbon emissions from fossil fuels and other energy or industrial processes in their entirety; there will always be some greenhouse gases generated by a global economy sustaining billions of people. Taken by itself, this would indicate that the carbon budget should never be exhausted. Cumulative emissions might approach it asymptotically, however, getting closer and closer but never quite reaching the limit. If carbon removal is expected to be at a sufficient scale to fully offset these residual emissions, on the other hand, the budget *can* be exhausted, so there is no need for tiptoeing asymptotically as we approach it. The equal-years-to-exhaustion condition embodied in the formula is a way to implement that asymptotic feature, and it holds no value if we can simply aim at net zero instead.

Of course, most of the discussion about climate policy in recent years has been predicated on the belief that net zero emissions are or will soon be feasible. The horse-trading at the Glasgow Conference in 2021, for instance, was conducted largely in terms of net zero pledges, and IPCC climate modelers have, perhaps reluctantly, taken it on as a framework for analyzing mitigation scenarios. As far as I can tell, the process began with the realization, documented in Chapter 5 of *Alligators*, that massive carbon removal, most likely involving overshooting, was required for model runs to meet acceptable warming targets, and this evolved into a *reliance* on the future feasibility of carbon removal at such a scale in policy research and politics. This shift occurred without any concomitant change in the probability that carbon removal will actually live up to this billing.

Still, it is entirely possible that carbon removal will make zero or even negative emissions possible in the decades to come. Perhaps the most precise way to state the assumption underlying the approach taken in the book is to say that it is not conditional on the availability of massive, affordable carbon removal. If future technologies give us this option, fine; we can decarbonize less and more slowly. But for now it would be too risky to proceed as if we were sure that these technologies will develop as hoped.[1]

These two assumptions justify formulating an emissions path that stays under budget and never uses it up. But there are an infinity of such paths; why adopt the one generated by the formula used in the book? One answer is just that it is simple and therefore transparent. There is no mystery about how the rate of emission reduction is calculated, nor how changes in the available budget alter allowable annual emissions. That would be enough reason to use it. But there is an additional argument worth considering, that under a narrow but perhaps plausible set of conditions an equal-years-to-exhaustion path represents a form of intergenerational equity.

The conditions are these: (1) The required year-to-year percentage rate of emission reduction is a valid measurement of the burden these reductions place on those who must achieve them. (2) No technological changes are forecast that will make it easier for future generations to achieve a given percentage reduction in emissions than it is for us today. (3) No relative absence of technological change will make it more difficult for future generations to maintain this steady rate. If these hold, by establishing a constant rate of annual emission reduction that doesn't exhaust the mandatory carbon budget, we are treating all people, present and future, equitably.[2]

Even if we regard the argument from intergenerational equity as a bit overstretched, however, the simplicity of the book's benchmarking suggestion is sufficient to justify it. As discussed in Chapter 3, it is likely that societies will choose to either front- or backload their emission cuts based on what they see as the technological landscape, and this can be approximated by building up or drawing down the years-to-exhaustion ratio, after which it can be adhered to indefinitely.

Derivation: To understand the math behind these calculations, consider the condition we assume to hold between any two years of our benchmark mitigation program:

$$\frac{B}{e} = \frac{B - e}{e(1 - r)} \tag{1}$$

where B is the carbon budget in the first year, e is the amount of carbon emissions in the first year, and r is the rate of emission reduction.[3] On the left hand side we have the number of years to exhaustion in the first year, the carbon budget divided by that year's emissions. The numerator on the right hand side is the remaining carbon budget in the second year after the previous year's emissions have been deducted. The denominator is the level of emissions in the second year, which has gone down by the rate r. The two are equal: years-to-exhaustion remains constant.

Solve this for r and you get

$$r = \frac{e}{B} \tag{2}$$

The benchmark rate of emission reduction is simply the inverse of the number of years to exhaustion. In Chapter 3 we looked at the case of a budget of 890.5 Gt of carbon dioxide at the outset of 2024, where 2023's emission had been 36 Gt. Years-to-exhaustion is thus 890.5 ÷ 36, or just under 25; the inverse of this – our rate of constant annual emission reduction – is just under 3.9%.

To understand better what the constant years-to-exhaustion condition implies, consider Figure 1. The solid line shows the cumulative global carbon emissions under the benchmark decarbonization program, beginning with 34.6 Gt in 2024; the dashed line is the

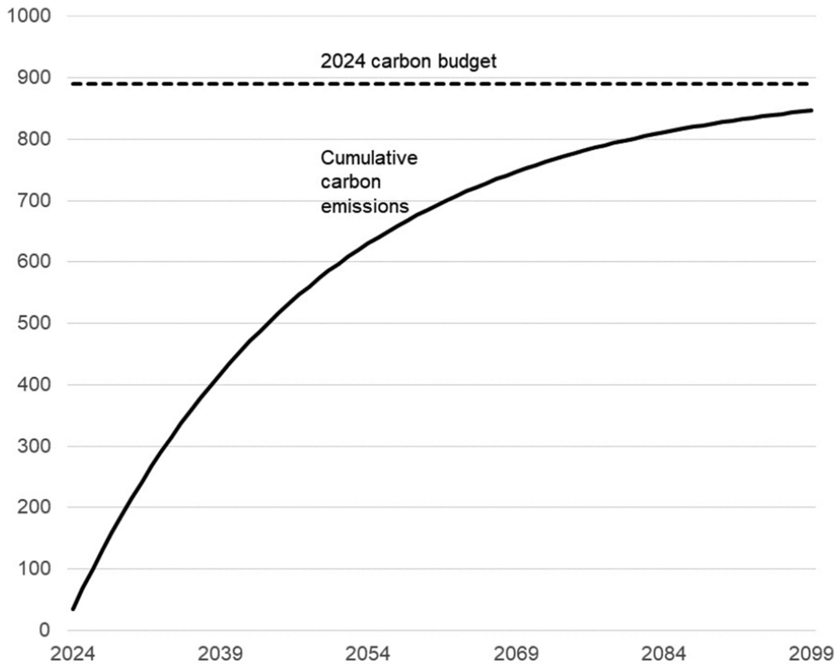

Figure 1 Hypothetical 2024 global carbon budget and cumulative emissions in gigatonnes of carbon dioxide.

available carbon budget as of the start of 2024. Note that, instead of reaching zero in just a few decades, emissions taper off only gradually, and by 2050 about 318 Gt of the budget is still unused. Even at the end of the century, 44 Gt are available.

From the standpoint of making as full a use of the IPCC's carbon budget as possible, this approach is absurd, since we would be imposing substantial costs on ourselves to keep emissions down even though we could temporarily emit more without exceeding the cumulative limit. On the other hand, a path like that depicted in Figure 1 is exactly the sort we should take if we are unwilling to assume that either carbon removal or the complete replacement of fossil energy will be an option in the foreseeable future, for in that case we should leave ourselves enough space under budget so the pressure to reduce further in any given future year is no greater than it is today. Since each argument is valid in its own way, I would expect thoughtful policy deliberation to generate some intermediate path, using somewhat more of the budget over coming decades than the solid line in Figure 1, but still leaving enough buffer to avoid a horizon shock. (The myopic bias of politics alone would push us in that direction.)

The cost of delay: If we continue to use the benchmark formula above, we can shed light on other aspects of climate mitigation. For example, a commonly asked question is, What is the effect of delaying action on decarbonization? It is straightforward to calculate how the constant emission reduction rate, represented by r, responds to changes in cumulative budgets. Consider the long postponement we have already experienced in responding to the climate challenge. As of the publication of this book, it is 2022 and the job has barely begun. Yet thirty years ago, in 1992, the leaders of over 150 countries gathered in Rio de Janeiro and agreed to the United Nations Framework Convention on Climate Change. What if an effective emission reduction program had been instituted right then and there?

Recall that our working carbon dioxide budget is 1070 Gt as of the beginning of 2019. Between 1992 and 2019 the world emitted about 794 Gt of carbon dioxide, so add this to get the budget at the start of 1992: 1864 Gt.[4] Emissions in 1991 were 23.2 Gt, so years-to-exhaustion was a luxurious 80.4, and a constant annual reduction rate of 1.2% would have been sufficient. Of course, eighty years was much more than we would need to ensure a future buffer of safety, so it would have been possible to permit several years of modest carbon energy growth before applying the brakes. Our difficult situation today is the result of not a few years but decades of inaction.

We can also apply the same technique to see what it would mean to delay action even further. One often hears that we have only so many years to achieve a proposed level of emission reduction, like half, or the battle is over. Claims like this are well intentioned but misleading; there is no magic moment that separates just in time from beyond hope. Instead, our arithmetic shows a continually worsening cost of procrastination.

For example, in the calculation above, we found that that, with plausible assumptions about the next few years of carbon emissions, the world would need to adopt a path of 3.9% annual emission reductions beginning in 2024. Suppose, however, that policies are just as paralyzed then as they are now, and it takes a further six years to get decarbonization on track. Suppose also that in each of those years emissions remains at 36 Gt. Our budget on New Year's Day, 2030, would then be 674.5 Gt, leaving us less than nineteen years to exhaustion, and our new required annual reduction rate would be 5.3%. To put it as clearly as possible, the cost of delaying action from 2024 to 2030 would be an increase in 1.4% in the rate we would have to cut emissions. If that sounds small, consider that this difference represents more than an additional third of the pace we would have been committed to without the added delay.

The effect of changing the national share of a global budget: This same math also applies to individual countries based on the share of the global carbon budget they take responsibility for. We can use the United States as a case in point. If the global budget is 890.5 is 2024, how much might be allocated to the United States? One possible answer is to take the US share of global emissions as the basis for its budget. If every country did that, of course, the required annual rate of reduction would be the same everywhere – in India and Ethiopia as in the United States and the EU. That would fail the most minimal test of fairness: it wouldn't take into account either the vastly unequal amount of the Earth's carbon space already used up by the wealthiest countries or the claim of lower-income countries to a greater share of future resources for development.

So a different principle we might adopt is that there should be equal carbon space for each human being. In that case, we would begin with the observation that the United States accounts for about 4.25% of the world's population. If we apply *that* share to the global carbon budget, the United States would be allocated just 43.8 Gt at the start of 2020: 4.25% of the global 1,033 Gt. With the United States emitting almost 5.3 Gt in 2019, the required annual rate of reduction would be a staggering 12%. (I use 2019 data because of the distorting effects of the pandemic in 2020, and to avoid assuming hypothetical future emissions I begin reductions as of the end of 2019.)

What does such an impossible number mean? Essentially, it would add to the general task of decarbonization all countries face an additional obligation for the United States to reverse its historical "overemission" relative to the world average. It has a certain abstract ethical appeal, but it is not remotely feasible. What such a calculation does demonstrate, in my opinion, is a context for thinking about how much of the global carbon budget the United States can claim without embarrassment. If an equal rate of emission reduction across the planet is 3.9% a year, and if the corresponding rate on the basis of one person, one share, of carbon is 12%, the 50% increase in the rate for a country like the United States suggested in Chapter 3 can be seen as something of a bargain.

Notes

[1] Carbon removal at a scale that does not allow for net zero or negative emissions does not require changing the formula so long as it is not anticipated. If there is a "surprise" leap in technology that permits, say, 5 Gt of carbon dioxide to be removed each year, we can simply add a corresponding amount of emissions, so that net emissions remain on course.

[2] This abstracts from the argument discussed in the Appendix to Chapter 2 according to which allowance should be made for the greater well-being expected of future generations compared to the present. Although it is routinely invoked by economists analyzing the discounting problem, this is simply an assumption, and it is called into question by the effects of climate change themselves.

[3] We could also express this condition in the context of a continuous growth function, but there really isn't any point, since carbon emission accounting is conducted in units of time duration such as years.

[4] These emission data, like the others in this addendum and in Chapter 3, come from the Global Carbon Project (2021).