Continuing my discussion of the report of the Climate Overshoot Commission released last week, today I dig into their recommendations on mitigation. As you may recall, the Commission's informal (but serious) job description was to speak of elephants in the room and unclothed emperors: to say things that are true and important about climate risks and responses that other, more political constrained bodies cannot. If you take this job description for statements and apply it to recommendations, it would suggest recommending things that are not politically feasible – at least not now – or that even lie outside the range of current debate. This does not mean making recommendations so outlandish or implausible that they can readily be ignored or arbitrarily rejected, of course. But if the job is to move the range of acceptable arguments and proposals – moving the Overton window, as the political scientists say – the most effective recommendations may well lie beyond the boundary of what could be adopted now. This perspective is especially relevant to the Commission's recommendations on mitigation.

Mitigation – deep rapid cuts to worldwide emissions – is the first, essential element of effective climate response. I don't think there's anyone thinking seriously about climate change who disagrees with this. In the Commission's words, mitigation is the "foundational strategy." Yet when the Commission began its work, it first planned not to speak about mitigation – not because they didn't recognize its primary importance, but because they thought there wasn't much for them to add to what's already being said, particularly given the tight time limit on their work. But partway through, the Commissioners realized that not speaking on mitigation would risk them being mistakenly seen to not accord it the needed priority, so they changed course – correctly, even necessarily, in my view. But in making this decision, they also resolved that their messages on mitigation had to cut through the noise and move the debate, and thus sought to make their recommendations radical. I think they succeeded at this, although it's not clear from the initial reactions to their report that their radicalism has been

noticed – yet.

Their mitigation recommendations include calls to adopt stronger national and international accountability mechanisms for emissions cuts; policy and financing innovations to promote faster deployment of zero-emissions technologies; and for countries to recognize each other's climate policies and reflect them in trade measures. They also call for cutting short-lived climate forcers even faster than now being pursued. These are strong recommendations, persuasive and well conceived. But they also could plausibly be adopted within a few years if governments are serious about ramping up their ambition, so do not necessarily meet the aim of proposing something radical enough to move the debate.

So, where's the radicalism?

It's in their very first mitigation recommendation, for a "graduated, differentiated phaseout" in production and consumption of fossil fuels. Wait a second, you might say, what's so radical about that? Isn't it obvious that the world needs to get rid of fossil fuels, and haven't a bunch of people called for it? Well, yes. But the Commission's proposal is vastly stronger than either the weak language adopted at Glasgow – which calls on parties to "… accelerat(e) efforts towards the phasedown of unabated coal power and phase-out of inefficient fossil fuel subsidies …" or the language now being discussed for the coming COP28, which speaks of phasing out ***unabated*** fossil fuels. The word "unabated" has been used frequently in recent months by Sultan al-Jaber of the United Arab Emirates who is overseeing this year's COP; it was included in a draft document by EU countries; and it appears in the mitigation findings of the global stocktake released earlier this month. The Commission's proposal is also substantially stronger, and at the same time more practical, than the most ambitious fossil-fuel proposals being promoted by activist nations: the Fossil Fuel Non-Proliferation Treaty and the Beyond Oil and Gas Alliance.

What makes the Commission's proposal so radical is the combination of its ambition; its inclusion of key design elements that make it plausibly operationalizable; and the stature of the Commission. No mitigation proposal remotely this strong has been advanced in policy debate, certainly not by any body with stature similar to the Commission's.

It is these elements taken together that make the Commission's proposal radical. Even better, they might position the proposal in that sweet spot that actually makes

change, where the “radical and transformative” overlaps with the “close enough to feasible that it can’t be arbitrarily dismissed or quietly ignored.”

There are three points of the proposal that support this claim. First, it’s expressed in terms of production and consumption of fossil fuels, not emissions or net emissions. Second, it calls for a phaseout of these. And third, in a few key design elements it is modeled on the provisions to cut CFCs in the Montreal Protocol – the only international environmental regime that has achieved a socio-technical transformation of remotely the scale required for greenhouse gases. I’ll discuss the first two points briefly, then unpack the third in a bit more detail.

Targeting fossil fuels not emissions.

First, the Commission’s proposal does not target greenhouse-gas emissions (or even worse, net emissions), but instead targets a concrete, readily observable proxy for emissions, fossil fuels – and not just some fossil fuels, all of them. Fossil fuels are responsible for the dominant share of emissions – more than two-thirds of anthropogenic emissions, although the arithmetic to precisely estimate this fraction is confused by the fact that some human emissions cool as well as heat. (And for even more confusion, those cooling emissions are linked to fossil-fuel combustion when the fuel contains sulfur– those fine particles I discussed in my last post that are giving us a cooling bonus, while also killing millions of people a year.) In addition to being a strong proxy for greenhouse-gas emissions, fossil-fuel production and trade is observed, documented, and regulated in nearly all jurisdictions.

To understand why targeting fossil fuels instead of emissions is so important, you need to simultaneously hold two ideas in your head that may seem contradictory, but are both true. First, what the Earth system sees and what actually matters for climate change is net emissions. This isn’t even net anthropogenic emissions: it’s net emissions from all sources and sinks, including human ones, natural ones, and all the feedbacks that link the two. Second, if you are serious about actually making large changes in human environmental burdens, in the context of a massively complex linked human-natural system, “net emissions” is an utterly impractical metric to use for control. Even total emissions from a country or legal entity (without the “net”) can’t be observed directly, but must be inferred from inventory processes that are chock-full of too-coarse, out-of-date, suspect, manipulable assumptions, and dependent on weak, mostly voluntary, and inconsistently controlled reporting processes. Take the additional step to “net emissions” (my

enterprise's, my industry's, my city's, or my country's), and it's everyone make your own adventure – an accounting artifact chock-full of opportunities for double-counting, unobservable and inflated counter-factuals, and enormous effort (and expense) spent not to actually reduce or remove emissions, but to ensure that the removals get credited to my books and the emissions debited to somebody else's. That net emissions has emerged as the dominant metric to measure and claim credit for climate action – when it is so full of holes, so impractical for counting anything you actually care about – is mystifying. Nothing in this rant negates the first proposition above, that net emissions really is the measure that matters for the climate. It is. It's just impossible to measure and control, and likely to remain so for as long as it matters for actually getting emissions down. The Commission's proposal is the first from a body of its stature that holds the promise of cutting through this mountain of delusion.

A phase-out, not a phase-down.

Second, the proposal calls for these fossil fuels to be phased out – not "phased down," as in the weak final language of the COP26 Glasgow Climate Pact; nor a phaseout of "unabated" fossil fuels, as is now being discussed in the run-up to COP28. A "phasedown" is intentionally ambiguous about how far it goes. A phaseout of "unabated" fossil fuels shifts the target away from the fuels themselves toward some undefined subset that are "unabated" (with plenty of slop over what "abated" actually means). This in turn centers the argument back on carbon capture and removal instead of reductions, and shifts the control metric back toward "net emissions" – even though it's ostensibly about fossil fuels. In contrast, the Commission's proposal puts the zero target on the fuels themselves, without escape language built into the target. Knowing that something is really going away has a wonderfully clarifying effect. As Samuel Johnson said of knowing you are going to be hanged in a fortnight, it "concentrates the mind wonderfully." It doesn't matter that the zero endpoint is not immediate. It doesn't even matter much if the timeline to get there is pretty long relative to what optimally prudent climate control might call for. It can, for example, be long enough to allow capital recovery for many investments made before adoption of the phaseout, as the Commission notes. The crucial factor is adopting a credible early commitment to move to zero, even if it takes a while to get there. This will immediately convey that the sector is now a terrible business proposition and chill new investment accordingly. It will thus achieve the top-line aims of both the Fossil-Fuel NPT proposal and the Beyond Oil and Gas Alliance – quickly stopping development of new or expanded production –

but in a way that harnesses rather than fighting against market forces.

Invoking the successful Montreal Protocol.

Third, the proposal is modeled on, and adopts a few key design elements from the extraordinarily successful provisions of the Montreal Protocol to cut CFCs and other ozone-depleting chemicals. Picking up these elements invokes the powerful symbolism of the only global environmental regime to have achieved a socio-technological transformation of remotely the scale required for greenhouse gases. In addition, and more concretely, these elements give credibility and force to the proposal, because they are the main points that enabled the Montreal Protocol to achieve what many thought impossible, and do it faster than even the most optimistic expected.

There is a lot to unpack in the parallels between the Commission's proposal and the Montreal Protocol, far too much to do justice to in one post. And there are some real problems in porting the model from ozone-depleting chemicals to fossil-fuels, which will need serious effort to work out and may ultimately turn out to be unresolvable. But there are good reasons to hope that with creative design and strategy, these elements may be transportable. To the extent they are, they may make the proposal something that governments actually can do – and thus hold the promise of driving bigger and faster changes than any mitigation proposal currently being considered.

With apologies for putting numbered lists within numbered lists, I'll unpack and discuss four of these design elements that the Commission proposal takes from the Montreal Protocol.

- First, the proposal phases out both production and consumption of fossil fuels, on the same schedule. Other current proposals target only production. But trying to eliminate a socially destructive product by squeezing only on the supply side is a self-defeating strategy, as decades of failed US policy to control illegal drugs from the supply side have shown. This "both-sides" strategy also reflects a key insight gained partway through the Montreal Protocol negotiations, that fighting over whether responsibility for a damaging product – and the initial duty to cut – lies with the producer or the consumer can be a long, fruitless battle, particularly when you know the eventual endpoint is zero. After all, at that point no one will be either producing or consuming, and all the burden-shifting fights along the way just delayed you

getting there. Moreover, to avoid additional fruitless fights over how to define and measure consumption (which is in fact harder than for production), the Protocol adopted the simplest possible definition: national consumption equals production plus imports minus exports in a single year. Yes, this neglects change in inventories, but they judged it close enough and it worked. This aspect of the Commission's proposal cuts through the entire thicket of problems and conflicts associated with emissions accounting and inventory procedures. It eliminates incentives to move emissions around rather than reducing them. And it makes it impossible for fossil-producers – whether firms or exporting jurisdictions – to claim that the bulk of the emissions from their products are their customers' problem because it shows up on their books. Nope, now it's both of your problems.

- Second, the proposal makes gradual phased reductions, which start out easy – a freeze at current levels – then get harder on a practical time-scale. I've already mentioned the bracingly clarifying effect of explicitly making the eventual target zero. But with that established, you can reduce near-term disruption by making the first control step real and binding, but not particularly difficult. Taking that first real control step – not just setting up a process, not promising future cuts from some manipulable baseline – has a motivating and action-focusing power almost independent of its quantitative level. This insight comes from something the negotiators of the Montreal Protocol got wrong. For months they haggled over whether the controls in the first-generation 1987 treaty should be a freeze, a 20% cut, or a 50% cut (they settled on 50%), but in retrospect the difference between these mattered little. The correct schedule, or the maximum feasible schedule, can't be known in advance: these are moving targets that change along the way as you make the reduction efforts and see the results. Rather than the initial control level or schedule of cuts, the crucial feature – of both the Protocol and the Commission proposal – is the requirement to repeatedly revisit and adjust the schedule based on new knowledge and technical progress. In the Montreal experience – which I suspect will likely be replicated for fossil fuels – the work to meet the earlier, less demanding cuts repeatedly made meeting the later, tighter ones easier. I guess "seek and ye shall find" applies to environmental innovation as well as spiritual insight. This adjustment mechanism in the Protocol is often called a "ratchet," but I find the analogy a bit misleading. When you're pulling something with a cable and ratchet, it often gets harder to pull with each click. But under the Protocol later steps sometimes turned out easier than earlier ones, because innovation around the controlled materials outpaced the

tightening limits. Similar processes may well operate here.

- Third, the proposal is differentiated by development status. The reduction and phaseout schedule builds in a substantial grace period for developing countries, even allowing near-term expansion. This is necessary both for reasons of equity and pragmatism. So long as the global total comes down, it's OK to share room under the near-term constraint with lower-income countries. This does of course imply the need for even faster cuts in the industrialized countries. The Commission calls for rich countries to provide this headroom by expanding their removals so rapidly that they move far into net-negative territory by mid-century. (Note: I'm not forgetting the serious accounting and credibility problems with net-emissions accounting that I complained about above. But the worst thing about these is that they provide potentially unlimited room for fossil-fuel continuance and expansion, problems that will be less severe in the context of ongoing deep cuts to these fuels.) Rightly wary of the risk that his headroom may create a new class of entrenched Petro-states, the Commission stresses that this differentiation is not a blank check: it is a time-limited opportunity to access resources for development, after which everyone is on the same train to zero.
- Finally, perhaps most important, the phaseout comes with an exemption for essential uses. This is modeled on the Montreal Protocol process managed by the Technology and Economics Assessment Panel (TEAP). It provides the pressure-relief valve that keeps the program viable as cuts start to bite. That pressure-relief valve in turn protects the program from overwhelming pushback, and thus keeps the continued movement toward zero credible. But it achieves this relief in a fundamentally different way from the current approach on GHGs, which changes the goals from emissions to net emissions via "net-zero" targets. Net-emissions targeting and scoring has no mechanism to reliably tell if net-zero really is net-zero, and no criteria to judge how many emissions, and which ones, get to continue because they are (supposedly) offset by real, long-term-stable, not double-counted removals somewhere else. In place of this free-for-all, an essential-use-exemption process makes elimination the default and requires making an affirmative case that specific uses are in fact essential. It puts these determinations in an impartial, technically competent, international process. And crucially, it make exemptions time-limited, not perpetual, so exempted uses must periodically make the case again that they are still essential, and no viable alternatives are available, as technology advances.

The Commission’s proposal is a sketch of a promising new approach. Filling in the details and making it viable will present a ton of challenges – as is clear just from a glance at the enormous differences between fluorochemical markets in the 1990’s and 2000’s and fossil-fuel markets today. I see three challenges as biggest, which are quite tightly linked to each other. First, if the proposal is first adopted by a highly motivated subset of nations, how can its terms be designed to give others incentives to join, so the system grows, fast enough, toward global or near-global participation? The Protocol succeeded brilliantly on this score, but under such starkly different conditions that it’s not clear the tools of this success can be directly adapted. Second, how can the proposal handle international trade, in fossil fuels and perhaps in fossil-intensive products? The Montreal Protocol included tight trade restrictions on controlled chemicals, mainly as incentives for new countries to join. It also included carefully drawn threats of much more extreme trade restrictions, which many observers argued were not really credible – except that they worked, so their credibility was not actually put to the test, so they must have been sufficiently credible. (Yes, there’s a little circularity in that, or perhaps a little magic – Thank you, Tom Schelling, we miss you.). But the trade at issue was vastly smaller and simpler than global trade in fossil fuels today. The Commission’s proposal discusses incorporating reciprocal recognition of climate policies into trade measures. But it does not specify how to design accompanying trade measures that would be feasible, enforceable (or sufficiently credible as threats), and not unjustly obstruct development. And that of course is the third linked challenge: differentiating the commitments and trade measures in a way that is compatible with development needs, acceptably equitable, and still achieves the targeted global cuts.

An additional big challenge, less linked to the others, concerns administrative capacity and scale. The Montreal essential-use process represented an enormous effort. Its workload was eased by the similarity of many uses and alternatives across the world – although these were often similar but not identical, especially once the Protocol had to deal with methyl bromide, an ozone-depleting chemical used in agriculture. Fossil uses are vastly larger and more diverse, although the same few, relatively homogeneous uses are repeatedly proposed as hardest to cut: e.g., some transport modes (especially aviation) and some industrial processes (steel, cement, a few others). And even for these, there is vigorous research underway to innovate around the need for fossil fuels. It’s not beyond imagining that an adapted essential-use process could make sound technical judgments about uses and alternatives, even at the needed volume. The scale, stakes, and incentives for

multiple actors to tweak the process to their advantage will be serious challenges: but we won't know how serious, or how they might be surmounted, until someone does the work to try to design a viable process.

These challenges are complex and severe. Some may be insurmountable and make the Commission's proposal ultimately unworkable. But this is not obvious *a priori*. The work needs to be done, with the right mix of extreme ambition, creativity, and analytic rigor.

I hope I've persuaded you that the Commission's proposal is constructively radical, in that sweet spot where what looks initially outrageous in fact merits a close look, and might turn out to be workable, in whole or in part – and to move the domain of the possible. As a proposal for a concrete, high-impact, coordinated action adopted by multiple states, it does sit rather awkwardly with the Paris process. This does not mean it's incompatible with Paris, of course. Stronger actions like this could be coordinated under the Paris umbrella, so long as the leading nations willing to move get to do so and adopt elements to induce others to follow, without being held back by lowest-common-denominator consensus requirements that in effect give vetoes to opponents of strong action. But it is rather culturally incompatible with the process-heavy weak tea that the Paris process has actually delivered thus far.

All in, the most puzzling thing about this Commission proposal is that no one seems to have noticed its radicalism. With a couple of notable exceptions, initial coverage and reaction to the Commission has mainly focused on their discussion of SRM – which is a more obvious advance in the Commission's work, and which I will discuss in my next post. The most critical reactions to the Commission have, unsurprisingly, focused exclusively on its discussion of SRM. Strangely, these critical reactions all state that deep emissions cuts are the first essential climate response — an observation with which the Commission report obviously agrees. But they also advance the familiar argument that talking about SRM can be a cover or excuse for weakness on mitigation – a point which is also carefully made in the Commission report — while seeming not to notice the radical ambition of the Commission's recommendations on mitigation. Perhaps this recommendation will be a sleeper, which does have a transformative impact but takes a while to build up to it. I hope so. Go back, give it a close read, and see what you think.