[SLIDE]

I will begin by placing the issue of energy and climate in an ethical context. I'll then briefly review our current understanding of the issue. I will draw conclusions about the implications for the electricity sector, and I will review the possible strategies to deal with those implications. Finally, I will summarize what I believe is the only responsible way forward.

Two hundred years ago, Thomas Jefferson reflected upon the awesome inevitability of Justice when he wrote:

"I tremble for my country when I reflect that God is just; that His justice cannot sleep forever."

[SLIDE]

And the stated mission of this University identifies Justice as the object through which learning becomes devotion to God when it says:

"The aim is to create a sense of human solidarity and concern for the common good that will bear fruit as learning becomes service to justice."

Anthropogenic climate disruption is an issue of justice. Our failure to take appropriate measures is an insult to Justice, made all the more acute by our collective failure to recognize it as such. Our global energy system is the primary engine of anthropogenic climate change, therefore *Justice* demands that we change the way we produce energy – all energy – as quickly as it is humanly possible to do so. Priority number one is coal.

What is the case for climate change as an issue of justice? It is rests on at least two principles that are not only core Judeo-Christian values but that were central to the thinking of the Founding Fathers at the birth of this Nation – the limitations on individual rights imposed by the very source of those rights, natural law; and intergenerational equity. It is also an issue of justice by virtue of its implications for the continued advancement of those in the developing world.

Jefferson, in the letter I quoted, was referring to the institution of slavery, but he could just as well have been referring to our over-exploitation of the commons, an insult to justice that he could barely have imagined 200 years ago.

He and the other Founders made repeated reference to the principles captured in the ancient legal principle of "usufruct" – Jefferson, for instance, in a letter to James Madison two weeks before Congress passed the Bill of Rights, asserted that, "the earth belongs *in usufruct* to the living" (emphasis added). This is the principle that the Earth belongs to God, that we are entitled to the fruitful use of the Earth's bounty, but only to the extent that our exercise of that right does not "injure the substance" of that bounty, in so doing impinging unjustly on the rights of others to enjoy the same usufructuary rights. This principle, which was evoked by John Locke and Adam Smith among others, draws on a rich tradition going back through Aquinas, to Augustine, to the New Testament and to Rome, to Greece and to the Old Testament, and it lies at the very heart of the idea of democracy as envisioned by the Founders – that every person has a right to pursue his or her share of the fruits and benefits of

Earth's natural bounty – the unalienable right to "the pursuit of happiness". While the principle deals with access rather than outcomes, the question of equity of outcomes arises directly if the result of our "pursuit" and the that of preceding generations is *likely* to deny to a vast proportion of those alive today access to their fair share of ecosystem services – if, in Locke's words, as a result of our actions and those of our predecessors, there is not "enough, and as good left". In other words, the very same gift of God – nature and its laws – from which we divine our unalienable rights, imposes natural limitations on our exercise of those rights.

These Enlightenment principles – and the rich philosophical and theological tradition from which they derive – were quite explicitly meant to apply not just to the present commons but also to the intergenerational commons, the community both of the living and of those yet to be born. The Old Testament Book of Leviticus, for instance, sets out the law of the Jubilee Year, by which at the end of each forty-nine-year cycle what we would call nature's services (in short, land) were to be surrendered intact to their "original owner", to ensure that all subsequent generations had the same opportunity to enjoy the fruits of Nature's bounty. Adam Smith, the patron saint of capitalism, wrote, "The Earth and the fullness of it belongs to every generation, and the preceding one can have no right to bind it up from prosperity". Thomas Paine in *Common Sense* wrote, "In order to discover the line of our duty rightly, we should take our children in our hand…that eminence will present a prospect, which a few present fears and prejudices conceal from our sight." The Preamble to the Constitution declares that its purpose is to "secure the Blessings of Liberty to ourselves and our Posterity". Indeed the whole idea of American democracy rests on the Enlightenment doctrine that no generation has the right to infringe the natural rights of future generations – that natural justice demands that the living be indebted to the dead only for improvements to the inheritance granted by God. In the language of a more recent and more secular version of contractarian thought from John Rawls, one would not choose from the Original Position to belong to a society in which one might be born into a generation hopelessly impoverished by previous generations.

So these dual principles – that we have a self-evident right to pursue the fruits of God's bounty only within the limits imposed by the natural rights of *all men and women* to pursue those same benefits; and that those rights and obligations extend not only to all of those alive today but also to those yet to be born – are both profoundly Judeo-Christian and quintessentially American.

[SLIDE]

These slides provide a graphic depiction of the geographic distribution of CO_2 emissions from fossil fuel combustion since the Industrial Revolution,
[SLIDE]
along with the distribution of the wealth that's been created by doing so.
[SLIDE]
We in the developed world have had a pretty good run. At least since the 1980s (and frankly well before that), evidence has been accumulating that there is a natural limit to rate at which the biosphere can absorb anthropogenic greenhouse gases without overwhelming the natural carbon cycle, and that we are rapidly approaching that limit. If so, the developed world has consumed a disproportionately large share of the biosphere's natural tolerance for fossil fuel combustion and has profited handsomely from doing so. Finger-pointing is hardly necessary – it's a matter of simple economic interests to point out that those who have reaped, and

continue to reap a disproportionately large share of the bounty available from the use of a shared limited resource, have a disproportionately large interest and a disproportionately large capacity to develop and deploy alternatives that will maximize the value of the remaining resource.

Now, many people point out that the theory of anthropogenic climate change is unproven; that much uncertainty remains about exactly what the impact on our ecosystem will be from continuing with business as usual. This, they would say, renders all these fancy ideas about distributive and intergenerational justice moot. Assessing that claim requires that we move from ethics to the philosophy of science. Karl Popper, in the first half of the last century, revolutionized our understanding of the scientific process, and while his hypothetico-deductive model has undergone extensive revision since then, his basic insight remains unchallenged. All scientific knowledge is by its nature conditional knowledge – in essence, it consists of collections of hypotheses that have yet to be proven false. Let's be optimistic and say that science approaches truth asymptotically, though it's entirely possible that the nature of the scientific enterprise is far less certain – every new theory accumulates unexplained phenomena or untestable hypotheses as it progresses, like single socks piling up in a drawer, and we can never be completely sure that we're closer today to the truth about how the world works than we were yesterday. We proceed under the illusion of the certainty of scientific knowledge, when in fact we act constantly under conditions of uncertainty, even uncertainty about just how uncertain we actually are. To cut to the chase, the claim that we haven't "proven" the theory of anthropogenic climate change is true but uninteresting. Because of the momentous ethical and existential questions involved, the only interesting questions are how far we've reduced the uncertainty, and what the consequences might be of acting on the assumption that it's true, versus the consequences of acting on the assumption that it's false.

I'm not going to dwell on the science, first because I'm not a scientist, I'm just a simple engineer, but second because what I really want to talk about is the eminent feasibility and affordability of actually doing the right thing. But here's where the preponderance of the scientific evidence appears to put us.

[SLIDE]

Let's be clear - there's little debate about the severity of impacts to the ecosystem under various warming scenarios – the uncertainty at issue is which warming scenarios we'll find ourselves in under various emission scenarios. The best estimates we have are that adapting to a 2°C rise by 2100 would cost 20% or more of global GDP; with warming of 4°C or more by 2100, the likelihood of run-away climate disruption increases dramatically, with dire consequences for the fundamental conditions upon which modern human civilization is based.

[SLIDE]

Meinshausen et al. have produced an analysis of what our total budget might be for liberating naturally sequestered carbon between now and 2050, integrating the latest work on climate sensitivity with carbon-cycle feedbacks, and how much of that budget is left. According to this analysis, continuing with business as usual – a rise of about 1.5% per year – puts us here [*], where we have a mean chance of about 90% of exceeding a 2°C warming by 2100. Peaking global emissions in 2015 and reducing them in absolute terms from there – roughly equivalent to a global cut of 80% below 1990 levels by 2050 – puts us here [*], where we have a mean chance of about 30% of exceeding 2°C warming by 2100. Each emission scenario, of course has a range of possible outcomes.

[SLIDE]

Taking a different approach, according to the IPCC4 the probability distribution of climate sensitivity values – the expected warming from a doubling of pre-industrial concentrations to 550 ppm – looks like this, with likely warming by 2100 of about 3°C and a worryingly fat tail. But these curves are three or four years old now, and maybe even then the IPCC got it wrong. Where do the most recent empirical observations point to on this continuum? Nearly every key indicator – observed sea level rise, polar ice dynamics, ocean acidity levels – is trending beyond the worst-case scenarios in IPCC4 – in other words, things aren't going well. MIT's Joint Program on the Science and Policy of Global Change, previously a relatively cautious voice on this issue, in January revised their assessment dramatically based on new information, increasing the risk – without urgent policy action – of a rise of more than 7°C by 2100 to 9% and reducing the chance of staying within 3°C to 1%, with the risk of a rise of 5°C or more by 2100 at over 50%. But maybe the warming changes we're observing are just natural – maybe we had little to do with it and can do little to stop it.
[SLIDE]
Upper left is the current fit of the range of models to actual over the past 110 years – the dark black line is actual and the light grey line is the midpoint of the various model results. Below right is a graph of the same models versus actual, with human forcing mechanisms stripped out of the model runs.

Each of these curves, even business as usual, includes some probability that we don't go beyond a 2°C rise, and that's what the coal industry and the U.S. Chamber of Commerce would really like us to focus our attention on.
[SLIDE]
But from a moral perspective what's most interesting is the rest of the curve, and particularly the "fat tail", seen here in the probability distribution I showed previously. This is the distribution for a 550 ppm stabilization level.
[SLIDE]
This graph shows our current trajectory – reaching 866 ppm by 2100. The most aggressive versions of policy measures on the table around the world (including Waxman-Markey) would, if implemented, put only put us back on the 550 ppm trajectory. To have a 50% chance of staying within a 2°C warming, we need to be on this trajectory – technically achievable according to recent studies by McKinsey but requiring much more aggressive policy action.
[SLIDE]
Looking at the "fat tail" of the probability distribution, the "best recent proposals" trajectory – 550 ppm by 2100 – carries a greater than 25% chance of a rise of 4°C or more – a 25% chance that the biosphere is changed beyond all recognition – and not in a good way – within the lifetime of the Notre Dame Class of 2013. What if the actual probability of catastrophic disruption is even half that number? What about a fifth of that number? If you showed up for a flight and were told there was a 5% chance the plane would crash, would you board anyway? Probably not. If you were buying a ticket for your child and were told there was a 5% chance the plane would crash, would you buy the ticket anyway? I hardly need to ask the question.

In short, science may not have taken us as far as it can, but it has certainly taken us as far as it needs to. From a moral standpoint, this is no longer a scientific issue - it's an issue of justice.

So what would we need to do?

[SLIDE]

The interesting thing from the perspective of the energy sector, is that any of the scenarios I've talked about that bring the P90 outcome within a morally responsible range, involve at a very minimum something that you don't hear very much about from the fossil fuels industry, even the "good guys" in the fossil fuels industry – the power sector must be fully de-carbonized by 2050. Zero or effectively zero greenhouse gas emissions. What often passes for an enlightened view – incremental transition to a "green" energy system sometime toward the end of the century – is in reality an extremely high-risk strategy.

But what if, as some sceptics maintain, the cost to reduce CO2 emissions from the power sector to zero would be ruinous, would destroy the quality of life in the developed world and deny the emerging countries the right to develop their economies? Maybe we should spend our money dealing with more solvable problems and just assume we'll adapt to what might come. Viewed in proper perspective, at least in the case of the power industry, this is an outrageous misrepresentation, to put it mildly. Consider a few data points for comparison. We invested 4% of U.S. GDP [every year for ten years] in the Apollo Program – in today's money, that's [$520 billion] [$5 trillion]. We've been willing to spend something like $700 billion (and counting) on the second Iraq war ostensibly because Iraq *might* have had weapons of mass destruction that they *might* have given to someone else who *might* have used them against the United States. Are we at least *that* convinced that continuing our current rate of greenhouse gas emissions could result in profound, even catastrophic changes to the ecosystem that our children and grandchildren will inherit? You bet we are – not even the most sceptical climatologist would claim otherwise.

Let's look at another data point, one that's less divisive but perhaps more telling than citing the cost of invading Iraq. The U.S. consumer electronics market in 2008 was $172 billion. Most of that is spent upgrading from existing equipment. It's estimated that more than 75% of all the computers ever sold are stockpiled in basements, closets, garages and warehouses. The average cell phone in the U.S. is replaced every 18 months. In other words, Americans spend at least $130 billion a year (and climbing) replacing things that remain perfectly capable of serving the purpose for which they were originally purchased – because we want a better one, or a cooler one, or simply because we're bored. We do this with cell phones, computers, TVs, Kindles, X-Boxes…so why do we treat the plants that power all of these all-important electronics any differently? The answer, of course, is that we don't, or more precisely, we treat them the same way when it suits our purposes to do so. Four times in the past 30 years regulators and legislators have willingly stranded hundreds of billions of dollars in energy assets – (1) in the late 70s and early 80s state utility regulators across the country, faced with huge rate increases from overbuilding of nuclear capacity that ended up costing, on average, three times the original estimates, ripped up the regulatory compact they'd tacitly operated under for over fifty years and stranded hundreds of billions of investor assets, bankrupting several utilities and nearly bankrupting over a dozen more; (2) less than a decade later, FERC Order 88 liberalized the interstate pipeline system, again stranding hundreds of billions of dollars in investment in under-utilized pipeline capacity; in the mid-90s the wholesale electricity market was liberalized under the Energy Policy Act of 1992, again stranding tens of billions of dollars in under-utilized or inefficient capacity;
[SLIDE]
and finally this – between 1997 and 2005, 250 GW of new gas-fired capacity – the equivalent of over three complete UK power systems in eight years – was built in response to low gas prices and the expected retirement of much of the existing coal fleet under the New Source Review provision of the Clean Air Act. In the event, the Bush Administration refused to

enforce NSR and gas prices shot up, so it was the new gas plants that were stranded – some $150 billion worth – but that simply confirms my point. This time it would be comparatively easy – the average coal plant in the U.S. is over 40 years old and fully depreciated, so while many of these plants remain technically useful, the financial asset base is relatively modest. So the next time you hear some coal industry executive or one of their lackeys in the legislature say that we have to wait until the existing plants are retired to replace them, don't believe them.

Looking at electricity – which accounts for a quarter of the world's non-LUCF greenhouse gas emissions, and which in coming decades could fuel, directly or indirectly, most or all of the world's transportation needs – what would it take to replace our current system with a zero-carbon system within the next 40 years, one that can deliver the same quality of energy services we receive today? The fossil fuel industry says it's not possible, that we'd bankrupt our economy trying. According to them, we'll still need coal, oil and natural gas – nearly as much as we're using today – to keep the lights on and keep the economy growing. The real answer is that the technical means are at hand; the cost of commercializing them rapidly at very large scale are relatively modest, at least by the standards set by the examples I've given here; and with some smart decisions in the next few years we can preserve and indeed improve upon the quality and quantity of energy services delivered not only to ourselves, but to everyone on the planet.

A politically feasible approach to full de-carbonization starts with consumption. It is well settled that energy consumption in the existing building stock could be reduced by 30-50% - with a net positive impact on quality of life – at a cost of less than 3 cents per kWh, or half or less of the cost of building a new power plant. Implementing proven regulated business models that would drive utilities to invest in efficiency rather than supply is ridiculously simple – it's the ultimate no-regrets policy decision. Financing mechanisms can and should be funded from carbon revenues, thus greatly multiplying the effectiveness of carbon pricing programs in actually reducing carbon. And best of all, even under the most dire forecasts of the impact of climate action on unit prices of electricity, household electric bills would be dramatically lower than they are today.

Having reduced the size of the supply challenge, what about the challenge of transitioning to a zero-carbon supply? We're engaged in a major project with McKinsey, the Dutch Energy Research Centre ECN, and transmission consultants KEMA to answer that question.
[SLIDE]
Nuclear will play a role, but cost, lead time and lack of operational flexibility will limit nuclear, at the very most, to replacing today's plants when they are retired. And any scenario that relies on significant new nuclear construction activity must take into account the possibility – quite high in my view – that an incident like Three Mile Island or, God forbid, Chernobyl occurs again sometime between now and 2050, with predictable consequences for continued deployment. Fossil plants with carbon capture and storage will probably play a role, but the uncertainties surrounding cost and the availability of economically feasible storage sites, as well as lack of operational flexibility, will limit the scope of that as well. Most importantly, none of the capture technologies available today are able to capture more than 90% of CO_2, which makes any scenario in which a significant amount of coal-fired generation remains on the system in 2050 highly problematic whether fitted with capture or not. Natural gas can play a transitional role, but by 2050 it would need to be off the system entirely (with the possible exception of rarely-used emergency back-up facilities) or fitted with capture, and it's not clear that carbon capture will ever make sense for gas-fired

combined cycle plants. That leaves renewables, which I'll get to in a moment. Every one of these options, including large-scale renewables, poses significant public acceptance challenges at the scale required to de-carbonize the power sector. These challenges can be overcome by concerted public action, but toward which options should that public action be directed? The answer to that question changes dramatically when you move from a regime of aggressive but incremental increases in renewables production and incremental reductions in carbon intensity, to a commitment to fully de-carbonize the power sector in the next forty years or less. The difference is not so much a technology issue or a cost issue; it is rather a system issue. To illustrate, let's look at the country with the highest penetration of non-hydro renewable supply, Denmark.

Denmark is said to derive over 20% of its electricity from non-hydro renewables, primarily wind but also biomass-fired cogeneration plants.
[SLIDE]
This is a graphic of hour-by-hour demand on the West Denmark grid in January, 2003, when renewables production had reached nearly the level it's at today.
[SLIDE]
When we introduce the actual wind production during that month (a very good month for wind production, by the way), effectively a reduction in demand for other forms of generation, we get a picture of what is referred to as the "intermittency" of resources like wind – small run-of-river hydro, for instance, has similar attributes but usually represents a tiny fraction of system supply. The green represents demand that must be met by other forms of generation, in Denmark primarily by coal.

[SLIDE]
When we then introduce actual production from combined heat-and-power plants (which have little or no real operational flexibility), the effect looks like this. Without some way to buffer the impact, the operational demand for fossil plants to meet demand during low-wind hours would be highly erratic, far too erratic to be met by conventional coal- and gas-fired generation, and during many hours of the month when wind production is high there is a significant surplus production, which under normal circumstances would require the wind turbines to be curtailed at a far higher rate than owners could sustain financially without compensation. But the West Denmark system works quite well – how?
[SLIDE]
West Denmark is blessed with two highly unusual advantages – first, it is synchronized with the much larger German grid and is strongly connected via high-voltage direct current cables with the rest of Scandinavia, with a total export/import capacity that is nearly identical to its installed wind capacity; second, the Nordic system it's connected to is very hydro-intensive, in effect a storage battery about seven times the size of the West Denmark system (Norway's system is over 95% large hydro). Thus West Denmark can and does send on average approximately 85% of its monthly wind production to be stored behind dams in Norway or to Germany – where the system is large enough to absorb that amount of excess production – and brings it back as hydro and coal generation when the power is needed. The point is not that it doesn't work – it does – but rather in understanding how it works. This in turn can inform what might and might not work if we want to derive 80% of the supply for an entire continent from renewables and 100% from zero-carbon sources. Again, the problem in the first instance isn't technology or resources – wind turbines are already highly efficient at converting kinetic energy to electricity, and the technical potential of the wind resource in the North Sea alone is 30 times the EU's total electricity consumption – it's a system issue.

The answer lies in shifting our focus urgently to two issues that barely ever get mentioned in the "green energy" conversation – a greater emphasis on sources that can be ramped up or down on demand (in industry parlance, firm dispatchable capacity, sometimes mislabelled baseload generation, which contrary to popular wisdom has little intrinsic value), and deployment of information and control technologies on the transmission and distribution grids that would allow demand to be much more closely coupled to supply (sometimes called a "smart grid", though that term can mean a lot of things).

Let's talk first about firm dispatchable capacity. Current renewables support policies, with very few exceptions, ignore the problem of firm capacity altogether and focus solely on increasing raw energy production. Why is firm dispatchable capacity such a high priority? Coal plants produce 80% of the CO_2 emitted by the power sector, so if you're not displacing unabated coal plants, you're not really making progress. And displacing unabated coal plants means replacing firm, dispatchable capacity. Wind power is a wonderful thing, but it takes at least ten MW of installed terrestrial wind capacity to replace one MW of coal capacity – offshore wind is much better, about twice as good. Recall that Denmark synthetically creates firm dispatchable capacity from wind farms by combining them with a large quantity of high-voltage direct current transmission and with storage in the form of hydroelectric potential. Except for the wind turbines, most of those assets were pre-existing – to build that from scratch, at continental scale, is quite a different proposition. Solar PV is in some ways more predictable, but it is still highly intermittent. Two strategies are often suggested for dealing with these intermittency issues. The first is building enough new transmission capacity to move electricity instantaneously from where there's too much of it to where there's not enough – the theory that "the wind is always blowing and the sun is always shining somewhere", also called the "supergrid". The second is to postulate a technological breakthrough that will enable storage of electricity (or compressed air, or hydrogen) in extremely large quantities, for extended periods of time, with very rapid charge and discharge rates. The first strategy – the "supergrid" – has much to recommend it, but the smoothing effect of interconnecting intermittent resources over very large areas is good but far from perfect, and the cost and public acceptance issues involved in constructing enough new high-voltage transmission to actually accomplish that are truly awe-inspiring. The second strategy – mass storage – is currently economically feasible only with pumped storage hydro, and the potential alternatives are, to my knowledge, decades away from commercial application. If you want to get really rich and save the world at the same time, solve the storage problem. But as we sit here today, the options for managing intermittency are imperfect at best, and without a major technological breakthrough on electrical storage, intermittent sources will play an important but limited role in a zero-carbon power supply.

Nuclear and coal-with-carbon-capture both suffer from a similar problem. While they're not intermittent, they provide little or no dispatchability because of their lack of operational flexibility. A reliable power system doesn't actually need inflexible baseload generation, and too much of it is as big a problem as too much intermittent supply. A zero-carbon system made up entirely of intermittent renewables, nuclear and coal-with-CCS would possess neither the flexibility nor the dependability required to reliably match current demand patterns.

So we need to look for resources that are plentiful, flexible, firm and dispatchable.
[SLIDE]
As the current push for intermittent renewables like wind and other more limited low-carbon opportunities like municipal solid waste incineration begins to reach it's maximum

penetration, a different set of resources will need to start kicking in in a big way post-2020. It will take that long for any new nuclear capacity to come online, for carbon capture and storage to be commercially demonstrated, and for renewable resources with higher firm capacity value to reach commercial scale. In the near term, swing capacity will come mostly from gas-fired generation, but gas-fired plants are far from zero-carbon. It is not at all certain that carbon capture will make economic sense for any gas-fired generation, and it simply does not work at all with the flexible simple-cycle gas turbines that will best fit the needs of a system with significant amounts of intermittent supply.
[SLIDE]
Fortunately, there are a number of renewable options, and three of them – concentrating solar thermal, geothermal and biomass – have the potential to be large-scale solutions well before 2050. Biomass-fired generation is a fully commercial technology today, and biomass fuel-processing technologies like torrefication will significantly expand the accessible feedstock supply and improve the economics. Conventional geothermal is a fully commercial technology with limited resource availability, but enhanced geothermal systems – in which water is injected into deep wells to fracture hot dry rock formations and then extracted as steam to drive a turbine – are a promising pilot-stage technology with vast potential. The U.S. Geological Survey released a massive study in 2008 that put the median U.S. potential for new conventional geothermal at over 30 GW and the median potential for enhanced geothermal at 518 GW, compared to about 800 GW of currently installed generation. Concentrating solar thermal power, or CSP, has effectively unlimited potential. Parabolic trough is proven technology, and central tower, compact linear Fresnel reflector, and Sterling engine-based technologies are rapidly moving out of the demonstration phase into commercial operation. Molten salt storage systems, which convert concentrating solar thermal power plants into firm, fully dispatchable resources, are under demonstration right now with up to fifteen hours of economic storage capacity. A hypothetical concentrating solar field 18 miles long and 18 miles wide in the Nevada desert – about 1% of federal lands in Nevada – could replace the entire U.S. coal fleet. That's far less than the acreage destroyed so far by mountaintop removal coal mining in Appalachia, to put it more starkly in perspective.

All three of these options are capable of providing large amounts of firm, dispatchable capacity.
[SLIDE]
This is a graph comparing the output on a cloudy day from a PV plant and a solar thermal plant about 30km away from each other in Spain. The difference in their ability to match the daily load profile – their relative dispatchability – is quite dramatic, and that's before adding thermal storage to the CSP plant.
[SLIDE]
These technologies are still more expensive than the high-carbon alternatives, but they're not that expensive, and with only 800 MW of capacity in commercial operation today, concentrating solar thermal in particular has enormous potential for improvements in cost and performance over the next ten years.
[SLIDE]
A study released by the Center for Global Development in Washington last December concluded quite credibly that with a total incremental investment of approximately $40 billion in government support, concentrating solar thermal power with thermal storage could reach full commercial competitiveness with coal (including a CO_2 price of [$25]/tonne) by 2020, the pre-condition for mass deployment of the technology in both developed and emerging economies. Their analysis is consistent with other recent assessments of the

technology. The high-voltage direct current transmission necessary to bring bulk power long distances to major demand centers has been in commercial use since the 1950s, and recent innovations in high-voltage switching technology have made networked HVDC applications possible. Perhaps most importantly, most existing generation will have to be replaced between now and 2030 in any case, so the illusion of cheap electricity we currently enjoy from fully depreciated old plants will come to an end. None of this comes cheap – replacing the current coal fleet entirely with concentrating solar thermal plants with thermal storage, and building the associated transmission, might cost $500 billion more over the next twenty years than replacing them with new coal plants, but once they're built they'll never have to pay a fuel bill, buy carbon credits, or force another mountaintop to be blown off and dumped into rivers in order to access "cheap" coal. The fact is that the difference between the cost of the options proposed here and the cost of business-as-usual is nowhere near as large as the fossil fuels industry would have us believe, and there may be no incremental cost at all on a lifecycle basis.

So firm, dispatchable renewable generation is technically feasible, with effectively limitless potential, and can be made economically competitive within ten years. But it will still be necessary, in planning for a zero-carbon power supply by 2050, to assume that such a system will have substantially more uncontrolled variability in supply than is the case today, with expanded but limited storage options. This brings us to the second dimension urgently requiring increased focus, which I will only touch on here. The master-slave relationship between demand and supply will have to change substantially in the coming two decades. Demand will have to become much more closely coupled with supply, such that uncontrollable fluctuations in supply, up or down but most importantly down, are instantaneously matched by controlled changes in demand. This will involve the tight integration not only of locally installed production, such as PV panels on rooftops and local biomass-fired combined heat and power plants, but also real-time integration of all major consumers of electricity in homes and businesses. Appliances, electronics and lighting will respond to system conditions in real time in ways that do not noticeably affect their delivery of services. The technology to do this largely exists today, but the regulated business models for the delivery of energy services are stuck in the 1990s, or worse. This is another no-regrets policy change that can and should be implemented.

So to bring it all together:
First, all cost-effective efficiency measures and aggressive deployment of current intermittent renewables technologies in the best resource areas, including in particular high-quality offshore wind.
Second, some deployment of CCS in favourable locations and replacement of at least some of the retiring nuclear production. Third, accelerated commercialization of firm, dispatchable renewable technologies like concentrating solar thermal with thermal storage.
Finally, accelerated deployment and integration of distributed generation and smart grid technology.
A pre-requisite for all of these measures, and one that needs to happen within the next few years, is a radical restructuring of regulated transmission and distribution business models and siting procedures to enable investment in the new infrastructure that will be required.
What might it look like when we're done? The best part of all is that it looks pretty darn cool.
[SLIDE]
We keep hearing from the coal industry, the US Chamber of Commerce and others, that we can't do this, that a zero-carbon electricity system will be too hard or technologically

impossible. How about a little "shock and awe" for the defenders of the status quo? How about some good old-fashioned American "can-do" attitude in response to the prophets of incrementalism? We've done much harder things for prosperity – or at least, our forebears did. Now it's our turn. Thank you.