Hurry or Wait?

Pacing the rollout of renewable energy in the face of climate change risk

November 11, 2016

THE Brattle GROUP

Jürgen Weiss

Dean Murphy

Jurgen Weiss is a Principal at The Brattle Group and a Senior Nonresident Fellow at Boston University's Questrom School of Business. Dean Murphy is a Principal at The Brattle Group. The ideas expressed in this paper do not represent positions by The Brattle Group. Any remaining errors are the authors'.

Abstract

Climate change risk will likely require the decarbonization of our economy, in large part through investment in renewable generation. Since technological progress will likely lower the future cost of renewable generation capacity, however, delaying decarbonization could save money by benefitting from lower installation costs later. On the other hand, waiting to decarbonize would also impose costs: it would allow long-lived greenhouse gases to accumulate in the atmosphere, increasing long-term climate impacts. Reducing emissions sooner would cut the cumulative concentration of CO2 and reduce expected climate impacts. In addition, early reductions would provide insurance against the possibility of greater climate sensitivity, positive feedback loops, or points of no return and the resulting risks of catastrophic climate outcomes. In this paper, we develop a simple model of investment and technology cost evolution to shed light on the financial costs and carbon reduction benefits of two stylized strategies for deploying renewable generation to transform the electricity sector: “Hurry” vs. “Wait.” Our fundamental question is whether it may be worth accelerating the replacement of existing fossil generators with renewables, when weighing the additional financial costs against the carbon benefits. Comparing the cost of accelerated renewable deployment and associated carbon savings with the lower cost but higher emissions approach of delayed deployment, we find that the incremental costs of Hurrying, expressed in terms of the cost per ton of avoided carbon emissions, is below most estimates of the social cost of carbon, even without considering the insurance value of earlier reductions. This result may be somewhat surprising because at present, unsubsidized renewable costs in many parts of the United States are significantly above the cost of existing fossil generation. Our result occurs because learning effects reduce future renewable costs, and costs decline faster still if renewable deployment is accelerated on a large scale. As renewable additions scale up from their initial small base, costs fall and the majority of additional renewables will be installed later at lower cost. This means the carbon savings accumulate faster than the incremental costs.

Keywords: climate change; electricity sector; environmental risks; fat tails; greenhouse gas emissions.

I. Introduction

The risks associated with climate change are likely to require the ultimate de-carbonization of our electricity sector and the broader economy, involving very large investments in long-lived assets using a variety of relatively new and emerging technologies, primarily renewable generation technologies.[1] Given observed rates of technological progress, it is likely that the cost of these technologies will continue to decline and that their performance will continue to improve for the foreseeable future. Thus, investing early and rapidly, say by installing large amounts of solar photovoltaic (PV) generation today, may foreclose the ability to install lower-cost PV later.[2] Although some of the potential cost reductions may result from learning and scaling effects that depend on earlier deployment and thus may not be captured by waiting, at least some cost improvements would likely be achieved even at low deployment rates. These might be due to "natural" improvements that come from basic technological advances, and/or by learning from other countries' deployments.[3]

But in the context of climate change, arguments for delaying deployment to exploit future cost reductions must be balanced against the potential costs of delay. Apart from the possibility that delayed deployment might also delay the hoped-for cost declines (by deferring learning-by-doing and scaling benefits), there are other costs to waiting. Since CO_2 remains in the atmosphere for many decades, waiting to reduce emissions increases the cumulative concentration and thus the resulting expected climate impacts. Also, the relationship between GHG concentration and climate impact is complex. The distributions themselves are unknown, and quite possibly have "fat tails" – i.e., may have higher probabilities of extreme outcomes, due to positive feedback loops, etc. There may also be complex timing issues, such as tipping points beyond which climate impacts and costs of adapting could become catastrophically large. Deploying carbon reduction measures early and quickly could therefore also have a second important benefit in the form of providing insurance against the unknown but significant risk of catastrophic damage.[4] Our analysis suggests that, especially given the latter uncertainties, Hurrying is likely preferable to Waiting. The additional costs imposed by Hurrying are relatively modest, and in fact are generally below estimates of the social cost of carbon, even without consideration of "fat tail" risks or catastrophic damage.[5]

In this article we discuss, in a simple numerical model of the United States electricity system, the potential magnitude of these benefits and costs of accelerated vs delayed deployment of renewable power generation.[6] Our basic question is whether it is worth accelerating the replacement of existing fossil generators with renewable sources, when weighing the additional costs that would be incurred against the benefit of lower cumulative CO_2 emissions. We also analyze key sensitivities on our assumptions before putting the costs of Hurrying into a broader historic context and providing some concluding thoughts.

II. Characterizing Accelerated Renewable Deployment

To estimate the potential magnitude of both benefits and costs associated with accelerating renewable energy deployment, we developed a simple model of the United States electricity system that distinguishes several broad classes of generation.[7] Fossil generation is characterized as Gas or Coal, and renewables as Wind or Solar PV.[8] We use this model to characterize two alternative renewable deployment paths—a "Hurry" path that characterizes an immediate, aggressive schedule of accelerated deployment of renewable energy sources, and an alternative "Wait" path, in which renewable deployment is delayed (i.e., continues on a path comparable to recent historical rates).

Under each path, we estimate carbon emissions and total system cost, accounting for the evolution of renewable costs and fossil fuel prices. Importantly, renewable costs include the full cost of construction and operation, but fossil costs include only the going forward cost of existing fossil generators, ignoring sunk capital costs. We then compare the present value of aggregate costs and the carbon emissions of Hurry with those of Wait, to yield our Base Case estimate of what it would cost to reduce cumulative carbon emissions by accelerating renewable deployment and retiring existing fossil generation sooner. We also perform a number of sensitivity analyses on the model inputs—renewable costs, learning rate and cost decline rate, fossil fuel prices, discount rates, etc.—to explore how the cost of reducing emissions might change from this Base Case estimate.

Starting with actual 2015 generation from all sources, Figure 1 illustrates the electricity generation mix over time on the Hurry path and the Wait path. Both paths start in 2015 with the current U.S. generation mix, classified broadly as Coal, Gas, Wind, PV, or Other (primarily nuclear and hydro). Under each of the paths, fossil generation is retired as new renewables are deployed so that total generation remains equal to load (assumed to grow at 1% annually). Both paths assume no further increase in coal-fired generation, but rather that any demand growth not met by increasing renewable generation could be provided by increasing the output from existing natural gas-fired capacity. Until 2030, the "Wait" path in the left panel of Figure 1 assumes modest annual growth rates of 2.5% for wind and 10% for solar PV, representing "natural" growth, perhaps deliberately delayed to take advantage of future cost declines. Only after 2030 does the Wait path accelerate renewable deployment (from a then-higher base), so that it takes until 2040 for renewables to replace all fossil generation. Under the Hurry path, illustrated in the right panel of Figure 1, deployment rates are substantially higher starting immediately in 2016. The installed capacity of wind and PV both grow quickly, so that by 2030 all fossil power generation is replaced by equal parts wind and solar PV. On either path, once accelerated deployment is assumed to begin (in 2030 under Wait, and 2016 under Hurry), wind generation grows by 30% per year and solar PV generation grows by 40% per year, continuing at this accelerated rate until all fossil generation has been replaced by an equal mix of the two.

Figure 1: Production Mix Assumptions Under Wait and Hurry Renewable Deployment Paths

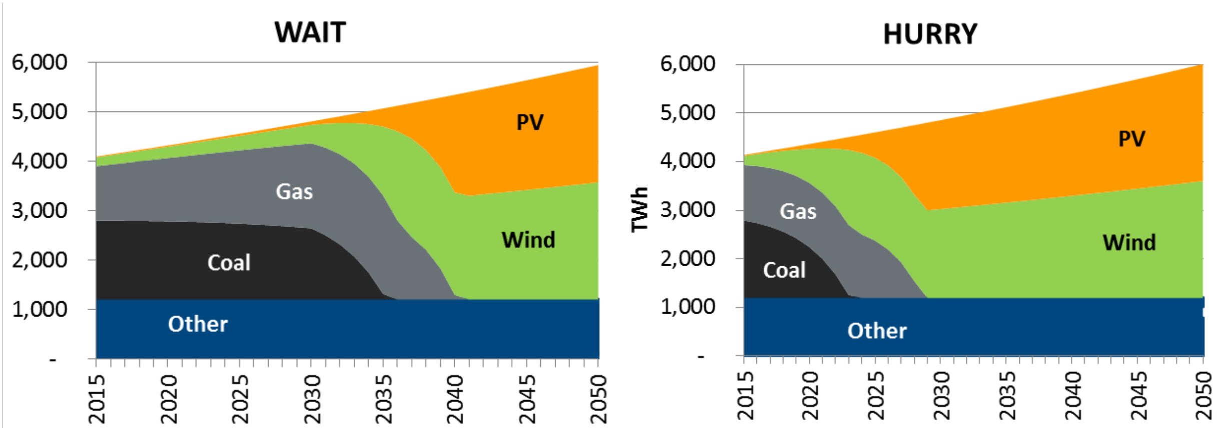

These renewable growth rates are slightly above the global averages observed since 2000, as illustrated in Figure 2.[9]

Figure 2: Historic Trend in Global Installations of Wind and Solar PV

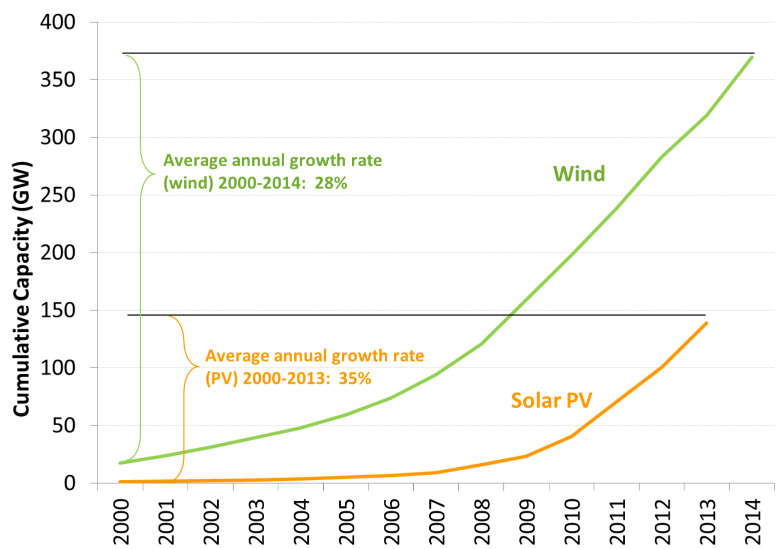

Sources: EPIA, Global Market Outlook for Photovoltaics 2014-2018, Figure 1; GWEC, Global Status of Wind Power in 2014.

To calculate the costs of the Hurry and Wait paths, we developed estimated costs for each generation type. As mentioned, fossil cost estimates include only going-forward costs (i.e., such avoidable costs as fixed and variable operating costs and fuel that are required to maintain and operate existing units, excluding the recovery of sunk capital costs); this assumes that existing fossil capacity would be able to provide as much power as is necessary.[10] The going forward costs of gas and coal generation are based on EIA estimates, using an average heat rate for gas and coal plants with EIA's projected gas and coal prices. The coal and gas price projections used are shown in Figure 3. Note that gas prices in the EIA Reference case, which we use in our Base Case projection, are assumed to rise materially over the next 20 years, reaching $6/MMBtu (real 2015 dollars) in the early 2030s. To test the impact of this assumption, we also examined as sensitivities, described below, the EIA's alternative low gas price case, the High Oil and Gas Resource case, in which gas prices rise to only about $4/MMBtu by 2035, as well as other even lower gas price assumptions.

Figure 3: Coal and Gas price projections, EIA Annual Energy Outlook 2015

Gas Price

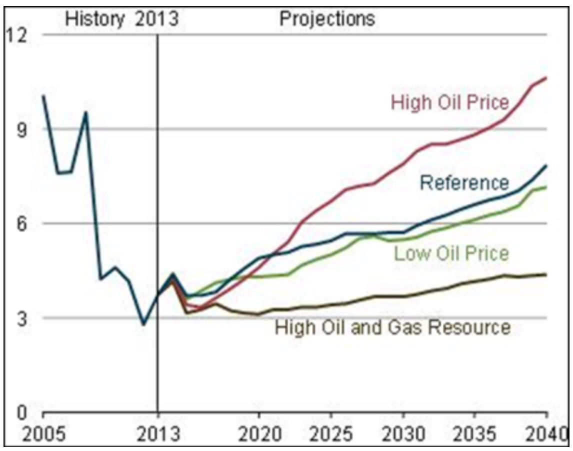

Coal Price

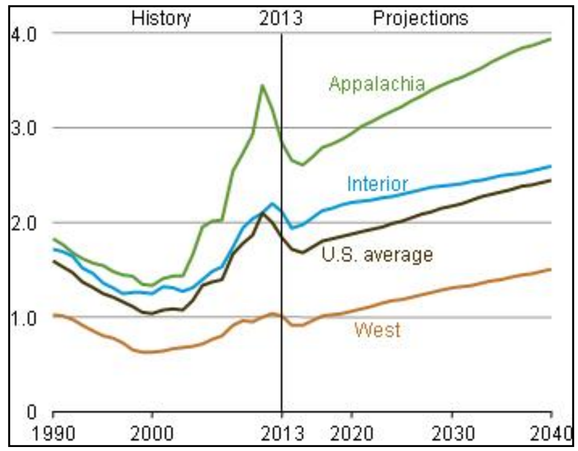

Our cost estimates for new renewables include the full initial capital cost, levelized over time, as well as operating costs. We used a recent report estimating unsubsidized costs of wind and PV as our 2016 Base Case cost assumptions.[11] Averaging across multiple types of solar PV installations, we assume 2016 costs of $143/MWh for solar PV and $87/MWh for wind.[12] By using only going forward costs for fossil resources, but full costs including capital costs for new renewable generation, we reflect the economics of our fundamental question: is it worth accelerating the replacement of existing fossil generators with new renewable generation, weighing the additional costs that would be incurred against the benefit of lower cumulative CO_2 emissions?

While difficult to forecast, future cost trajectories for renewable energy technologies such as wind and solar PV can be and often are estimated using learning rates, which typically characterize how costs fall with experience ("learning by doing," or LBD). Learning rates are often estimated using a one-factor model in which costs fall by a specified fraction with each doubling of cumulatively installed capacity. Thus, it becomes more difficult to achieve further cost reductions as the installed base increases, since each successive doubling requires the installed base to increase by twice as much. The learning rate may differ by technology and is typically estimated from the historic relationship between cost and deployment for a technology.[13] Estimates of learning rates vary widely for both wind and solar PV, depending on the time period and method used to estimate the rate, but a recent review of learning rates found that onshore wind averages a 12% cost reduction per doubling of installed capacity, and solar PV averages 23%, using one-factor models.

An alternative to a one-factor model is a two-factor model, which assumes that learning is driven only partly by cumulative deployment/LBD. In a two-factor model, cost improvements are also influenced by factors independent of deployment. Conceptually, these other factors are often characterized as cumulative research investment or knowledge stock and can be driven by targeted research, or simply by advances in related fields over time (e.g., improvements in semiconductor manufacturing may reduce PV costs). In the case of a domestic industry as we consider here, vicarious experience gained from international deployment, but not captured by domestic deployment, might also be considered as one of these factors. A one-factor model that relates cost reductions only to cumulative deployment (LBD) tends to magnify the benefits of a Hurry path with accelerated deployment – doubling the deployment rate doubles the cost reductions, and the only way to get cost reductions is through more deployment. In contrast, a 2-factor model that attributes only part of the cost decline to deployment dampens the benefits of a Hurry path – some cost reductions occur just by waiting, which will cause a Wait path to perform relatively better. We use a two-factor model that combines the effect of experience gained from cumulative U.S. deployment, and a separate factor based on elapsed time, independent of deployment, which can be thought of as a combination of research, advances in related fields, and international experience not captured by cumulative domestic deployment.

Our two-factor model assumes a constant annual cost decline of 1.5% for both wind and solar PV, attributing a portion of observed cost declines to accumulated R&D spending and technological advances rather than to learning by doing. To achieve roughly the same overall cost decline as has been observed historically, the learning rate for cumulative deployment is reduced to 7% for wind and 12% for solar PV (vs. the observed 12% and 23% overall rates for one-factor models).[14] Thus renewable costs decline with time and decline more rapidly with faster deployment, and cost decreases occur earlier on the Hurry path, shown in Figure 4. But ultimately, costs reach approximately the same point on both paths. The cost of wind power is projected to fall to about \$40/MWh by 2050 – about half its current level. Solar PV costs are projected to fall to roughly the same level, or even slightly lower, despite starting considerably higher, because solar PV has a higher learning rate, and also starts from a smaller base which implies that there is ultimately more opportunity for learning. Of course, the <u>average</u> cost of the full renewable portfolio at any point in time is above the cost of <u>new</u> renewables, since the portfolio includes older vintages with higher costs.

Figure 4: Cost of New Renewables under Hurry and Wait Renewable Deployment Paths

(Cost of nev … nerator life)

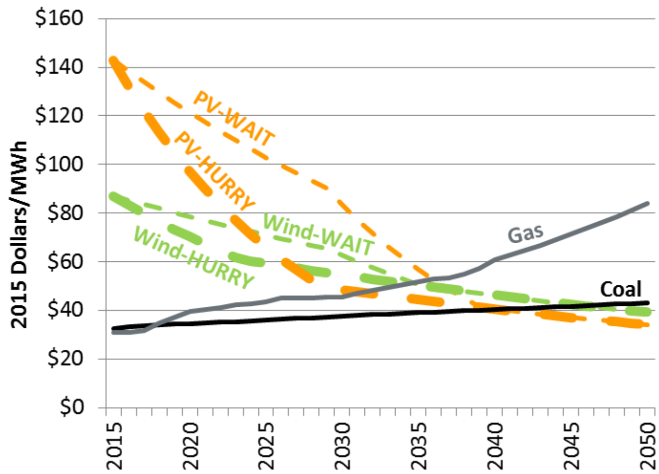

As we will discuss below, this approach also allows us to test the sensitivity of our results to different assumptions about the relative importance of the experience-dependent and the time/technology-dependent components of cost declines.

These assumed cost reductions may seem ambitious. However, it should be noted that our assumed final costs approaching $40/MWh (unsubsidized) for both wind and solar PV by somewhere after 2040 represent costs already achieved today in certain locations. For example, recent auctions for renewable projects in Chile resulted in otherwise unsubsidized power purchase agreement prices of $29.10/MWh for solar PV and $38.10/MWh for wind.[15] These low prices are likely the result of a combination of scale of installation (especially for solar PV), favorable renewable resources (a lot of sun and wind), low capital/financing costs and relatively low local costs for land and labor. Our cost assumptions are averages across all installations in the United States, which on average tends to have higher labor and land cost, worse renewable resources and a mix of project scales. However, future cost reductions could be achieved in part without any further technological progress by developing renewable energy resources disproportionately in regions with high quality resources, which are plentiful in the United States, and/or at larger scale. For example, the Lazard study that we used for cost assumptions for all technologies estimates the current unsubsidized levelized cost of wind in Texas to be $36-$51/MWh, i.e. already near or below our long-run cost assumption. Estimates for the Midwest are lower, at $32-$51/MWh. Similarly, estimates for the current unsubsidized cost of solar PV are as low as $57/MWh in Texas and $53/MWh in the Southwestern United States.[16] In sum, while our assumptions about the learning and time induced cost reductions for wind and solar PV in the United States may appear aggressive, they should be realistically achievable, based on already observed current costs and opportunities for further cost reductions through technological progress, scaling and targeting resource-rich locations.

III. Results: Financial Costs and Carbon Benefits

Using the assumptions described above, we simulated the renewable build-out under the Hurry and Wait paths. We call this our Base Case analysis, to distinguish it from several alternative sensitivity analyses we discuss below. Figure 5 shows the costs incurred, by technology, for the Base Case Hurry and Wait paths, and Figure 6 compares their total annual costs. The annual costs of the Hurry path are initially above those of the Wait path, and the difference grows as the amount of installed renewables on the Hurry path increases, displacing greater amounts of relatively lower-cost fossil. By 2030, the Hurry path eliminates fossil and its rate of cost increase slows to track load growth. After about 2035, costs on the Hurry path level off and then begin to fall slightly as the first generation of relatively more expensive renewables begins to reach the end of its life and is replaced with considerably cheaper second-generation renewables, outweighing the effect of load growth on costs. The costs of the Wait path rise more slowly in the early years, driven by load growth and rising fuel costs, with a modest contribution from the Wait path's slower renewable deployment. Beginning in 2030 when the Wait path's renewable deployment accelerates, its costs increase more sharply, exacerbated by gas prices that continue to rise (the Wait path still implies significant gas use at this point). Starting around 2034, the cost growth rate on the Wait path diminishes, as the falling capital cost of new wind and rising gas prices make all-in wind costs comparable to the dispatch cost of gas generators (wind makes up the bulk of renewables being deployed during these years). After about 2045, the annual cost on the Wait path, still rising slowly with load growth and not yet far into the second generation of cheaper renewables, rises above the cost of the Hurry path, which by then is falling slowly.

Figure 5: Annual Generation Costs by Technology under Hurry and Wait Paths

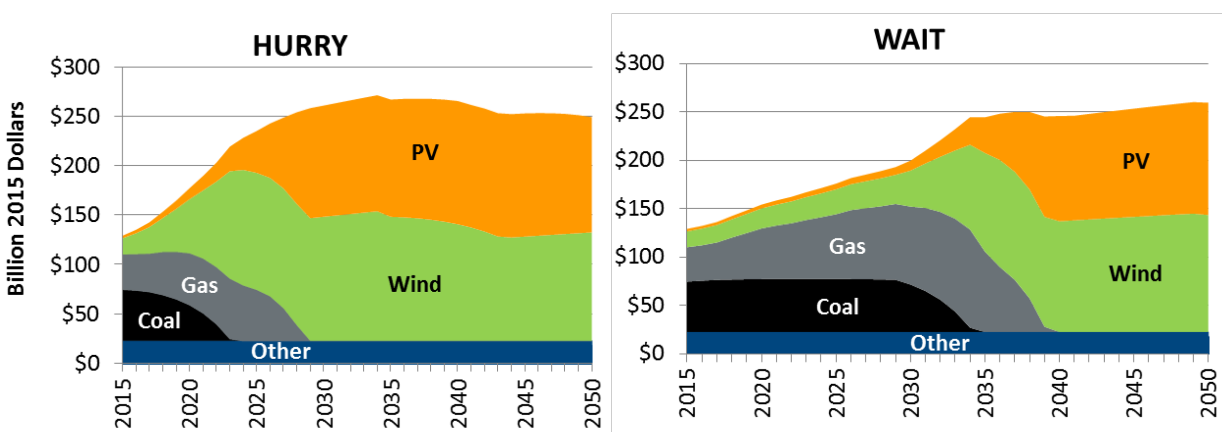

Figure 6: Total Annual Generation Costs under Hurry and Wait Paths

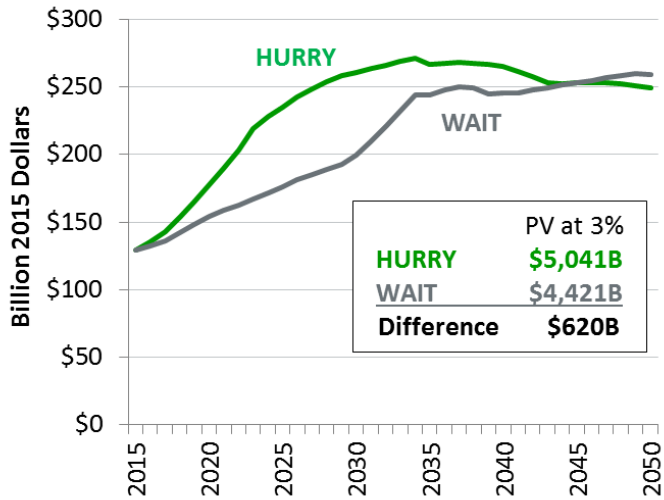

Despite Hurry costs being above Wait costs for most of the horizon, the present value of costs on the Hurry path is only about 14% higher than the costs on the Wait path. In any given year, the difference in electricity costs between the paths would be between 0.1 cents/kWh and 1.4 cents/kWh, averaging approximately 0.4 cents/kWh. Even ignoring the value of lower cumulative greenhouse gas emissions under the Hurry path, these increases are modest relative to observed historical rate fluctuations. For example, Figure 7 shows residential retail rates in Texas and Massachusetts between 2000 and 2014; the results for other states are similar. Rates have fluctuated by about 5 to 7 cents/kWh, driven primarily by underlying fluctuations in natural gas prices. The incremental costs we estimate for an accelerated renewable buildout are nearly an order of magnitude smaller than these recent historical rate fluctuations.[17]

Figure 7: Evolution of monthly average residential rates for electricity

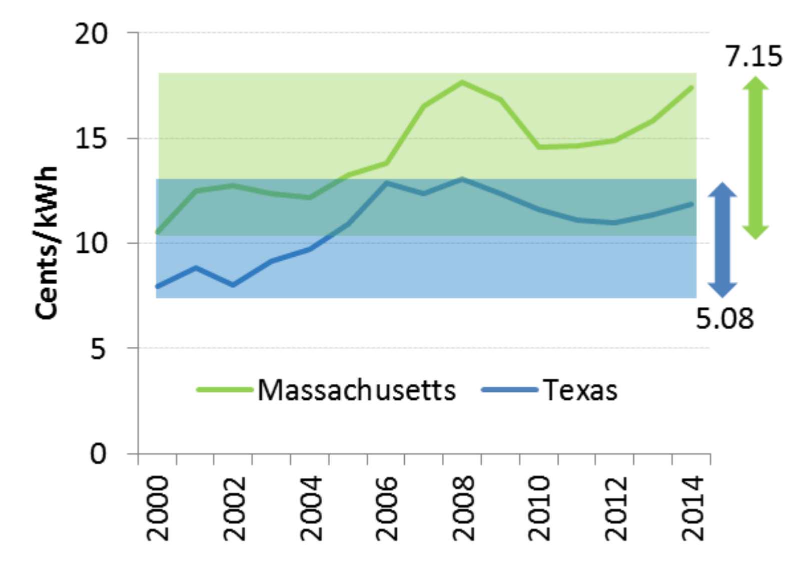

Source: U.S. Energy Information Administration

The benefit achieved by the higher electricity cost on the Hurry path, of course, is that it reduces cumulative CO_2 emissions considerably – from 45 billion tons over the time horizon in Wait, to 18 billion tons in Hurry, a reduction of 27 billion tons, or 60%, as shown in Figure 8. This is about three-quarters of an entire year's worth of current global carbon emissions (global CO_2 emissions in 2014 were estimated at 35.7 billion tons).[18]

Figure 8: Cumulative CO2 Emissions under Hurry and Wait Paths

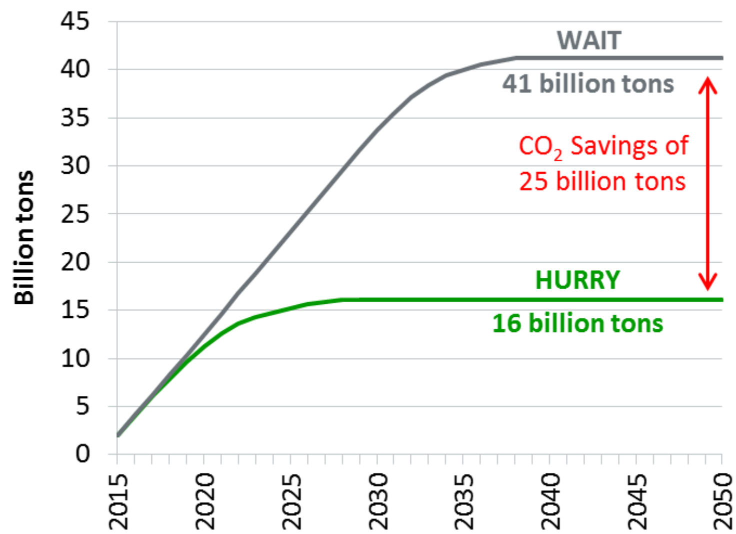

IV. Discussion

A key metric for evaluating whether the Hurry path is preferred to Wait is how the incremental cost of accelerating renewable deployment compares with the cumulative CO2 emissions it avoids – we refer to this as the Avoided Cost of Carbon (ACC) for the Hurry path. With the Base Case assumptions described here, the implied ACC is about $25 per ton; this is the present value of the difference in costs ($620 billion as shown in Figure 6 above), divided by the difference in CO2 emissions (25 billion tons, from Figure 8), between the Hurry and Wait paths.

One of the first things to note about this ACC value is that it is considerably below what current intuition might suggest about the cost of accelerated deployment of renewables. Absent subsidies, renewables currently cost on the order of $100/MWh (more for solar PV, less for wind), whereas fossil generation going forward costs are around $30 for both coal and gas – implying an additional $70 or so for each renewable MWh that is substituted for fossil. As a rough rule of thumb, each MWh of fossil power emits about 0.7 tons of CO2 (an equal mix of coal at 1 ton/MWh, and gas at 40% of that). This suggests an avoided cost of carbon of about $100/ton CO2 (= [$70/MWh]/[0.70tons/MWh]). It is important to explain why our $25/ton ACC result is so different from this tempting but ultimately misleading intuition.

There are several forces at work here. Discounting the incremental costs of Hurrying to present value has some effect, though it is modest because of the relatively low 3% discount rate we used. A somewhat bigger contributor is the fact that fossil fuel costs are expected to increase over time, which reduces the cost gap with renewables. But the primary factor is that current renewable costs are not representative of what renewables will cost on an accelerated build-out path that dramatically increases the amount of installed renewables, because renewable costs are expected fall with time and with experience.

There is indeed a considerable amount of additional experience to be gained, which will likely drive significant additional cost reductions. Wind, and especially solar PV, currently provides only a small share of U.S. generation – about 5% and 0.5% respectively. This implies that if costs follow learning rates similar to those observed historically, there will be significant reductions in costs through both of the mechanisms we consider – learning by doing, and even time-based cost reductions – prior to reaching the point where renewables are being installed in large quantities. That is, renewable penetration is so low initially, especially for solar PV, that there will be several doublings of installed capacity, and concomitant cost reductions, prior to installing the large majority of renewable capacity needed to decarbonize the U.S. electricity sector.[19] The parameters used in our Base Case imply that by the time renewable deployment is complete on the Hurry path in 2030, the cost of new renewables will have fallen to just over $50/MWh, and even the weighted average of all installed renewable capacity, i.e., accounting for the higher cost of renewables installed earlier, will be reduced to about $65/MWh. This makes the gap to fossil generation costs much smaller than at present, significantly reducing the avoided cost of carbon. And the cost of new renewables will fall further beyond that point (albeit slowly), while natural gas costs are projected to rise.

As a result, at some point in the future the all-in cost of new renewable generation is expected to match and subsequently fall below the going forward cost of gas-fired generation. On the Wait path, the later deployment of renewables means that there is still significant gas-fired generation operating at and beyond this break-even point, which increases Wait costs and reduces the ACC. On the Hurry path, in contrast, the earlier renewable build-out means that there is no gas-fired generation remaining by the time renewables and gas-fired generation are at cost parity, even though parity is reached earlier due to faster experience-induced cost reductions for renewables.

To evaluate the policy question of whether it is worth pursuing the Hurry path rather than Wait, we compare the economic cost of avoiding CO2 emissions, with the non-economic costs that would be incurred by not preventing those emissions. This latter cost is the "social cost of carbon," or SCC; it measures the economic damages associated with an incremental increase in CO2 emissions, and is often used to assess public policy measures aimed at reducing carbon emissions. Our ACC of $25/ton reflects emissions avoided primarily in the period from the early 2020s to the early 2030s, expressed as the present value (at 3%) to 2015. The SCC is estimated by the U.S. government, as well as other governments and entities, and is typically calculated and expressed as a present value to the time of emissions.[20] Several different SCC estimates are developed and used by the U.S. government; the central value, which is based on a 3% discount rate in 2007 dollars, is $42/ton in 2020 and $50/ton in 2030, continuing to rise thereafter. (The SCC is also reported based on discount rates of 2.5% and 5%. A 95th percentile value, discounted at 3%, is also reported as a metric to reflect the potential risks associated with climate change; this value is $123/ton in 2020 and $152/ton in 2030.) Importantly, these values are not expressed as the present value to 2015, as is our ACC. To make these SCC and ACC measures directly comparable, we re-state the SCC as the present value to 2015, for emissions that occur in the early 2020s to the early 2030s. Re-stated in this way, the central value of the U.S. SCC is about $37/ton, and the 95th percentile value is $113/ton. So our primary result is that the $25/ton ACC we find for Hurry (relative to Wait) is below the $37/ton central value of the SCC, and is far below the $113/ton 95th percentile value. This implies that the additional economic cost of Hurry would be more than justified by the expected benefit of the avoided emissions and social cost, and is quite small relative to this 95th percentile measure of climate change risks.

V. Sensitivity Analyses

In order to better understand how to interpret these results in a policy context, we performed several sensitivity analyses designed to test the robustness and sensitivity of our main conclusion. Figure 9 below summarizes the sensitivity analyses we performed. In particular, we tested changes to natural gas prices, the renewable learning parameters (which affect the magnitude of cost declines and how they depend on learning effects vs time trend), the discount rate, the start year of the Hurry versus Wait paths and the time delay between them, and the degree of decarbonization achieved. We also looked at pessimistic cases that combine a negative outlook on both of the most important parameters – gas prices and renewable learning parameters.

Changing our input assumptions generally affects the implied ACC in the expected direction. The magnitude of the sensitivity effects is perhaps the most interesting and surprising result from this analysis. Below, we discuss the effect of sensitivity analyses on some of the key inputs, and in the following section, we discuss their implications.

Figure 9: Sensitivity Analysis Results

		Time Trend		Learning Rate						
Scenario	Discount Rate	Wind (%/yr)	Solar (%/yr)	Wind (%/dbl)	Solar (%/dbl)	Gas Price	Decarb. Level (%)	Avoided CO2 (B tons)	Incremental Cost ($B, NPV)	**Avoided CO2 Cost ($/ton)**
Base Case	3.0%	1.5%	1.5%	7.0%	12.0%	EIA Ref.	100%	25.2	$ 620	**$ 24.63**
EIA Low Gas	3.0%	1.5%	1.5%	7.0%	12.0%	EIA Low	100%	25.2	$ 806	**$ 32.03**
$3 Gas	3.0%	1.5%	1.5%	7.0%	12.0%	$3 gas	100%	25.2	$ 880	**$ 35.01**
Half Learning Rates	3.0%	0.8%	0.8%	3.5%	6.0%	EIA Ref.	100%	25.2	$ 1,105	**$ 43.95**
Low LBD/Hi Time	3.0%	3.5%	5.0%	3.5%	6.0%	EIA Ref.	100%	25.2	$ 794	**$ 31.58**
No LBD/All Time	3.0%	4.0%	7.0%	0.0%	0.0%	EIA Ref.	100%	25.2	$ 1,041	**$ 41.40**
All LBD/No Time	3.0%	0.0%	0.0%	11.0%	15.0%	EIA Ref.	100%	25.2	$ 437	**$ 17.38**
No Learning (LBD or time)	3.0%	0.0%	0.0%	0.0%	0.0%	EIA Ref.	100%	25.2	$ 1,753	**$ 69.71**
2.5% Discounting	2.5%	1.5%	1.5%	7.0%	12.0%	EIA Ref.	100%	25.2	$ 658	**$ 26.15**
5% Discounting	5.0%	1.5%	1.5%	7.0%	12.0%	EIA Ref.	100%	25.2	$ 491	**$ 19.54**
Wait = 2050	3.0%	1.5%	1.5%	7.0%	12.0%	EIA Ref.	100%	54.6	$ 423	**$ 7.75**
Delay Hurry 1 year	3.0%	1.5%	1.5%	7.0%	12.0%	EIA Ref.	100%	23.4	$ 553	**$ 23.66**
Delay Wait 1 year	3.0%	1.5%	1.5%	7.0%	12.0%	EIA Ref.	100%	26.9	$ 639	**$ 23.77**
Half Decarbonization	3.0%	1.5%	1.5%	7.0%	12.0%	EIA Ref.	50%	15.7	$ 436	**$ 27.85**
Pessimistic ($3 Gas, Half Learn)	3.0%	0.8%	0.8%	3.5%	6.0%	$3 gas	100%	25.2	$ 1,366	**$ 54.33**
Ex Pessimistic ($3 Gas, No Learn)	3.0%	0.0%	0.0%	0.0%	0.0%	$3 gas	100%	25.2	$ 2,014	**$ 80.09**

Lower gas prices make it relatively more costly to switch quickly to renewable energy sources. However, in our model the ACC is not very sensitive to lower gas prices. Using a much lower gas price projection (EIA's High Oil and Gas Resource case, where gas price reaches just $4 by 2034), the ACC rises only to about $32 (vs. $25/ton in our Base Case). And if gas prices remained at $3/MMBtu (real) forever, the ACC is only $35/ton. Of course, a higher gas price trajectory would lead to an ACC below $25. This relative lack of sensitivity to gas price is due to the fact that the bulk of renewable energy sources will be installed after significant additional experience has been gained and renewable costs have thus fallen significantly (through a combination of learning and other factors). Discounting also plays a role, since the largest cost differences occur some years in the future after the generation mix has diverged significantly, but since the discount rate used is low (3%), this is not a major effect.

One of the more important assumptions is the extent and pace of **renewable cost declines** over time, and how dependent they are on the level of deployment. All else equal, smaller cost declines, or cost declines that are less dependent on deployment, will tend to increase the ACC and thus favor the Wait path. If costs decline by only half as much overall (with the same proportion deployment-related), meaning that wind and PV costs fall only to about $60 and $70/MWh, respectively, the implied avoided cost of carbon would nearly double ($44), which is not surprising. If costs decline by about the same amount overall, but are less dependent on deployment, then avoided carbon cost rises to $32/ton, or to $41/ton if costs depend only on time and not at all on deployment. The higher ACC is primarily the result of the Wait path performing better if costs decline even absent accelerated deployment. On the other hand, if cost decline is related only to deployment, as with one-factor learning models, Hurry is relatively more attractive, and the ACC falls to $17/ton. Finally, assuming that (even though unlikely) there will be no learning at all, the Hurry path would look much less attractive, with a $70/ton ACC. This value is similar to the apparently intuitive but ultimately misleading intuition described above, but it does confirm the important role of learning and cost declines over time.

We also tested the sensitivity of our results to the choice of **discount rate**. Reflecting an approach closer to social rather than private discounting, the discount rate used in our Base Case is a modest 3%, the same rate used by the U.S. government for its central value SCC estimate. The SCC is evaluated at alternative discount rates of 2.5% and 5%, and so we tested the sensitivity of our ACC results to these rates also; it is important to maintain consistency in the discount rate when comparing the ACC and the SCC. A lower discount rate generally increases the present value of avoided CO2 cost, and a higher rate decreases it. The same happens when discounting social costs, but since the SCC is evaluated over a much longer timeframe – as far out as 2300, rather than just the 2050 time horizon of our ACC analysis – the SCC is much more sensitive to the choice of discount rate. This means that at lower discount rates, those far-distant future social costs loom larger, whereas at higher discount rates, the nearer-term generation costs are relatively more important. At the lower 2.5% discount rate, the ACC increases very slightly to $26/ton, but the corresponding SCC value is much higher, about $58/ton, leaving the ACC farther below the SCC than in our Base Case. At the higher 5% rate, the ACC drops to just below $20/ton, but the relevant SCC value at 5% falls by more, to $9/ton. At lower discount rates, the ACC increases but the SCC increases by considerably more, leaving the ACC farther below the SCC. At higher rates, the ACC falls but the SCC falls by more, and the relative values may cross.

We also compared the **relative timing of the renewable build out**, examining a much delayed Wait path where accelerated renewable deployment would only begin in 2050. Given our time horizon (looking through 2050), this effectively implies no decarbonization at all. In this sensitivity, the ACC falls dramatically to below $8/ton, primarily driven by the fact that cumulative emissions would differ by much more between the Hurry and modified Wait path, compounded by the fact that high projected gas prices in the late 2030s make fossil resources more costly than renewables. This shows that decarbonizing will ultimately reduce costs (since fossil fuel costs rise, and renewable costs will fall somewhat even without accelerating their deployment, they eventually cross) and is desirable based on its prima facie economics, even without considering the social cost of carbon. Our study was designed to test a stronger question – whether it may be worthwhile to accelerate renewable deployment even faster than its direct economics would suggest, to reduce cumulative carbon emissions.

Our simple characterization largely ignores potential renewables integration challenges that may arise at high levels of renewable deployment, a topic we discuss separately below. But to acknowledge that integration worries might lead to concerns about full decarbonization, we also tested a sensitivity in which the Hurry and Wait paths lead to half rather than full decarbonization. This might reflect a two-step approach, with the second half of decarbonization proceeding only once/if integration challenges can be addressed satisfactorily. Since half decarbonization misses the cost reductions associated with the final doubling of renewable generation, average renewable costs fall by less, which increases the average ACC to $28/ton. While this is not a dramatic increase, it does illustrate the point that the Hurry path offers increasing benefits as the renewable rollout progresses. Once significant deployment, much learning and the associated cost reductions have occurred, it will be possible to install even larger amounts of renewables at even lower costs to build on the earlier progress.

Finally, we also looked at a very pessimistic case that combines a negative view on both of the two most important factors that affect the implied cost of avoiding carbon by accelerating the renewable build-out: future natural gas prices, and the rate at which renewable costs decline with time and experience. This pessimistic case combines persistent low gas prices, modeled at $3/MMBtu indefinitely, with both the time and learning components of renewables cost declines assumed to be half our base case assumptions and results in an ACC of $54/ton, somewhat above the estimated Social Cost of Carbon. In a much more extreme case that assumes that there will be no further renewable cost reductions beyond their current level, combined with $3 natural gas indefinitely, the ACC reaches $80/ton, significantly above most estimates of the social cost of carbon. Although this is admittedly based on extremely negative and unrealistic assumptions, it is instructive because it illustrates the level of pessimism that is required to counter the conclusion that an accelerated renewable deployment strategy would probably reduce carbon emissions at a cost below the social cost of carbon.

In summary, our sensitivity analyses show that under a broad range of assumptions, including negative perspectives on the potential for renewable cost reductions and the price of competing natural gas, the cost of avoided CO2 emissions from a Hurry path is near or below the U.S. government's estimate of the SCC. Even combining two of the most pessimistic assumptions – halving of learning rates, and $3 gas indefinitely – pushes the ACC to $54/ton. While this is above the $37/ton central value estimate of the SCC, it is not dramatically so for such a pessimistic scenario, and it is still well below the $113/ton 95th percentile value used to characterize climate change risk.

Conservatisms in the Analyses

It is worth noting that we make a number of conservative assumptions in our analysis, which strengthen our conclusion that the avoided cost of carbon by Hurrying is likely to be low relative to the SCC. First, our sensitivity analyses have focused on negative effects – low gas prices, lower renewable learning rates – that make the Hurry path relatively less attractive. But these factors could go in a positive direction instead – gas prices may rise higher or sooner, and/or renewable costs may fall faster due to unforeseen technological innovation. Either of these would reduce the ACC farther below the SCC. Also, truncating our analysis at 2050 tends to overstate the actual ACC. It leaves out the fact that Hurry costs would remain below Wait costs beyond 2050; since cumulative renewable deployment is always higher with Hurry, the greater experience at any point in time means Hurry costs continue to be somewhat below Wait costs for second and later generations of renewables as well. Finally, the fossil generation cost estimates we use may be understated, particularly for the Wait path. We assume that fossil costs consist of only the current to-go costs of existing plants, but if those plants are to remain in service for several decades longer, as they do on the Wait path, their costs could be higher. Aging fossil plants may have increasing maintenance needs or new environmental requirements that could substantially increase their operating costs or capital expenditure needs, and additional new gas plants (whose capital costs would contribute to overall costs) are likely to be needed to replace retiring coal plants and meet load growth on the Wait path, where gas generation increases by 75% from 2015 to 2034. Further, by assuming that average fossil costs are the same on the Hurry and Wait paths, we ignore the fact that less efficient, higher cost fossil capacity may be retired first as renewables are added, reducing the average cost of remaining fossil and doing so more quickly in the Hurry case, where fossil retires more quickly. All these factors mean that our cost assumptions are conservative and actual costs under the Wait case may be higher, reducing the ACC below our modeled result. On the other side of the equation, the SCC values are based on climate models that omit some impacts, and thus they likely understate the full social cost of carbon.

Further, since the SCC measures discussed above are based on the average outcomes of climate models, they do not capture the complexity of uncertainties about how GHG concentrations affect climate. The true distributions themselves are unknown and quite possibly have "fat tails" – i.e., higher probabilities of extreme outcomes – due to positive feedback loops, or tipping points beyond which climate impacts and the costs of adapting could become catastrophically large. To account for these possibilities and the potential insurance value of avoiding them by reducing carbon emissions, the U.S. government's interagency working group on SCC recommends considering the 95th percentile SCC value from the models, in addition to the average value. This 95th percentile value (restated in 2015$, present value to 2015) is $113/ton, far above not only our Base Case ACC, but also well above any of the negative sensitivity analysis results. Recall the pessimistic case above that halved the learning rates and also assumed $3 gas indefinitely resulted in an ACC of $54; while this is above the $37/ton expected costs of climate damage, it is less than half of this 95th percentile measure that the government suggests for evaluating extreme climate risks. And even assuming no renewable cost improvement, combined with indefinite $3 gas, the ACC of $80/ton is still well below this 95th percentile measure. Beyond this, some recent research suggests that the true social cost of carbon may be far above what these climate models estimate for other reasons. This may be because climate damage persists through time, having a long-term downward impact on economic growth rates, especially in poor countries. Others have suggested that a slow start to cutting carbon emissions could result in a "hard landing" for the economy that exacerbates systematic financial risks: forcing sudden and costly changes in energy use, causing a sudden devaluation of carbon-intensive assets which could destabilize debt markets, or causing natural catastrophes that impact insurers' liabilities.[21]

All this implies that even taking into account only the expected social cost of greenhouse gas emissions but not considering the catastrophic potential of extreme climate change risks, and accounting for some negative factors that could increase the cost, a Hurry approach is likely to be justified on economic grounds for the expected social costs it avoids. The insurance value that Hurrying provides against potential extreme climate risks strengthens this conclusion further; none of the scenarios we evaluated yielded an avoided cost of carbon as high as the 95th percentile SCC. Essentially, the insurance that Hurry offers against extreme outcomes is likely to be essentially "free", and can be had by just choosing the strategy with highest expected value (including the social cost of carbon). This implies that at the very least, any effort to slow the pace of renewable deployment in order to take advantage of future cost reductions is likely counter-productive. While arguments about "over-paying" for the current generation of renewables or high integration costs may appear to be attractive, they cannot assume "all else equal", since cumulative greenhouse gas emissions in this case will not be equal.

VI. Renewable Integration at high levels of Renewable Energy

One potential shortcoming of our analysis is the fact that we have not explicitly included the cost of "integrating" renewable generation at very high penetration levels. The challenges and costs of renewables integration have been the subject of numerous studies, with most U.S. studies to date focusing on renewable (wind and/or solar PV) shares of 30% or less, and finding integration costs in the range of $5-12/MWh of renewable generation.[22] Accordingly, our analysis does include an estimated integration cost of $6/MWh of renewable generation[23], which we assume to be subject to the same time and learning trends applied to the costs of renewable resources themselves. This assumption is consistent with the observation that, while to our knowledge the issue has not been (as) scientifically studied, it appears that estimates of the levels of renewables that can be safely integrated have increased, and the estimated cost of doing so has decreased, both over time and with increasing actual renewables penetration.

This is not to say that integrating very large shares of renewable energy and perhaps ultimately operating a system comprised almost entirely of non-emitting resources will be easy or cheap. For example, as larger amounts of renewables are added and as dispatchable fossil generation is retired, integration likely becomes more challenging, and the cost may increase significantly once renewable penetration rises beyond some threshold – e.g., if significant transmission and energy storage solutions become necessary. On the other hand, the available techniques and technologies for integrating intermittent generation are likely subject to similar types of learning, scaling and time effects as renewable generation itself, which means that integrating a given amount of renewables should get easier and cheaper with time and experience. And at high renewable penetration, entirely different approaches may be warranted, which may have different cost implications than the approaches used at lower penetration levels.

The relatively low ACC we estimate, and the conservatisms we have included, suggest that even if integrating very high shares of renewables becomes significantly more costly than we have assumed, there is substantial "head room" for larger integration costs that would still leave the ACC below the corresponding social cost of carbon. While it is very difficult to estimate future renewable energy integration costs at high penetration, one way to proxy for this effect in our model would be to include the integration costs implicitly in the cost of renewables, by using lower learning rates to characterize their cost improvement over time. This could represent a higher learning rate for the actual renewable generation, but with integration costs that increase with cumulative deployment. As our sensitivity analysis shows, assuming only half the base case learning rate would increase our average cost of carbon to around \$43/ton, somewhat above the \$37/ton central social cost of carbon estimate but still far below the 95th percentile cost. And given that we found the ACC is also quite low for half-decarbonization, the possibility that integration costs might increase sharply beyond some high penetration threshold does not argue for waiting. Hurrying to at least half-decarbonization would be more beneficial than Waiting from the outset.

VII. Putting the cost in historical perspective

We turn now to an alternative framing of the question, to ask how much society has been willing to spend in other situations that involved catastrophic risks or large scale society-wide transformational efforts. Defense efforts in time of war are likely examples of how societies react to existential – and unlike climate change, more imminent and obvious – threats. A clear example is the threat of Nazi Germany in the early 1940s, which was certainly perceived by the Allies as a major threat to global civilization. According to some historic estimates, the United States devoted between 30-40% of real GDP to the war effort between 1942 and 1944. Over the same period, the Soviet Union devoted 53-61% of its GDP, and the United Kingdom dedicated about half of GDP to war mobilization in each of the years 1940-1945.[24, 25] Without speculating about how many years the Allied Forces would have been able or willing to maintain these levels of mobilization, these figures show that existential threats can prompt societies to mobilize incredible resources – on the order of 50% of GDP – for an arguably unknowable chance of avoiding catastrophe.

Alternatively, programs like the Manhattan project that are designed to create technology to win wars, or the Apollo program, at least partly inspired by military/political competition with the Soviet Union but perhaps also intended to provide other large societal benefits, have involved significantly smaller outlays. The cumulative cost of the Manhattan project was approximately \$25 billion over five years and the cost of the Apollo program approximately \$110 billion over 13 years (both in 2015 dollars). In any given year, neither required outlays in excess of 0.4% of GDP.[26]

By comparison, in our Base Case, the incremental cost to pursue the Hurry path rather than Wait is \$620 billion in present value. This is about 3% of a single year's GDP, a much smaller fraction than the annual efforts expended during crises like World War II, and the cost is actually spread over many years. On an annual basis, the years with highest additional cost would require on the order of \$50 billion more, less than 0.3% of 2015 GDP which is below the GDP share of the Apollo program. In our pessimistic sensitivity case of halved learning rates and indefinite \$3 gas, the additional costs of Hurrying could total \$1.4 trillion in present value, about \$150 billion per year in the highest-cost years, 0.8% of 2015 U.S. GDP which is roughly twice the share of the Apollo program. While obviously imperfect, these comparisons provide some indication that the greatly accelerated deployment of renewable generation would likely involve costs at the low end of what societies have been willing to spend on reducing existential threats and/or creating important options for the future.

VIII. Conclusions

Increasing climate change concerns create a compelling reason to reduce greenhouse gas emissions deeply and quickly. Fortunately the cost of emissions-free renewable energy technologies has been falling rapidly. However, falling costs also intuitively suggests that waiting until costs have fallen further might allow us to decarbonize at significantly lower cost. We developed a simple modeling framework to analyze the potential incremental cost and the carbon emissions that would be avoided by decarbonizing quickly versus waiting for further reductions in technology cost. We compared the resulting avoided cost of carbon to the estimated social cost of carbon used by the U.S. government in evaluating policies designed to lower greenhouse gas emissions. Our model incorporates a simple approach to projecting the future evolution of renewable generation costs, using a combination of a pure time trend and a learning rate that relates cost reductions to cumulative installed capacity. Our Base Case results show that the incremental costs of beginning immediately to accelerate decarbonization would lead to relatively modest additional costs compared to delaying the decarbonization effort until 2030. Specifically, we find that the net present value of additional cost would be about $620 billion (this value could be overstated due to conservatisms in our analysis). While this figure is large in absolute terms, it represents only about a 14% increase in the total cost of electricity generation over the next 35 years, and even in the highest-cost years, is less than 0.3% of GDP. Not counting any of the climate benefits of a Hurry path, this would mean dedicating approximately one year of economic growth to early and rapid decarbonization.

Hurrying would significantly lower cumulative carbon emissions, at a cost of about $25/ton of CO2. This is significantly below estimates for the social cost of carbon, which center around $37/ton. We performed a number of sensitivity analyses which show that our results are generally robust, even to significant negative variation in important input parameters. There are also a number of conservatisms in our analysis and in the comparison with the SCC that further strengthen the result. This suggests that early and rapid decarbonization is likely both cost-effective and prudent.

Finally, this comparison of the avoided cost per ton of CO2 with the expected social cost of carbon does not reflect another important benefit of early decarbonization, namely reducing the risk of ending up with a catastrophic result in the "fat tail" of the uncertainty distribution on climate change outcomes. Such an outcome may be triggered by positive feedback loops or points of no return in the climate system. Acknowledging that this would amount to an unknown reduction in a poorly understood risk, it is nonetheless worth noting that the magnitude of this risk could rival that of any previous challenge faced by humanity. The financial cost of Hurry can be thought of as insurance against such a risk of catastrophic climate outcomes, though it appears that the cost would be justified solely by the expected value of avoiding the sort of climate change damages that scientists believe they already understand.

[1] The same climate risks will likely require significant reductions in greenhouse gas emissions from other sectors, but here we focus on the electricity sector, which is widely believed to be the "low hanging fruit". In fact, electrification may be an effective way to decarbonize other sectors (e.g., via electric vehicles, electric heating, etc.), which makes the pace of decarbonizing the electric sector even more important.

[2] There is some evidence that this type of waiting is already happening. Having witnessed rapid declines in the prices obtained in recent long-term procurements of PV power, the City of Austin, TX decided to postpone additional procurements of PV under long-term contracts so as to be able to benefit from even lower prices in a few years. See http://www.utilitydive.com/news/as-solar-prices-drop-fast-austin-energy-seeks-to-delay-chunk-of-600-mw-pur/406237/ .

[3] The value of waiting in this context has been pointed out by William H. Hogan, Clean Energy Technologies: Learning By Doing And Learning By Waiting, Energy Policy Research Seminar, September 29, 2014. Of course, waiting for others to deploy renewable energy represents a free rider problem – those who wait for technology costs to fall benefit from the early investments made by others, creating a temptation for all to follow that strategy and potentially undermining the desired deployment-related cost reduction benefits.

[4] For a discussion of fat tails and associated catastrophic risks, see for example Martin Weitzman , *Fat-Tailed Uncertainty in the Economics of Catastrophic Climate Change*, *Review of Environmental Economics and Policy 5 (2)*, 2011, p. 275–292. For a more general overview about incorporating uncertainty about climate outcomes into decision making, see Geoffrey Heal and Antony Millner, "Uncertainty and Decision Making in Climate Change Economics," *Review of Environmental Economics and Policy 8 (1)*, 2014, p. 120-137.

[5] The social cost of carbon estimates developed for use in benefit-cost analysis by U.S. government agencies do include a 95th percentile cost (using a 3% discount rate) to illustrate some of the risks associated with climate change above expected levels.

[6] This article is an extension of an earlier and shorter article: Eleanor Denny and Jurgen Weiss, "Hurry or Wait – The Pros and Cons of Going Fast or Slow on Climate Change," *The Economist's Voice 12 (1)*, 2015, p. 19-24.

7 Even though our numerical estimates are for the United States, the implications are likely applicable more broadly. While the United States enjoys some of the world's lowest fossil fuel costs, it also benefits from some of the world's best renewable resources (sun in the Southwest, wind in the Midwest, relatively low land costs, etc.).

8 The remaining currently-existing generation, primarily nuclear and hydro, is classed as "Other" and is assumed for simplicity to remain constant through 2050.

9 Global annual growth rates for both wind and solar PV vary significantly year by year and have exceeded our assumed accelerated growth rates historically on several occasions. Also, other entities have reported higher compound annual growth rates. For example, see Fraunhofer ISE, "Photovoltaics Report", March 11, 2016, which reports a global compound annual growth rate of solar PV of 44% between 2000 and 2014 (p.4), with significantly increased pace of deployment after about 2006. It might be argued that growth at the rates we simulate, while feasible on a small base, may become a challenge to maintain on a

much larger base. That is, a year-on-year 40% increase in installed PV may be possible when PV is 0.5% of total U.S. generation, but to continue to grow at that same relative rate when the installed base becomes 20% of generation is a different matter. We are not arguing that any particular growth rate is possible or could be sustained as the installed base becomes large, but there is clearly some flexibility in how quickly a renewable build-out proceeds, which may be influenced by public policy. We are looking at the cost implications of faster vs slower expansion, and using two essentially arbitrary deployment schedules to investigate this issue.

10 This and several other cost assumptions make our results conservative, as discussed later.

11 Lazard, "Lazard's Levelized Cost of Energy Comparison - Version 9.0", 2015, p. 2.

12 Lazard reports cost ranges for 5 types of solar PV installations. We assume no thin-film PV, a 40% share of utility scale projects, a 30% share of residential PV projects and a 15% share each of rooftop PV for C&I customers and community solar projects respectively. Lazard reports a range of low-high. For solar PV, we assumed a 50% share of low cost and a 50% share of high cost projects. For wind, we assumed a 100% share of high cost wind.

13 For a recent review of the use of learning rates in the energy field, see Edward S. Rubin, Inês M.L. Azevedo, Paulina Jaramillo and Sonia Yeh, "A review of learning rates for electricity supply technologies," *Energy Policy 86*, 2015, p. 198–218.

14 We decomposed the observed average one-factor model learning rates for wind (12%) and solar PV (23%) into time and learning components so that our assumed accelerated deployment rates would lead to approximately equal contributions from the time and learning components over a 10-year period, while maintaining approximately the same overall cost declines as seen with a single-factor model that uses these same average learning rates.

15 See http://www.renewableenergyworld.com/articles/2016/08/solar-sold-in-chile-at-lowest-ever-half-price-of-coal.html and http://cleantechies.com/2016/08/18/wind-energy-scores-big-in-chiles-electricity-auction/.

16 Lazard (2015), p. 9.

17 Renewable deployment provides a hedge against long-term fossil price volatility (evidenced by historic rate fluctuations driven by underlying fossil fuel price volatility). Earlier deployment under the Hurry path provides an additional and earlier hedge. We do not consider this value in our analysis, but historic rate fluctuations show (and our sensitivity analysis confirms) that total costs of hurrying could be significantly smaller under higher gas prices.

18 See http://edgar.jrc.ec.europa.eu/news_docs/jrc-2015-trends-in-global-co2-emissions-2015-report-98184.pdf

19 In the illustrative scenario outlined here, wind capacity increases to seven times the current installed base; this is two more doublings before installing half of the ultimate wind capacity. Solar PV deployment grows even more. It will increase seventy-fold, meaning the current installed base will double five more times before installing half of the PV capacity that would ultimately be installed.

20 For the current U.S. EPA SCC estimates, see the EPA Fact Sheet, Social Cost of Carbon, December 2015. For a more complete overview of the methodology used to calculate the SCC, see Technical Support

Document: Social Cost of Carbon for Regulatory Impact Analysis - Under Executive Order 12866, Interagency Working Group on Social Cost of Carbon, United States Government, February 2010. The UK uses a central 2020 SCC estimate of £90/ton, rising to £120/ton by 2050. See Paul Watkiss, The Social Cost of Carbon, which also cites SCC estimates derived from other Integrated Assessment Models (IAMs), which tend to be somewhat lower than the UK estimate for 2020 but higher after 2030.

21 See, for example, F. Moore and D. Diaz, Temperature impacts on economic growth warrant stringent mitigation policy, *Nature Climate Change* ***5***, 2015, p. 127-131. Also, European Systemic Risk Board, "Reports of the Advisory Scientific Committee – Too late, too sudden: Transition to a low-carbon economy and systemic risk," No 6, February 2016.

22 See for example Milligan et al., "Review and Status of Wind Integration and Transmission in the United States: Key Issues and Lessons Learned," NREL, March 15, p. 2. See also Ahlstrom et al., "Relevant Studies for NERC's Analysis of EPA's Clean Power Plan 111(d) Compliance," NREL, June 2015, Table ES-1, which lists recent renewables integration studies and the renewables penetration levels examined.

23 Specifically, we use the mid-range of integration costs reported by Lazard, which range from $2-$10/MWh.

24 See Mark Harrison, "The economics of World War II: an overview," Table 1-8.

25 See Mark Harrison, "Resource mobilization for World War II: the U.S.A., U.K., U.S.S.R., and Germany, 1938-1945," *Economic History Review 41 (2)*, 1988, p. 171-192, Table 3.

26 Deborah Stine, "The Manhattan Project, the Apollo Program, and Federal Energy Technology R&D Programs: A Comparative Analysis," Congressional Research Service, June 30, 2009.