



Electricity Markets & Policy
Energy Analysis & Environmental Impacts Division
Lawrence Berkeley National Laboratory

Plant-level performance and degradation of 31 GW_{DC} of utility-scale PV in the United States

Mark Bolinger, Will Gorman, Dev Millstein, Dirk Jordan¹

¹National Renewable Energy Laboratory

March 2022



This work was supported by the Solar Energy Technologies Office of the U.S. Department of Energy's Office of Energy Efficiency and Renewable Energy under Lawrence Berkeley National Laboratory Contract No. DE-AC02-05CH11231.

Disclaimer

This document was prepared as an account of work sponsored by the United States Government. While this document is believed to contain correct information, neither the United States Government nor any agency thereof, nor The Regents of the University of California, nor any of their employees, makes any warranty, express or implied, or assumes any legal responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by its trade name, trademark, manufacturer, or otherwise, does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or any agency thereof, or The Regents of the University of California. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or any agency thereof, or The Regents of the University of California.

Ernest Orlando Lawrence Berkeley National Laboratory is an equal opportunity employer.

Copyright Notice

This manuscript has been authored by an author at Lawrence Berkeley National Laboratory under Contract No. DE-AC02-05CH11231 with the U.S. Department of Energy. The U.S. Government retains, and the publisher, by accepting the article for publication, acknowledges, that the U.S. Government retains a non-exclusive, paid-up, irrevocable, worldwide license to publish or reproduce the published form of this manuscript, or allow others to do so, for U.S. Government purposes.

Acknowledgements

The work described in this study was conducted at Lawrence Berkeley National Laboratory and supported by the Solar Energy Technologies Office, Office of Energy Efficiency and Renewable Energy, U.S. Department of Energy under Contract No. DE-AC02-05CH11231.

Table of Contents

Acknowledgements.....	i
Table of Contents.....	ii
Table of Figures.....	ii
List of Tables.....	ii
Executive Summary.....	iii
1. Introduction.....	4
2. Data Sample.....	5
3. Assessment of First-Year Plant-Level Performance.....	8
4. Assessment of Plant-Level Performance Degradation.....	10
5. Discussion and Conclusions.....	14
6. References.....	16

Table of Figures

Figure 1. Histogram of individual plant capacity within the sample.....	7
Figure 2. Solar curtailment history in CAISO and ERCOT.....	8
Figure 3. Actual vs. modeled first-year capacity factor by (a) year and (b) region.....	9
Figure 4. Age fixed effects and best-fit line for final model specification.....	12
Figure 5. Comparison of updated (blue) to original (orange) results.....	13

List of Tables

Table 1. Geographic descriptive statistics of sample.....	6
Table 2. Temporal descriptive statistics of sample.....	7
Table 3. Statistically significant comparisons of sub-samples using the fixed effects model.....	14

Executive Summary

In this updated study, which samples 50% more capacity than the original and adds two additional years of operating history, we assess the performance of a fleet of 631 utility-scale PV plants totaling 31.0 GW_{DC} (23.6 GW_{AC}) of capacity that achieved commercial operations in the United States from 2007-2018 and that have operated for at least two full calendar years. We use detailed information on individual plant characteristics, in conjunction with modeled irradiance data, to model expected or “ideal” capacity factors in each full calendar year of each plant’s operating history. A comparison of ideal versus actual first-year capacity factors finds that this fleet has modestly underperformed initial expectations (as modeled) on average, though perhaps due as much to modeling issues as to actual underperformance. We then analyze fleet-wide performance degradation in subsequent years by employing a “fixed effects” regression model to statistically isolate the impact of age on plant performance. The resulting average fleet-wide degradation rate of -1.2%/year ($\pm 0.1\%$) represents a slight improvement (seemingly driven by the oldest plants in our sample) over the -1.3%/year ($\pm 0.2\%$) found in our original study, yet is still of greater magnitude than is commonly found. We emphasize, however, that these fleet-wide estimates reflect both recoverable and unrecoverable degradation across the entire plant, and so will naturally be of greater magnitude than module- or cell-level studies, and/or studies that focus only on unrecoverable degradation. Moreover, when focusing on a sub-sample of newer and larger plants with higher DC:AC ratios—i.e., plants that more-closely resemble what is being built today—we find a more moderate sample-wide average performance decline of -0.7%/year ($\pm 0.4\%$), which is more in line with other estimates from the recent literature.

1. Introduction

The deployment of photovoltaic (PV) modules in large, utility-scale configurations is a relatively recent phenomenon. In the United States, the first two utility-scale PV plants—defined here to include any ground-mounted PV plant larger than 5 MW_{AC}—achieved commercial operations as recently as late 2007. Thus, at the time of writing (in late 2021), the oldest utility-scale PV plants in the United States had only been operating for thirteen full calendar years (2008-2020), or less than half of their expected useful lives of 30 year or longer. Moreover, most utility-scale PV plants in the United States are much younger than these oldest plants: at the end of 2020, there were ~38.7 GW_{AC} (~50.9 GW_{DC}) of utility-scale PV operating in the United States, with an average plant age of less than four years (Bolinger et al., 2021).

With such a young and rapidly growing fleet of utility-scale PV plants supplying the majority of solar generation in the United States, it is crucial to understand how these utility-scale plants have performed to date. Bolinger et al. (2020) found that a fleet of 21.0 GW_{DC} of utility-scale PV plants built in the United States from 2007-2016 had generally lived up to *ex ante* expectations for initial performance, but that average plant-level degradation of -1.3%/year through 2018 had been worse than both *ex ante* expectations and results from other degradation studies. Specifically, Bolinger et al. (2020)'s review of power purchase agreements (PPAs) and the existing literature found that most PPAs codify an expected performance degradation rate of just -0.5%/year, based in part on earlier studies in the literature, many of which focus on unrecoverable degradation at the module, rather than plant, level. The minority of studies that have looked more broadly at plant-level degradation have tended to find higher degradation rates, closer to -1.0%/year (Bolinger et al., 2020). In the interim since the publication of Bolinger et al. (2020), the performance of utility-scale PV plants in the United States has come under increased scrutiny (kWh Analytics, 2020, 2021a, 2021b; Deline et al., 2021; Jordan et al., 2021; Kharait and Chan, 2021), in some cases with charges of significant downside deviations from so-called P50 (or median) pre-construction estimates, both initially and over time (kWh Analytics, 2020, 2021a, 2021b). It is within this environment of heightened interest that we offer this significantly expanded update to Bolinger et al. (2020).

Following the approaches detailed in Bolinger et al. (2020), this update from the same authors expands the plant sample by ~50% (i.e., by 220 plants and 9.86 GW_{DC}) and adds two additional years of operations (2019 and 2020). Specifically, we assess the plant-level performance of a fleet of 631 utility-scale PV plants totaling 31.0 GW_{DC} of capacity that achieved commercial operations in the United States from 2007-2018, and thus have been operating for at least two (2019 and 2020) and as many as thirteen (2008-2020) full calendar years. Using detailed information on individual plant characteristics, in conjunction with modeled irradiance data, we assess the extent to which actual first-year performance has lived up to expected (as modeled) performance. We then analyze fleet-wide degradation in energy output in subsequent years, by employing a “fixed effects” regression model to statistically isolate the impact of age on system performance.

In what is perhaps an indication of the robustness of our methods, we find results similar to our original study. On average, these plants' first-year performance has fallen short of modeled expectations to a modest degree. The fleet-wide degradation rate of $-1.2\%/year$ ($\pm 0.1\%$) represents a slight improvement over the $-1.3\%/year$ ($\pm 0.2\%$) found in our original study, yet is still of greater magnitude than is commonly assumed. We emphasize, however, that these fleet-wide estimates reflect both recoverable and unrecoverable degradation across the entire plant, and so will naturally be of greater magnitude than module- or cell-level studies, and/or studies that focus only on unrecoverable degradation. Moreover, application of the fixed effects model to a variety of sub-samples in an attempt to tease out potential degradation drivers suggests that newer and larger plants with higher DC:AC ratios—i.e., plants that more closely resemble recent plants—have experienced a lower magnitude of degradation— $-0.7\%/year$ ($\pm 0.4\%$)—that is more in line with other estimates from the recent literature (e.g., Deline et al., 2021).

2. Data Sample

The updated sample of utility-scale PV plants that we analyze consists of 631 plants totaling $31.0 \text{ GW}_{\text{DC}}$ ($23.6 \text{ GW}_{\text{AC}}$) installed across 37 states from 2007-2018 (Table 1, Table 2, Figure 1). In aggregate, these 631 plants contributed $>50\%$ of all solar electricity generated in the United States in 2019 (across all sectors—residential, commercial, and utility-scale—and including concentrating solar thermal power) and 40% of all solar electricity generated in 2020. Through 2020, they collectively offer 2,916 plant-years of operational experience, almost one-third of which are in California (Table 1). Operational history ranges from 2 to 13 full calendar years, with an average of 4.6 years—once again, indicative of the relative youth of the utility-scale PV sector.

Table 1. Geographic descriptive statistics of sample

State	# of Plants	# of MW _{DC}	# of MW _{AC}	MW _{AC} /Plant		# of Plant-Years	Years per Plant			% of US sample		
				Average	Median		Min	Avg	Max	Plants	MW _{AC}	Plant-Years
CA	186	12,321	9,605	52	20	935	2	5.0	11	29.5%	40.7%	32.1%
NC	59	2,625	1,938	33	20	233	2	3.9	8	9.4%	8.2%	8.0%
AZ	36	2,102	1,567	44	20	225	2	6.3	9	5.7%	6.6%	7.7%
NJ	38	402	328	9	8	190	2	5.0	9	6.0%	1.4%	6.5%
NM	27	606	484	18	10	161	2	6.0	10	4.3%	2.1%	5.5%
TX	34	2,182	1,692	50	23	144	2	4.2	10	5.4%	7.2%	4.9%
NV	21	2,157	1,623	77	50	125	3	6.0	13	3.3%	6.9%	4.3%
GA	24	1,351	984	41	30	109	2	4.5	7	3.8%	4.2%	3.7%
FL	28	1,935	1,332	48	61	104	2	3.7	11	4.4%	5.7%	3.6%
CO	14	510	404	29	18	79	2	5.6	13	2.2%	1.7%	2.7%
OR	22	329	251	11	10	66	2	3.0	4	3.5%	1.1%	2.3%
IN	13	147	109	8	8	66	2	5.1	7	2.1%	0.5%	2.3%
MD	11	250	190	17	12	52	2	4.7	8	1.7%	0.8%	1.8%
UT	12	1,049	810	68	80	50	4	4.2	5	1.9%	3.4%	1.7%
VA	16	577	435	27	20	50	2	3.1	4	2.5%	1.8%	1.7%
SC	17	331	240	14	10	48	2	2.8	4	2.7%	1.0%	1.6%
MN	14	363	252	18	7	43	2	3.1	4	2.2%	1.1%	1.5%
TN	8	187	147	18	16	35	2	4.4	8	1.3%	0.6%	1.2%
ID	8	323	240	30	20	26	3	3.3	4	1.3%	1.0%	0.9%
OH	3	51	38	13	10	22	3	7.3	10	0.5%	0.2%	0.8%
NY	5	103	81	16	10	21	2	4.2	9	0.8%	0.3%	0.7%
IL	2	33	28	14	14	19	8	9.5	11	0.3%	0.1%	0.7%
AL	6	266	198	33	13	18	2	3.0	4	1.0%	0.8%	0.6%
MA	4	52	39	10	9	17	3	4.3	6	0.6%	0.2%	0.6%
DE	2	26	22	11	11	17	8	8.5	9	0.3%	0.1%	0.6%
MI	4	87	72	18	18	10	2	2.5	3	0.6%	0.3%	0.3%
MS	3	215	155	52	52	9	3	3.0	3	0.5%	0.7%	0.3%
PA	1	11	10	10	10	8	8	8.0	8	0.2%	0.0%	0.3%
AR	2	122	94	47	47	7	2	3.5	5	0.3%	0.4%	0.2%
KY	2	25	19	9	9	7	3	3.5	4	0.3%	0.1%	0.2%
MO	2	22	16	8	8	5	2	2.5	3	0.3%	0.1%	0.2%
CT	2	56	40	20	20	4	2	2.0	2	0.3%	0.2%	0.1%
NE	1	7	6	6	6	3	3	3.0	3	0.2%	0.0%	0.1%
WY	1	98	80	80	80	2	2	2.0	2	0.2%	0.3%	0.1%
VT	1	26	20	20	20	2	2	2.0	2	0.2%	0.1%	0.1%
WA	1	28	19	19	19	2	2	2.0	2	0.2%	0.1%	0.1%
OK	1	13	10	10	10	2	2	2.0	2	0.2%	0.0%	0.1%
Total	631	30,990	23,579	37	20	2,916	2	4.6	13	100.0%	100.0%	100.0%

A histogram of plants within our sample by capacity (Figure 1) shows the majority falling into the 20-50 MW_{DC} capacity bin. More than 85% of plants are 100 MW_{DC} or less, but a number of plants feature several hundred MW_{DC} of capacity, with the largest being more than 750 MW_{DC}.

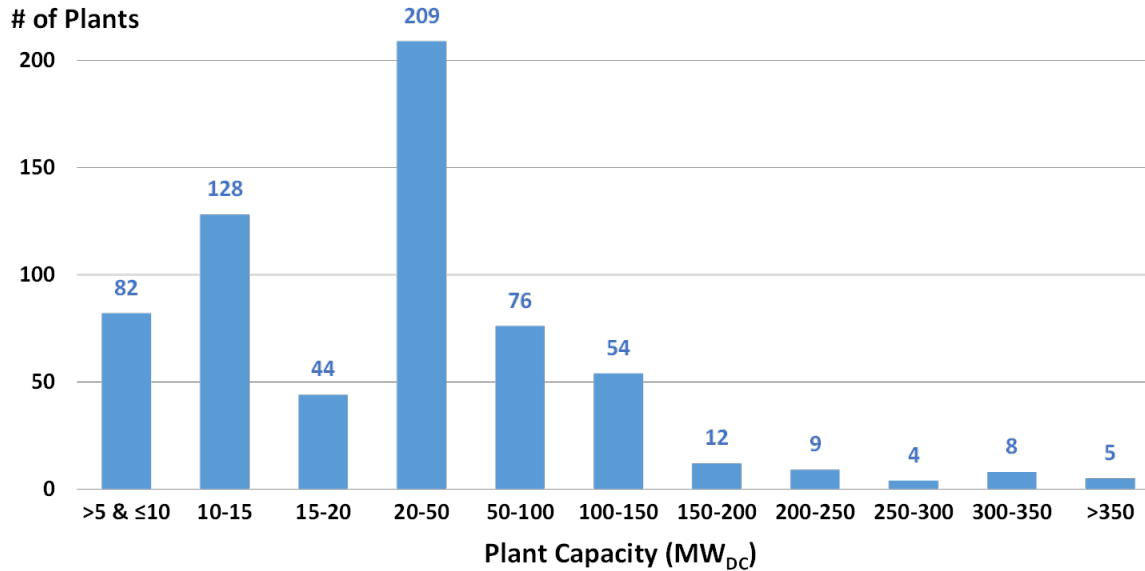


Figure 1. Histogram of individual plant capacity within the sample

Following the approach described in Bolinger et al. (2020), we normalize the performance of each individual plant in the sample by calculating its “capacity factor” in each full calendar year of operations, per Equation 1:

$$Capacity\ Factor_y = \frac{MWh\ generated\ in\ calendar\ year\ y}{(MW_{DC}\ in\ calendar\ year\ y * number\ of\ hours\ in\ calendar\ year\ y)} \quad (Eq.1)$$

Though this empirical measure of capacity factor requires only annual generation (which we source and cross-reference from a combination of Form EIA-923, FERC Form 1, FERC Electric Quarterly Reports, and the California Energy Commission) and capacity at the plant level, our performance analysis methods also require modeling “ideal” capacity factors, for which we need additional plant-level data. Most of these key plant characteristics (or “metadata”) are sourced from Bolinger et al. (2021), and include module type (Si vs. thin-film, primarily CdTe), module manufacturer, mount type (fixed-tilt, single-axis tracking, or dual-axis tracking), tilt (for fixed-tilt mounts), azimuth, coordinates (latitude and longitude), commercial operation date, capacity (MW_{DC} and MW_{AC}), and DC:AC ratio. Table 2 summarizes some of these characteristics across the sample over time.

Table 2. Temporal descriptive statistics of sample

COD Year:	2007	2008	2009	2010	2011	2012	2013	2014	2015	2016	2017	2018	Total
# of Plants:	2	1	3	7	29	36	47	50	82	148	143	83	631
# of MW _{DC} :	22	12	63	172	512	1,133	2,231	3,566	3,575	9,869	4,911	4,923	30,990
# of MW _{AC} :	19	10	54	144	432	886	1,746	2,728	2,745	7,517	3,695	3,602	23,579
Average DC:AC Ratio:	1.18	1.21	1.16	1.19	1.19	1.28	1.28	1.31	1.30	1.31	1.33	1.37	1.31
Average MW _{AC} /Plant:	9	10	18	21	15	25	37	55	33	51	26	43	37
% Plants with Tracking:	100%	0%	67%	14%	55%	50%	60%	62%	65%	73%	75%	61%	66%
% Plants with Si Modules:	100%	0%	67%	29%	79%	83%	81%	74%	76%	86%	84%	77%	81%

We also compile hourly solar curtailment data from the California Independent System Operator (CAISO) and the Electric Reliability Council of Texas (ERCOT), which are currently the only two (of seven) independent system operators (ISOs) in the United States that report curtailment of solar generation. Since all seven ISOs report curtailment of wind generation, we assume that no curtailment of solar generation is happening outside of what is reported in CAISO and ERCOT; this assumption extends to those plants located outside of the seven ISO regions, where we similarly lack data on curtailment. As solar curtailment has grown over time in both California and Texas (Figure 2), it has become increasingly important to control for curtailment when assessing plant performance. We do so by grossing up the empirical capacity factors of solar plants that have been curtailed in both states, using the approach described in Bolinger et al. (2020).

Quarterly and Annual Solar Curtailment




Figure 2. Solar curtailment history in CAISO and ERCOT

Finally, we pull irradiance data from 2008-2020 for each plant location from the National Solar Radiation Database (NSRDB), which uses NREL’s Physical Solar Model to provide solar radiation and meteorological data at 4-kilometer horizontal resolution across 30-minute intervals (NREL, 2021a).

3. Assessment of First-Year Plant-Level Performance

Based on the data described in the previous section, Figure 3 plots “actual” capacity factors (as grossed up to correct for curtailment in California and Texas, if appropriate) versus modeled or “ideal” capacity factors (simulated using NREL’s System Advisor Model (NREL, 2021b)) for each plant in our sample, focusing on the first full calendar year of each plant’s operational history. This first-year comparison (of $CF_{f,t}^{hist}$ to $CF_{f,t}^{ideal}$ at $t=1$, per Equation 2, below) provides an indication of how well the plants in our sample have performed *before* degradation can muddy the waters over time. The scatterplot is shown

two ways: (a) color-coded by calendar year (which is always one year greater than the year of each plant’s commercial operation date) and (b) color-coded by the region in which each plant is located.

The actual first-year capacity factor falls short of modeled for 61% of the plants in our sample, and the scatterplot visually confirms the intuition that there should be more underperforming than overperforming outliers (since there are relatively few reasons that a plant would ever outperform its ideal capacity factor to any significant degree). The median difference between actual and modeled DC capacity factor across the full sample is -0.4 percentage points (or -2.1%); though if considering only underperformers (those below the unity line), the magnitude of the difference increases to -1.1 percentage points of DC capacity factor, or -6%.

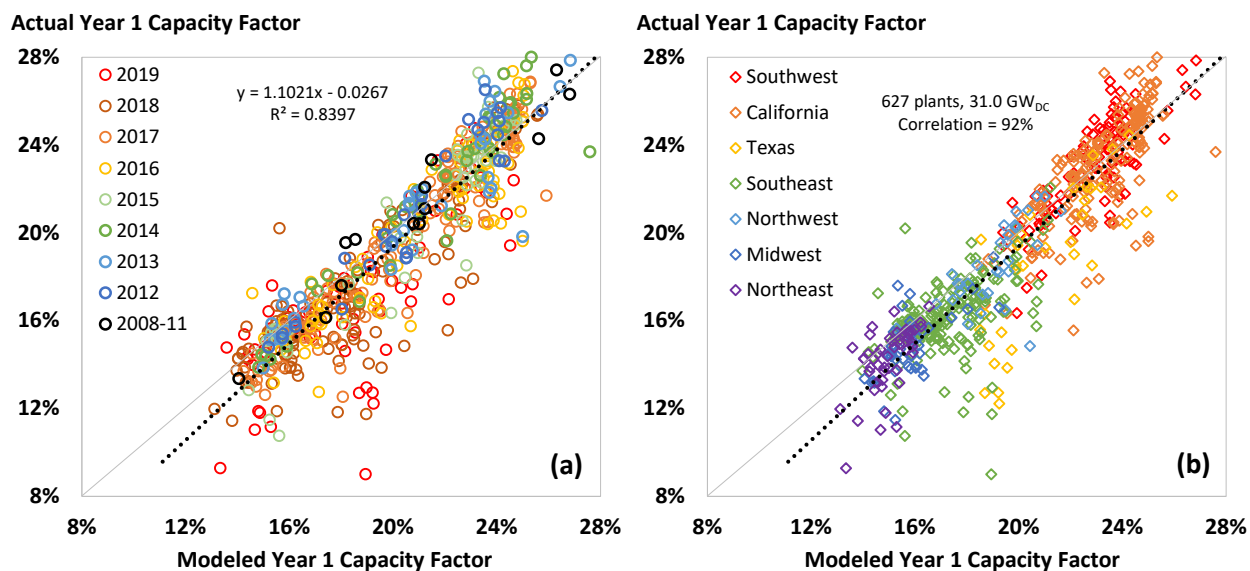


Figure 3. Actual vs. modeled first-year capacity factor by (a) year and (b) region

Turning to the temporal and regional aspects of Figure 3, the underperformance seems to be somewhat skewed to the lower end of the capacity factor range (per the dashed best-fit regression line), with the most significant outliers seemingly concentrated among both newer plants (e.g., those whose first full operational years were 2018 or 2019) and those located in Texas, California, the Southeast, and the Northeast. There are a number of potential explanations for these temporal and regional results—both of which are broadly consistent with findings from kWh Analytics (2021a). For example, newer plants are more likely than older plants to have experienced curtailment in their first year (as curtailment has increased with solar market penetration—see Figure 2), and although we correct for curtailment in both California and Texas, our correction is not perfect (particularly in California, where we use an algorithm to allocate system-wide curtailment data to individual plants—see Bolinger et al. (2020) for more details), and there may very well be unaccounted-for curtailment happening in other regions, such as the Southeast. Newer plants are also more likely to have higher DC:AC ratios (see Table 2), which accentuate the impact of sub-hourly clipping—a phenomenon that is not captured by industry-standard hourly performance modeling but that can reduce actual performance by ~1%-4% (Kharait and

Chan, 2021; kWh Analytics, 2020). Newer plants are also perhaps more likely to be sited on uneven terrain as the market expands and the availability of flat sites diminishes; recent research suggests that uneven terrain is a challenge for backtracking algorithms, as well as performance models, and can reduce plant output by 1%-2% (Kharait and Chan, 2021; kWh Analytics, 2021b). Newer plants are also more likely to have experienced catastrophic weather events in their first year, as the frequency and intensity of hailstorms, hurricanes, tornadoes, and wildfires have increased with global climate change. For example, smoke from wildfires—which is difficult to capture in performance models—reportedly reduced solar generation in the West by as much 6% in 2020 (Keelin and Perez, 2021; kWh Analytics, 2021b). Finally, underperformance in the Northeast in particular may be related to snow cover, which is also hard to model yet is one of the more significant contributors to losses (kWh Analytics, 2020). In short, some or maybe even much of the underperformance shown in Figure 3 may be attributable to inaccurate modeling (or modeling inputs) rather than poorly performing plants. Moreover, despite our attempt to limit the impact of degradation on the comparison in Figure 3 by focusing on just the first full year of operations, early first-year degradation (or even first-year “teething” issues, for plants that came online late in the prior year) could nevertheless be contributing to the apparent underperformance to a limited degree.

4. Assessment of Plant-Level Performance Degradation

Moving beyond the narrow comparison of actual ($CF_{f,t}^{hist}$ in Equation 2) to modeled ($CF_{f,t}^{ideal}$) performance at age one, we now turn to a longitudinal study of how actual performance has held up as plants age. For this purpose, we return to the “fixed effects” regression model used and described in Bolinger et al. (2020), and represented by Equation 2:

$$CF_{f,t}^{hist} = CF_{f,t}^{ideal} + S_f + A_t + \epsilon_{f,t} \quad (\text{Eq. 2})$$

Where,

$CF_{f,t}^{hist}$ = empirical capacity factor of plant f at age t (grossed up, as appropriate, for curtailment in California and Texas)

$CF_{f,t}^{ideal}$ = ideal capacity factor of plant f at age t (modeled using NREL’s System Advisor Model, based on physical plant characteristics and the solar resource at the site)

S_f = site-specific fixed effects for plant f (a dummy variable to control for differences in capacity factor across plants that are not already captured via the ideal capacity factor; expressed as an absolute deviation from a reference plant)

A_t = age fixed effects at age t (a dummy variable to control for differences in capacity factor within plants, over time, that are not already captured via the ideal capacity factor; expressed as an absolute deviation from the average historical capacity factor at age one)

$\epsilon_{f,t}$ = residual (random error for plant f at age t)

Fixed effects regression—so named because it holds constant or “fixes” the average “effects” of each variable, as shown below— is well suited to analysis of panel data like ours, which consist of both cross-

sectional (i.e., variation in capacity factor “across” plants of the same age) and time series (i.e., variation in capacity factor “within” each plant as it ages) data. Because our interest here is solely in the time series or “within-plant” variation (A_t), we need to control for all cross-sectional or “across-plant” variation (S_f). We do this, in part, by using what we know about each plant’s (and site’s) characteristics to model the ideal capacity factor for each plant at each age and include it as an explanatory variable ($CF_{i,t}^{ideal}$). Even with the inclusion of $CF_{i,t}^{ideal}$ in Equation 2, however, there most likely remains some unobserved heterogeneity across plants and/or plant sites that we need to control for if we are to isolate the impact of age on performance. For this reason—i.e., the likelihood of omitted variables that are correlated with one or more of the explanatory variables included in the equation—ordinary least squares (OLS) regression will likely suffer from endogeneity problems in the form of omitted variable bias. Fortunately, fixed effects regression eliminates omitted variable bias, via the transformation of Equation 2 illustrated by Equations 3 and 4.

First, Equation 3 calculates the average over time for each variable in Equation 2. Because S_f does not vary over time in Equation 2, the average of S_f over time in Equation 3 is simply equal to S_f .

$$\overline{CF_f^{hist}} = \overline{CF_f^{ideal}} + S_f + \bar{A} + \bar{\epsilon}_f \quad (\text{Eq. 3})$$

Subtracting Equation 3 from Equation 2 yields Equation 4:

$$CF_{f,t}^{hist} - \overline{CF_f^{hist}} = (CF_{f,t}^{ideal} - \overline{CF_f^{ideal}}) + (S_f - S_f) + (A_t - \bar{A}) + (\epsilon_{f,t} - \bar{\epsilon}_f) \quad (\text{Eq. 4})$$

In Equation 4, the site-specific fixed effects (S_f) cancel, dropping out of the regression and leaving only those explanatory variables that vary with time. In other words, by subtracting the means, we eliminate all unobservable “across-plant” variation—a key source of omitted variable bias—and limit all variation to “within-plant” variation (i.e., which tells us how performance changes over time with age). As such, Equation 4 can now be solved without violating OLS constraints.

We emphasize that the age fixed effects (A_t) are applicable only to the entire sample of plants being analyzed, and are not specific to any individual plant. Adding the age fixed effects to the average historical capacity factor of the sample at age one results in an annual time series of average capacity factors for the sample as a whole, which we normalize by indexing the first year (i.e., age one) to 1.0. As a result of this model construct, the fixed effects model yields a single curve that illustrates the average impact of age on plant performance for the entire sample. Though this curve need not be linear, in practice it is approximately so; as a result, we take a best-fit line across the normalized curve, weighted by the number of plants at each age, to yield a single, linear average degradation rate, with confidence intervals, for the entire sample.

Figure 4 shows the results of the fixed effects regression. The solid blue circles show the average fixed effects at each age, with the shaded blue area representing the 95% confidence interval around each age fixed effect (with the uncertainty directly related to sample size, which declines with age). The

dashed line represents the regression slope of -1.2%/year, which is also weighted by sample size, such that higher ages with smaller sample (e.g., age 10-13) have very little influence on the slope. For example, dropping the fixed effect at age 13 (which, at least visually, appears to be an outlier) from the slope regression only changes the slope from -1.23%/year to -1.17%/year (given the very little weight that age 13 has in the regression, due to its small sample size of just two plants).




Figure 4. Age fixed effects and best-fit line for final model specification

Figure 5 replots the data from Figure 4, along with data from the original analysis in Bolinger et al. (2020), shown in orange. The updated -1.2%/year performance decline represents a slight improvement from the -1.3%/year found in the original analysis. Further exploration suggests that the slight improvement in slope is more likely attributable to the two additional years of operations than to the ~10 GW_{DC} of additional sample added with this update (e.g., analyzing the original sample from Bolinger et al. (2020) through 2020, rather than through 2018, also shows this improvement to -1.2%/year). Figure 5 visually supports this notion, particularly among older plants with higher ages (e.g., ages 6-10 in the original analysis, which correspond to ages 8-12 in the current analysis), which in most cases have clearly performed better on average through 2020 than they had through 2018. Though our data and model are not capable of explaining this apparent performance improvement among older plants, potential explanations could include: (1) a change of ownership after the 5-year recapture period for the federal investment tax credit (ITC)—or even just a post-ITC “flip” in revenue allocation to favor cash equity investors over tax equity investors—triggering a renewed emphasis on performance; or (2) it being easier to recognize—and then, more to the point, address—performance issues that persist over longer time periods.




Figure 5. Comparison of updated (blue) to original (orange) results

While the minor improvement from $-1.3\%/year$ in the original analysis (Bolinger et al. 2020) to $-1.2\%/year$ in this update is encouraging, the magnitude of the performance decline remains quite large—and significantly larger than most other studies have found. We emphasize once again that, due in large part to the low resolution of our data—i.e., annual, plant-level performance data—we are effectively measuring “total” performance decline with age, which includes both recoverable (with the exception of curtailment, for which we do control where we have the data to do so) and unrecoverable forms of degradation, and at the plant level rather than at the module or inverter level. As such, the magnitude of our numbers should be less surprising (and might also be slightly overstated, to the extent that we may not be fully capturing solar curtailment in those regions that do not report it).

While Figures 4 and 5 analyze the full sample of 631 plants, we also used the fixed effects model to analyze and compare various sub-samples, with the goal of identifying significant degradation pathways. Though we looked at numerous sub-sample comparisons, Table 3 lists only those that yielded statistically different degradation rates. The results are generally intuitive. For example, newer plants (those with CODs from 2015-2018) seem to have aged better than older plants—consistent with a story of technology improvement. Larger plants ($\geq 25 MW_{AC}$), which perhaps receive greater attention from asset managers, seem to have aged better than smaller plants. Plants with higher DC:AC ratios (≥ 1.25) seem to have aged better than their counterparts have—perhaps due to some amount of degradation on the DC side of the inverter being masked through more frequent power clipping. Combining these three variables yields the starkest contrast in Table 3: the performance of newer plants with greater capacities and higher DC:AC ratios—in other words, plants that look the most like those being built today—has declined by only about half as much ($-0.7\%/year$) as that of plants that fall outside of that sub-sample ($-1.4\%/year$). It is worth noting that the $-0.7\%/year$ performance decline among this sub-sample is more in line with other recent estimates from the literature (Deline et al., 2021; Jordan et al., 2021). Finally, the remaining sub-sample comparisons shown in Table 3 find that

fixed-tilt plants age better than plants that use single-axis tracking (perhaps due to fewer moving parts), plants using thin-film (mostly CdTe) modules have aged better than those using crystalline silicon modules, and plants located at sites with a lower irradiance and/or average temperature have aged better than their counterparts.

Table 3. Statistically significant comparisons of sub-samples using the fixed effects model

Variable	Less Degradation	More Degradation
Vintage	Post-2014 COD -1.1%/year, ±0.3%	Pre-2015 COD -1.2%/year, ±0.1%
Capacity	≥25 MW _{AC} -0.8%/year, ±0.2%	<25 MW _{AC} -1.4%/year, ±0.1%
DC:AC Ratio	≥1.25 DC:AC -1.1%/year, ±0.2%	<1.25 DC:AC -1.3%/year, ±0.2%
Combination (Vintage, Capacity, DC:AC)	Post-2014, ≥25 MW _{AC} , ≥1.25 DC:AC -0.7%/year, ±0.4%	Pre-2015, <25 MW _{AC} , <1.25 DC:AC -1.4%/year, ±0.2%
Mount	Fixed-Tilt -1.2%/year, ±0.2%	Single-Axis Tracking -1.3%/year, ±0.2%
Module Type	Thin-film (mostly CdTe) -1.0%/year, ±0.2%	x-Si -1.3%/year, ±0.1%
Solar Resource	GHI <210 W/m ² -1.0%/year, ±0.2%	GHI ≥210 W/m ² -1.2%/year, ±0.1%
Average Temperature at the Site	<15° C -0.9%/year, ±0.2%	≥15° C -1.3%/year, ±0.1%

Although these sub-sample comparisons yield interesting and intuitive results, and perhaps shed light on various degradation drivers, it is important to recognize that some of these variables are correlated. For example, newer plants also tend to have higher DC:AC ratios (as shown in Table 2), making it difficult to interpret results from the “Vintage” and “DC:AC Ratio” comparisons in Table 3. In an attempt to better isolate some of these more-correlated variables, we also ran a simple multivariate regression; the results were rather underwhelming, but do support the notion that vintage and solar resource, and to a lesser extent average site temperature (which is correlated with solar resource) and plant capacity, are among the more-significant drivers.

5. Discussion and Conclusions

The utility-scale solar market in the United States has grown and evolved rapidly in recent years. Yet the results of this updated analysis—which increases our plant sample size by ~50% (or ~10 GW_{DC} of capacity) and adds two additional years of operations, in 2019 and 2020—are largely consistent with those of our original analysis, published in Bolinger et al. (2020). Specifically, we find that solar plants have been modestly underperforming modeled expectations in their first year (similar to what we found earlier), and that average fleet-wide plant-level performance has declined by 1.2%/year—a slight improvement from the -1.3%/year found in Bolinger et al. (2020). This consistency in results across studies—particularly in the face of market expansion to new regions, evolving technology, increasing

solar curtailment, and growing impacts from wildfires and other natural disasters—bolsters our confidence in our approach and methods.

This update highlights the increasing importance of accounting for curtailment when assessing plant performance. Our original analysis, presented in Bolinger et al. (2020), found that controlling for curtailment improved the degradation rate by only about 0.1%/year (i.e., -1.3%/year controlled versus -1.4%/year uncontrolled), suggesting that curtailment was—at that time—a minor contributor to overall performance decline with age. In this updated analysis, however, the difference has grown to 0.4%/year (i.e., -1.2%/year controlled versus -1.6%/year uncontrolled), reflecting the significant growth in curtailment (particularly within CAISO; see Figure 2) in the two-year interim between studies. Had we not controlled for curtailment in either study, we would currently be reporting a slight worsening of total performance decline with age (i.e., -1.6%/year versus -1.4%/year in the original study) instead of a slight improvement (-1.2%/year versus -1.3%/year originally). And while it is possible that we are underestimating curtailment—and hence overstating degradation—in those regions that do not report curtailment data, the potential impact of this possibility is likely to be modest. For example, applying CAISO’s curtailment rate to plants in the Southwest (a non-ISO region that does not report solar curtailment data, but where curtailment likely occurs on occasion) only improves the overall sample-wide degradation rate to -1.1%/year (from -1.2%/year).

Despite the slight improvement from Bolinger et al. (2020), our revised -1.2%/year overall degradation rate is still large, and generally worse than what most other studies find. That said, the low temporal (i.e., annual) and spatial (i.e., plant-level) resolution of our generation data prohibits us from filtering out maintenance events and other downtime (other than curtailment, for which we do attempt to control where data exist, but may nevertheless be underestimating in some regions), which in turn means that we capture both recoverable and unrecoverable—i.e., total—degradation across the entire plant. When viewed through this lens, the magnitude of our average fleet-wide degradation rate is less surprising (and, again, may be slightly overstated if curtailment is occurring outside of CAISO and ERCOT), and our results do not appear to be as far away from others’. This is particularly the case when focusing on a sub-sample of newer and larger plants with higher DC:AC ratios—i.e., plants that more closely resemble what is being built today—which exhibit a more moderate sample-wide average performance decline of -0.7%/year ($\pm 0.4\%$).

While the analysis of various sub-samples finds that newer plants have aged better than older plants, the slight improvement in the overall fleet-wide average degradation rate across the total sample appears to be coming mostly from the oldest plants in our sample—e.g., those aged 6-10 years in the original study, or 8-12 years in the current study. This finding is important for several reasons. First, the utility-scale fleet in the United States is still relatively young, with an average plant age of less than 4 years, and there have been relatively few studies of plants as old as 8-12 years. The apparent turnaround in the performance of these older plants since our original analysis affirms the notion that we are detecting recoverable (in addition to unrecoverable) degradation that can, indeed, be recovered. It also potentially highlights the role of policy in driving plant performance, to the extent that the 5-year ITC recapture period is a factor (as hypothesized earlier). Finally, and perhaps most

importantly, this turnaround in older plants instills optimism that plant owners do see value in maintaining solid performance over the longer term, and will not simply leave older plants to wither on the vine.

6. References

- Bolinger, M., J. Seel, C. Warner, D. Robson. 2021. Utility-Scale Solar, 2021 Edition: Empirical Trends in Deployment, Technology, Cost, Performance, PPA Pricing, and Value in the United States. Berkeley, Calif.: Lawrence Berkeley National Laboratory. <https://utilityscalesolar.lbl.gov> (accessed 12/27/2021)
- Bolinger, M., W. Gorman, D. Millstein, D. Jordan. 2020. "System-level performance and degradation of 21 GWDC of utility-scale PV plants in the United States." *Journal of Renewable and Sustainable Energy* 12, 043501 (2020) <https://doi.org/10.1063/5.0004710>
- Deline, C., K. Anderson, D. Jordan, K. Perry, M. Muller, L. Simpson, J. Desai, A. Walker, R. White. 2021. "PV Fleet Performance Data Initiative: 2021 Update." Proceedings of the PV Reliability Workshop Virtual Conference, February 23, 2021. <https://www.nrel.gov/docs/fy21osti/80055.pdf> (accessed 12/27/2021)
- Jordan, D., K. Anderson, K. Perry, M. Muller, M. Deceglie, R. White, C. Deline. 2021 (submitted). "Photovoltaic Fleet Degradation Insights." *Progress in Photovoltaics: Research and Applications* (submitted in 2021).
- Keelin, P. and M. Perez. 2021. "2020 regional solar deviation summary and North American wildfire smoke impact report." Clean Power Research blog, 3/26/2021. <https://www.solaranywhere.com/2021/2020-solar-deviation-summary-na-wildfire-smoke-impact-report/> (accessed 12/27/2021)
- Kharait, R. and C. Chan. 2021. Addendum to the 2019 Solar Energy Assessment Validation for Utility-Scale Projects. DNV, 9/28/2021 <https://www.dnv.com/Publications/addendum-to-the-2019-solar-energy-assessment-validation-study-207695> (accessed 12/27/2021)
- kWh Analytics. 2021a. "Solar Generation Index: Quantitative Insights from Leading Asset Owners, 2021 Volume." <https://www.kwhanalytics.com/sgi> (accessed 12/27/2021)
- kWh Analytics. 2021b. "Solar Risk Assessment 2021: Quantitative Insights from the Industry Experts." <https://www.kwhanalytics.com/solar-risk-assessment> (accessed 12/27/2021)
- kWh Analytics. 2020. "Solar Risk Assessment 2020: Quantitative Insights from the Industry Experts." <https://www.kwhanalytics.com/solar-risk-assessment> (accessed 12/27/2021)
- National Renewable Energy Laboratory (NREL). 2021a. National Solar Radiation Database (NSRDB). <https://nsrdb.nrel.gov/> (accessed 12/27/2021), Golden, CO: National Renewable Energy Laboratory.
- National Renewable Energy Laboratory (NREL). 2021b. System Advisor Model (SAM). <https://sam.nrel.gov/> (accessed 12/27/2021), Golden, CO: National Renewable Energy Laboratory.