


# System-level performance and degradation of 21 GW<sub>DC</sub> of utility-scale PV plants in the United States

Cite as: J. Renewable Sustainable Energy 12, 043501 (2020); <https://doi.org/10.1063/5.0004710>  
Submitted: 13 February 2020 . Accepted: 03 June 2020 . Published Online: 07 July 2020

Mark Bolinger , Will Gorman , Dev Millstein , and Dirk Jordan

## COLLECTIONS

 This paper was selected as Featured



View Online



Export Citation



CrossMark

## ARTICLES YOU MAY BE INTERESTED IN

[Wind and photovoltaic potential in Europe in the context of mid-term energy storage](#)  
Journal of Renewable and Sustainable Energy 12, 034101 (2020); <https://doi.org/10.1063/1.5131560>


[Reliable long-term performance assessment of commercial photovoltaic modules tested under field conditions over 5 years](#)  
Journal of Renewable and Sustainable Energy 11, 063501 (2019); <https://doi.org/10.1063/1.5128171>

[Analysis of the cloud enhancement phenomenon and its effects on photovoltaic generators based on cloud speed sensor measurements](#)  
Journal of Renewable and Sustainable Energy 12, 043502 (2020); <https://doi.org/10.1063/5.0007550>



**NEW!**

Sign up for topic alerts  
New articles delivered to your inbox



# System-level performance and degradation of 21 GW<sub>DC</sub> of utility-scale PV plants in the United States <sup>F</sup>

Cite as: J. Renewable Sustainable Energy **12**, 043501 (2020); doi: 10.1063/5.0004710

Submitted: 13 February 2020 · Accepted: 3 June 2020 ·

Published Online: 7 July 2020



View Online



Export Citation



CrossMark

Mark Bolinger,<sup>1,a)</sup>  Will Gorman,<sup>1</sup>  Dev Millstein,<sup>1</sup>  and Dirk Jordan<sup>2</sup>

## AFFILIATIONS

<sup>1</sup>Lawrence Berkeley National Laboratory, One Cyclotron Road, Berkeley, California 94720

<sup>2</sup>National Renewable Energy Laboratory, 15013 Denver West Parkway, Golden, Colorado 80401

<sup>a)</sup> Author to whom correspondence should be addressed: [MABolinger@lbl.gov](mailto:MABolinger@lbl.gov)

## ABSTRACT

We assess the performance of a fleet of 411 utility-scale (i.e., >5 MW<sub>AC</sub> and ground-mounted) photovoltaic (PV) projects totaling 21.1 GW<sub>DC</sub> (16.3 GW<sub>AC</sub>) of capacity, which achieved commercial operations in the United States from 2007 to 2016. This fleet of projects contributed more than 50% of all solar electricity generated in the United States in 2017. Using detailed information on individual project characteristics, in conjunction with modeled irradiance data, we assess the extent to which actual first-year performance has lived up to both modeled and stated expectations. We then analyze system-level performance degradation in subsequent years by employing a “fixed effects” regression model to statistically isolate the impact of age on system performance. We find that this fleet of utility-scale PV projects has generally lived up to *ex ante* expectations for first-year performance but that subsequent system-level degradation—found to be  $-1.3\%/year$  ( $\pm 0.2\%$ )—has, on average, been worse than both *ex ante* expectations (commonly  $-0.5\%/year$ ) and results from past studies (ranging from  $-0.8\%/year$  to  $-1.0\%/year$ ). We emphasize that  $-1.3\%/year$  is a *system-level* estimate that captures more than just module degradation (e.g., including soiling, balance of plant degradation, and downtime for maintenance and/or other events). A side analysis of a variety of project characteristics suggests that system-level degradation rates tend to be of lower magnitude among newer projects and larger projects and at sites with lower long-term average temperatures.

© 2020 Author(s). All article content, except where otherwise noted, is licensed under a Creative Commons Attribution (CC BY) license (<http://creativecommons.org/licenses/by/4.0/>). <https://doi.org/10.1063/5.0004710>

## I. 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 projects—defined here to include any ground-mounted PV plant larger than 5 MW<sub>AC</sub>—achieved commercial operations in late 2007, followed by another plant in 2008 and three more in 2009. Thus, at the time of writing (in late 2019), the six oldest utility-scale PV projects in the U.S. fleet have only nine to eleven full calendar years of operating history—not a long track record for a technology expected to have a useful life of 30 years or longer. Moreover, the majority of the U.S. fleet has significantly less operational experience: among all utility-scale PV projects built in the U.S. from 2007 to 2018, the median, mean, and capacity-weighted average commercial operation dates (CODs) all fall within the year 2016, implying just a few years of operating history.

Though relatively young, the market for utility-scale PV in the United States has grown rapidly in recent years. From the humble beginnings described above, with just six projects (totaling 97 MW<sub>DC</sub>) built from 2007 to 2009, utility-scale PV became the largest sector (in terms of installed capacity) of the overall solar market by 2012 and remained so through 2018, with more than 33 GW<sub>DC</sub> operating in the United States at the end of that year. Based on data from the Energy Information Administration (EIA), more than half of all solar electricity generated in the United States in 2018 (i.e., across all three sectors—residential, commercial and industrial, and utility-scale—including concentrating solar thermal power) came from PV plants with capacities greater than 5 MW<sub>AC</sub> (U.S. Energy Information Administration, 2019a, 2019b). Analyst projections suggest that utility-scale’s market dominance will continue for at least another five years (Bolinger *et al.*, 2019).

With such a young fleet of utility-scale PV projects supplying the majority of solar generation in the United States and with the utility-scale sector projected to continue to dominate the market in the future, it is crucial to understand how these utility-scale projects have performed to date. Even at today's much-reduced costs (Bolinger *et al.*, 2019), each utility-scale PV project still requires millions of dollars of investment capital. As both the price and duration of power purchase agreements (PPAs) continue to decline, more of the return on that investment capital will shift into later, often post-PPA, years (Norton Rose Fulbright, 2019), thereby increasing the importance of predictable performance and solid long-term reliability. In addition, the planned phase-down of the federal investment tax credit (ITC) from 30% to 10% over the next few years will leave a greater share of investor capital at risk, further elevating the importance of performance and reliability. In order to appropriately price their capital, investors in utility-scale PV projects need to understand the degree of long-term performance risk that they face, while insurers and other financial intermediaries must have an even better understanding of the same risks in order to accurately price the "solar revenue puts" and other hedge-like products that they offer investors (Clarion Energy, 2019; kWh Analytics, 2019).

Yet, to date, most analyses of PV project performance have occurred among smaller, distributed PV systems—understandable given the much longer operating history of that market segment. Moreover, many of these studies have tended to focus primarily or exclusively on module-level performance and degradation, ignoring the potentially significant effect of "balance of plant" or "balance of system" components on overall system-level performance. Finally, some of these past studies have been conducted by specific module manufacturers, comparing their own modules to other types or by certain project owners looking at their own portfolio of projects, rather than taking a broader fleet-wide view. In light of this history and the increasing importance of the utility-scale PV sector in the United States, there is a clear need for more analysis of the performance of the entire fleet of large- or utility-scale PV projects operating under real-world conditions.

Through a variety of approaches, this paper assesses the system-level performance of a fleet of 411 utility-scale (i.e.,  $>5$  MW<sub>AC</sub> and ground-mounted) PV projects totaling 21.1 GW<sub>DC</sub> (16.3 GW<sub>AC</sub>) of capacity, which achieved commercial operations in the United States from 2007 to 2016 and, thus, have been operating for at least two (2017 and 2018) and as many as eleven (2008–2018) full calendar years. Using detailed information on individual project characteristics, in conjunction with modeled irradiance data, we assess the extent to which actual first-year performance has lived up to both modeled and stated expectations. We then analyze system-level degradation in energy output in subsequent years, by employing a fixed effects regression model to statistically isolate the impact of age on system performance. We find that this fleet of utility-scale PV projects has generally lived up to *ex ante* expectations for initial performance but that system-level degradation has, on average, been worse than both *ex ante* expectations and results from past degradation studies.

## II. SYSTEM-LEVEL VS MODULE-LEVEL ANALYSIS

Before proceeding, it is worth emphasizing that all the data and analysis presented herein are at the system-level, rather than at the module-level. This distinction is important, particularly with respect

to the analysis of performance degradation, where much of the existing research and literature have focused on module-level degradation. Moreover, as will be demonstrated later in Sec. VI, PV investors, developers, and power purchasers often seemingly confuse the two, assuming module-level performance degradation rates in instances where system-level degradation rates of greater magnitude would instead be more appropriate (e.g., within power purchase agreements). Thus, one contribution of this article is to draw this distinction more clearly, in order to inform the use of more-realistic inputs to financial models.

Module-level degradation stems from a wide array of degradation pathways, including but not limited to delamination, backsheet adhesion loss, junction box failure, frame breakage, cell cracks, potential induced degradation, high ambient and cell temperature, and cell hot spots (Köntges *et al.*, 2014; Jordan *et al.*, 2016b; Deceglie *et al.*, 2019). Different module types (e.g., x-Si vs CdTe vs copper indium gallium selenide (CIGS)) may experience a higher or lower prevalence of certain pathways (Deceglie *et al.*, 2019; Jordan *et al.*, 2016b; Kraus *et al.*, 2019), as might modules of similar type made by different manufacturers (Anderson *et al.*, 2013; Hasselbrink *et al.*, 2013). This combination of multiple possible degradation pathways that manifest differently in different module types and among different manufacturers has led to module-level degradation being widely studied and dominating much of the extant literature on PV performance degradation (relative to system-level degradation). For example, Jordan *et al.* (2016b) conducted a meta-analysis of more than 11 000 degradation rates reported in almost 200 studies across 40 different countries, and—at least among the 1537 reported rates for x-Si technology that the authors considered to be of "high quality"—75% were module-level, compared to just 25% that were system-level. Though there is, of course, a range, a general rule of thumb is that module-level degradation is typically on the order of  $-0.5\%/year$  (Jordan *et al.*, 2016b), and module manufacturers will often warrant that the output of their modules will not decline by more than  $-0.5\%/year$  (which may be one reason why financial models often assume  $-0.5\%/year$  degradation in output, even though a system-level degradation rate—likely of greater magnitude—would be more appropriate).

Though critically important to study (given that modules are the backbone of PV power plants), module-level degradation is just one component of system-level degradation, which also includes any and all other degradation pathways that could erode the output of the *entire system* over time. For example, in addition to module-level degradation, system-level degradation might stem from problems with the balance of system—including trackers, inverters, wiring, fuses, and breakers—or from the growth of vegetation that contributes to increased shading over time. It might also reflect plant curtailment imposed by grid operators due to a supply/demand imbalance and/or transmission limitations (though, as described later, we take steps to control for curtailment within our analysis).

To illustrate this critical difference between module- and system-level degradation, Fig. 1 presents a fabricated, stylized buildup of system-level degradation from a number of individual system components experiencing some combination of efficiency loss, discrete outages, and/or other degradation-related events. Starting at the top of the figure, the dark blue shaded area represents typical module degradation of  $-0.5\%/year$ . Next, the lighter blue shading represents inverter downtime, which cuts output from part or perhaps even the entire plant (the latter illustrated by the week-long, 100% reduction in output

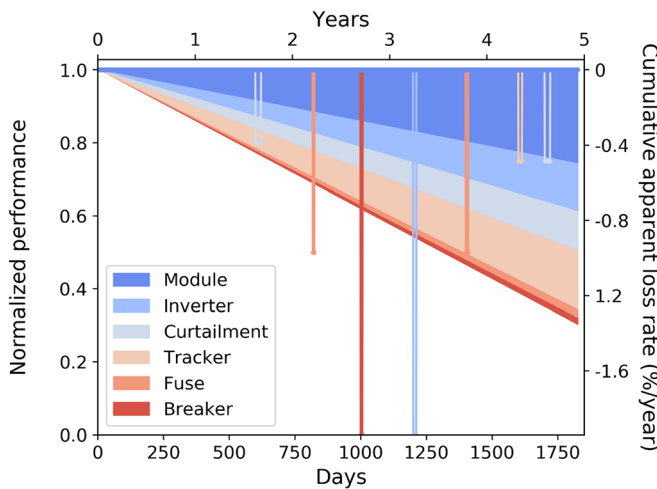


FIG. 1. Illustrative buildup from module- to system-level degradation.

shown by the light blue vertical line in year three). Next, in gray, is system curtailment, punctuated by two events that reduced system output to 70%–80% of normal for some period (vertical gray bars). Trackers might also occasionally get out-of-sync or stuck in sub-optimal positions, perhaps reducing output to ~80% of optimal until they can be fixed (pink area and vertical bar). In orange, fuses may fail, temporarily reducing the output of half of the system. Finally, in red, a breaker or transformer may go down, cutting the entire output of the plant for a week or more.

Though merely illustrative, and not intended to be exhaustive, the system-level degradation pathways shown in Fig. 1, along with the frequency and magnitude of each, are nevertheless loosely based on documented problems with PV systems that were funded under the U.S. Treasury’s Section 1603 cash payment program, as analyzed and reported in Jordan *et al.* (2020). As such, they provide a useful representation of the potential difference between module-level and system-level performance degradation over time—and also demonstrate how a typical module degradation rate of  $-0.5\%/year$  can potentially grow to more than  $-1.0\%/year$  once the entire system is considered.

Indeed, the comparatively few studies that have examined system-level rather than (or in addition to) module-level degradation have typically found degradation rates that exceed the commonly cited  $-0.5\%/year$  for module-level degradation. For example, the compendium of degradation rates assembled by Jordan *et al.* (2016b) found that x-Si module technology—which, as described later, is used in 80% of our sample of utility-scale projects in the United States (Table II)—exhibited an average degradation rate of  $-0.8\%/year$  to  $-0.9\%/year$  (with a median of  $-0.5\%/year$  to  $-0.6\%/year$ ). Module-level studies

generally fell at the lower/milder end of that range, while system-level studies were generally at the upper/worse end—whether due to additional degradation in the balance of plant or the knock-on effect of individual module degradation impacting entire strings (Jordan *et al.*, 2016b). More recently, Deceglie *et al.* (2019) analyzed more than 500 PV systems in the United States and found an average system-level degradation rate of  $-0.8\%$  among the non-residential systems (including some utility-scale projects) in their sample, while degradation among the residential sub-sample was worse, at  $-1.3\%/year$ . However, among the non-residential systems employing x-Si modules (again, 80% of the projects in our sample use Si modules), the average rate was  $-1.0\%/year$  (Deceglie *et al.*, 2019).

While module-level degradation is clearly important, it is overall system-level performance degradation that ultimately affects the financial performance of PV plants. Module-level degradation is, of course, a significant component of system-level degradation, but it does not tell the whole story. As we will demonstrate, investors who assume that the output of PV plants within their portfolio will degrade at rates that are based on module-level studies are potentially in for an unpleasant surprise as time passes.

### III. DATA SAMPLE

The sample of utility-scale PV plants that we analyze consists of 411 plants totaling 21.1  $GW_{DC}$  (16.3  $GW_{AC}$ ) installed across 28 states from 2007 to 2016 (Tables I and II, Fig. 2). In aggregate, these plants contributed  $>50\%$  of all solar electricity generated in the United States in 2017 (across all sectors—residential, commercial, and utility-scale—and including concentrating solar thermal power) and 40% of all solar electricity generated in 2018. They collectively offer 1536 project-years of operational experience, more than a third of which are in California (Table I). Operational history ranges from 2 to 11 full calendar years, with an average of 3.7 years—once again indicative of the relative youth of the utility-scale PV sector (Table II).

A histogram of projects within our sample by capacity (Fig. 2) shows the majority falling into the 20–50  $MW_{DC}$  capacity bin. Nearly 85% of projects are 100  $MW_{DC}$  or less, but a number of projects feature several hundred  $MW_{DC}$  of capacity, with the largest being nearly 760  $MW_{DC}$ .

We normalize the performance of each individual plant in the sample by calculating its “capacity factor,” which expresses the quantity of electricity generated by each plant over a certain period relative to the quantity that could have been generated if that plant were operating at full capacity over that entire period. Due to the seasonal nature of PV generation, we only calculate capacity factors over full-year periods. Furthermore, due to certain data quality issues—i.e., the fact that monthly generation profiles for some plants differ across different sources, even though calendar-year totals match—we limit the full-year periods over which we calculate capacity factors to calendar years. Hence, the capacity factor (CF) equation is

$$CF_y = \frac{\text{MWh generated in calendar year } y}{(\text{MW}_{DC} \text{ of capacity in calendar year } y \times \text{number of hours in calendar year } y)} \quad (1)$$

In order to calculate empirical capacity factors (and to simulate “ideal” capacity factors—i.e., controlling for weather and other variables), we compile the following data for each plant in our sample:

TABLE I. Geographic descriptive statistics of the sample.

States	No. of Projects	No. of MW <sub>DC</sub>	No. of MW <sub>AC</sub>	MW <sub>AC</sub> /project		No. of project-years	Years per project			% of the US sample		
				Average	Median		Min	Avg	Max	Projects	MW <sub>AC</sub>	Project-Years
CA	150	10 133	7958	53	20	539	2	3.6	9	36%	49%	35%
AZ	32	1915	1431	45	19	151	2	4.7	7	8%	9%	10%
NM	25	579	462	18	10	118	2	4.7	8	6%	3%	8%
NJ	28	297	243	9	8	113	2	4.0	7	7%	1%	7%
NC	31	1174	883	28	20	94	2	3.0	6	8%	5%	6%
NV	17	1841	1383	81	48	79	2	4.6	11	4%	8%	5%
TX	14	716	569	41	22	64	2	4.6	8	3%	3%	4%
GA	23	1240	953	41	30	61	2	2.7	5	6%	6%	4%
CO	12	420	338	28	16	50	2	4.2	11	3%	2%	3%
FL	8	436	283	35	19	41	2	5.1	9	2%	2%	3%
IN	11	116	87	8	9	40	2	3.6	5	3%	1%	3%
MD	7	109	84	12	13	27	2	3.9	6	2%	1%	2%
UT	12	1049	810	68	80	26	2	2.2	3	3%	5%	2%
MN	9	299	208	23	7	18	2	2.0	2	2%	1%	1%
IL	2	33	28	14	14	15	6	7.5	9	0%	0%	1%
OH	2	22	18	9	9	15	7	7.5	8	0%	0%	1%
OR	7	83	63	9	10	14	2	2.0	2	2%	0%	1%
TN	3	50	39	13	16	14	4	4.7	6	1%	0%	1%
DE	2	26	22	11	11	13	6	6.5	7	0%	0%	1%
MA	3	44	33	11	12	8	2	2.7	4	1%	0%	1%
VA	4	180	136	34	20	8	2	2.0	2	1%	1%	1%
NY	1	38	32	32	32	7	7	7.0	7	0%	0%	0%
PA	1	11	10	10	10	6	6	6.0	6	0%	0%	0%
SC	2	24	17	8	8	4	2	2.0	2	0%	0%	0%
ID	2	163	120	60	60	4	2	2.0	2	0%	1%	0%
AR	1	17	13	13	13	3	3	3.0	3	0%	0%	0%
AL	1	100	75	75	75	2	2	2.0	2	0%	0%	0%
KY	1	14	10	10	10	2	2	2.0	2	0%	0%	0%
Total	411	21 130	16 308	40	20	1536	2	3.73	11	100%	100%	100%

• Key plant characteristics are used to calculate actual historical, and to simulate ideal, capacity factors from 2008 to 2018 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<sub>DC</sub> and MW<sub>AC</sub>), and DC:AC ratio. These system-level characteristics are sourced from Bolinger *et al.* (2019) and, where possible, have been confirmed with project owners and/or visually verified through satellite imagery. Table II summarizes

TABLE II. Temporal descriptive statistics of the sample.

COD year:	2007	2008	2009	2010	2011	2012	2013	2014	2015	2016	Total
No. of Projects	2	1	3	7	31	36	49	50	83	149	411
No. of MW <sub>DC</sub>	22	12	63	172	532	1133	2249	3564	3596	9785	21 130
No. of MW <sub>AC</sub>	19	10	54	144	448	886	1761	2726	2764	7496	16 308
Average DC:AC ratio	1.18	1.21	1.16	1.19	1.19	1.28	1.28	1.31	1.30	1.31	1.30
Average MW <sub>AC</sub> /project	9	10	18	21	14	25	36	55	33	50	40
% Projects with tracking	100%	0%	67%	14%	52%	50%	57%	62%	66%	74%	64%
% Projects with Si modules	100%	0%	67%	29%	74%	83%	82%	74%	76%	88%	80%

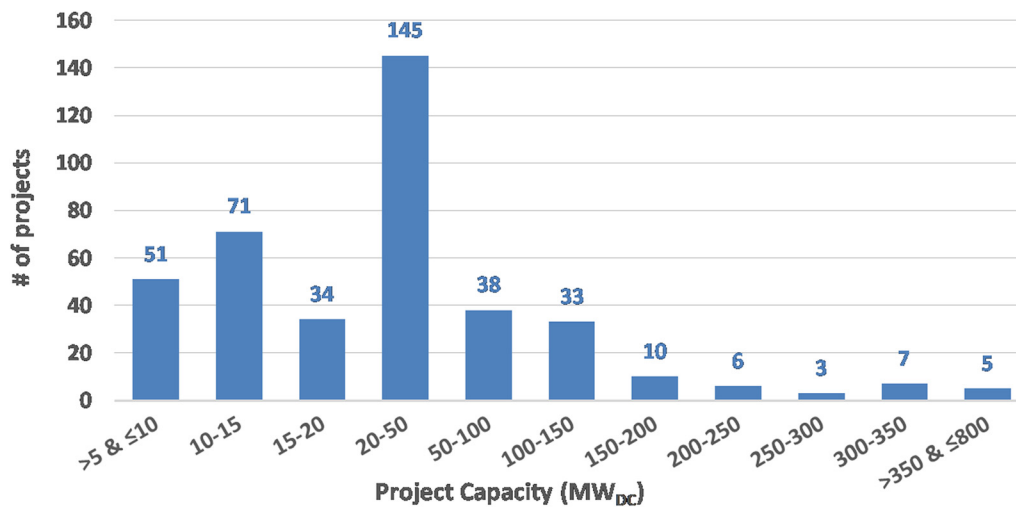


FIG. 2. Histogram of individual project capacity within the sample.

some of these characteristics across the sample over time. All projects in our sample use monofacial modules (no bifacial), and as of the end of 2018, only two had been paired with battery storage (with batteries added in early 2018 in both instances, thus minimizing the possibility that battery degradation could have a meaningful influence on overall system-level degradation).

- Net generation data for each project are used to calculate actual historical capacity factors and are compiled from a variety of sources, each offering different temporal resolution and plant coverage: Form EIA-923 (monthly), Federal Energy Regulatory Commission (FERC) Form 1 (annual), FERC Electric Quarterly Reports (quarterly, monthly, daily, hourly, or even sub-hourly, depending on the filer), and the California Energy Commission (annual). We generally default to using Form EIA-923 data—due to its nationwide coverage of all plants  $>1$  MW<sub>AC</sub>—but roll up the reported monthly generation data to full calendar years, in light of the problems described earlier with the monthly data in some cases [a finding that others have commented upon as well (DNV GL, 2019)]. We do, however, crosscheck the EIA data against the other sources listed above and, as necessary, substitute calendar-year data from one of those other sources in instances where the EIA calendar-year net generation data appear suspect.
- Irradiance data for each site location (based on project coordinates) are used to simulate the ideal capacity factor. Data for the years 2008–2018 come from the National Solar Radiation Database (NSRDB), which uses National Renewable Energy Laboratory’s (NREL) Physical Solar Model to provide solar radiation and meteorological data at 4-km horizontal resolution across 30-min intervals (NREL, 2019a).
- Hourly solar curtailment data are sourced from the California Independent System Operator (CAISO) and the Electric Reliability Council of Texas (ERCOT), are described in Bolinger *et al.* (2019), and are used to gross up the historical capacity factors of plants that have been curtailed in California and Texas. The other five independent system operators (ISOs) across the

United States do not yet report solar curtailment data (though all seven ISOs do report data on wind power curtailment, perhaps suggesting that solar curtailment has not yet risen to meaningful levels in these other ISOs). There are also a number of plants in the sample, that are located outside of ISO regions—we do not have curtailment data for these plants, either. Thus, other than in California and Texas, where we have data showing that curtailment of solar plants has occurred, we assume no curtailment. The implications of this simplifying assumption are likely minimal, given that—as shown later—whether or not we account for curtailment in California and Texas makes little difference in the implied sample-wide system degradation rate. Furthermore, even if it occurs, solar curtailment in other markets outside of California and Texas is likely to be relatively minor, given only modest market penetrations.

#### IV. ASSESSMENT OF FIRST-YEAR SYSTEM-LEVEL PERFORMANCE

Based on the data described in Sec. III, Fig. 3 plots first-year capacity factors from projects in our sample against two different representations of expected capacity factors: one simulated and the other empirical. The graph on the left [Fig. 3(a)] compares actual first-year capacity factors to simulated ideal capacity factors. Because, at least for this purpose, we are interested in comparing actual to expected capacity factors in a typical year, we normalize both the actual and modeled data by correcting for interannual variance in solar irradiance. This involves dividing the actual first-year capacity factor for each plant by a “solar index,” which is simply a ratio of the irradiance at that site in that particular year relative to the long-term average irradiance at the same site from 2007 to 2018 [irradiance values for each site over time come from NSRDB (NREL, 2019a)]. Similarly, for the modeled capacity factors, we use “typical meteorological year” (TMY) irradiance data rather than data for a specific year.

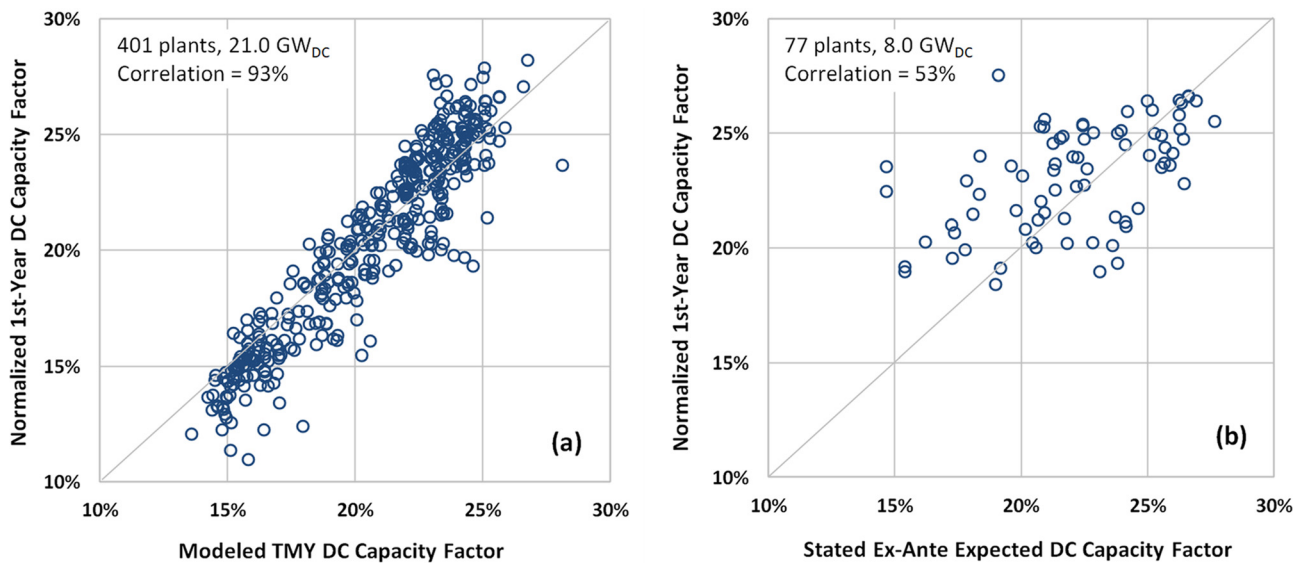


FIG. 3. Normalized first-year capacity factor vs (a) modeled and (b) stated expectations.

This normalized modeling approach allows us to approximate *ex ante* capacity factor expectations for almost our entire sample—and the correlation is strong (93%). Though clearly some projects exceed while others fall short of the unity line, in aggregate, the actual first-year capacity factors “underperform” the TMY-based simulations by 0.6% (median), 1.2% (simple mean), or 1.3% (capacity-weighted average), while 89% of projects generated more than 90% of their simulated TMY estimate in their first full calendar year of operations. This degree of shortfall is comparable to other estimates. For example, among a sample of 39 utility-scale projects totaling 1.2 GW, [DNV GL \(2019\)](#) found that actual performance lagged preconstruction estimates by 3.1%, with this performance gap declining to 1.7% if they discarded the first full year (rather than just the first full month) of production. Meanwhile, in successive analyses of PV systems supported by the Section 1603 cash payment program in the United States, [Jordan and Kurtz \(2015\)](#) found that about 90% of systems without reported issues generated more than 90% of their preconstruction estimates, while [Jordan et al. \(2020\)](#) found that 80%–90% of an expanded sample performed within 10% of expected output. Some of this “underperformance” could simply reflect the +3% to +5% global horizontal irradiance (GHI) bias in NSRDB identified by [Hansen et al. \(2015\)](#).

The graph on the right [Fig. 3(b)] looks at a smaller sub-sample of projects for which we have collected *ex ante* capacity factor expectations as published in power purchase agreements. The correlation in this case is significantly weaker (53%), but the errors seem to be mostly random, though slightly skewed toward actual first-year capacity factors outperforming stated expectations. Indeed, for this sub-sample of 77 projects totaling 8.0 GW<sub>DC</sub>, actual first-year performance outperforms the stated expectations by 3.4% (median), 5.8% (capacity-weighted average), or 6.3% (simple mean), while 88% of projects generated more than 90% of their stated expectations in their first full calendar year of operations. It is not particularly surprising to see

actual first-year capacity factors outperform expectations as stated in PPAs, given that these stated expectations typically serve as the benchmark for contractual performance guarantees, presumably causing developers to err on the conservative side when establishing them. On the other hand, most PPAs also penalize significant over-delivery by setting a lower price for each MWh delivered beyond a certain threshold (e.g., over 120% of expected annual generation), which serves to limit the degree of conservatism in setting performance expectations.

Together, [Figs. 3\(a\) and 3\(b\)](#) demonstrate that, across our sample, actual first-year capacity factors—which range widely from 11% to 28% (in DC terms)—have generally matched expectations fairly well (at least in the case of the TMY simulations) and without significant bias (in both cases).

## V. ASSESSMENT OF SYSTEM-LEVEL PERFORMANCE DEGRADATION

Having established that first-year performance across our sample has generally not strayed too far from expectations, we now turn to assessing system-level degradation in subsequent years. Following a description of our approach and methodology, we present the resulting average sample-wide system degradation rate and findings from a side-analysis of potential degradation drivers.

### A. Degradation methodology

As noted earlier, concerns over the quality of monthly generation data (i.e., different data sources show different monthly generation profiles for the same plant, even though their calendar year totals match) have forced us to resort to analyzing annual generation data, by calendar year. This, in turn, limits the approaches available to measure system-level degradation. For example, with annual data, there are simply not enough data points to successfully employ a year-over-year approach, as described by [Anderson et al. \(2013\)](#), [Hasselbrink](#)

*et al.* (2013), and *Jordan et al.* (2016a) and (2018), and which forms the basis for NREL’s open-source RdTools platform (NREL, 2019d).

Instead, we adopt a regression-based approach using a fixed effects model. Fixed effects regression is well suited to the analysis of panel data like ours, which consist of both cross-sectional (i.e., variation in capacity factor “across” plants in each time period) and time series (i.e., variation in the capacity factor “within” each plant over time) data. Because our interest here is solely in the time series or the “within-plant” variation, we need to control for all the cross-sectional or “across-plant” variation. 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 in each time period and include it as an explanatory variable [ $CF_{f,t}^{ideal}$  in Eq. (2)]. Even with the inclusion of  $CF_{f,t}^{ideal}$  in Eq. (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 illustrated in Eqs. (2)–(4). To our knowledge, fixed effects models have not previously been employed in prior studies of PV performance degradation but have been used in several recent studies of the long-term performance of wind power plants in various countries (Hamilton *et al.*, 2020; Germer and Kleidon, 2019; Olauson *et al.*, 2017; Staffell and Green, 2014; Hughes, 2012). Using annual data, this approach regresses actual historical capacity factors on ideal capacity factors, along with fixed effects for each plant,  $S_f$ , and for each whole year of plant age,  $A_t$ , according to the following model:

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

where  $CF_{f,t}^{hist}$  is the historical capacity factor of plant  $f$  at age  $t$  ( $t$  in whole calendar years),  $\beta$  is the coefficient of performance,  $CF_{f,t}^{ideal}$  is the ideal capacity factor of plant  $f$  at age  $t$ ,  $S_f$  is the site-specific fixed effects for plant  $f$ ,  $A_t$  is the age fixed effects at age  $t$ , and  $\epsilon_{f,t}$  is the random error for plant  $f$  at age  $t$ . Equation (2) is known as a fixed effects regression because it holds constant or “fixes” the average “effects” of each variable. We can illustrate this through two transformations of Eq. (2). First, Eq. (3) calculates the average over time for each variable in Eq. (2). Because  $S_f$  does not vary over time in Eq. (2), the average of  $S_f$  over time in Eq. (3) is simply equal to  $S_f$ .

$$\overline{CF_{f,t}^{hist}} = \beta \overline{CF_{f,t}^{ideal}} + S_f + \bar{A} + \bar{\epsilon}_f. \quad (3)$$

Subtracting Eq. (3) from Eq. (2) yields the following equation:

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

In Eq. (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 limit all variations to the within-plant variation and eliminate all unobservable across-plant variations—a key source of omitted variable bias. As such, Eq. (4) can now be solved without violating OLS constraints.

In populating the variables of Eq. (2) (i.e., the un-transformed version of the fixed effects regression, which is the version we implement), we calculate historical capacity factors ( $CF_{f,t}^{hist}$ ) by dividing the annual net generation data for each plant in each calendar year by the product of that plant’s capacity (in DC terms) and the number of hours in each calendar year, per Eq. (1). However, for projects located in Texas and California, we first gross up the annual net generation data as necessary to account for curtailment (in an attempt to control for curtailment, rather than allowing it to contribute to our estimate of system degradation). In Texas, this adjustment is straightforward, as we have hourly plant-level curtailment data over the full operating history of each plant and so know exactly which plants have been curtailed, when, and by how much. In California, however, we only have system-wide (as opposed to plant-specific) solar curtailment data and only back through 2015, requiring the development of a method to extrapolate system-wide curtailment back further in time and to allocate it across individual plants.

We extrapolate curtailment to years earlier than 2015 by applying 2015’s ratio between solar curtailment and installed solar capacity to installed solar capacity in those earlier years. Given that CAISO solar curtailment in 2015 was already modest (just 0.7%), this simple extrapolation yields low levels of curtailment prior to 2015 (and almost none prior to 2012). The first step in the plant-level allocation process is to associate each utility-scale PV plant in CAISO with a nearby pricing node for which we have historical locational marginal prices (LMPs) on an hourly basis. We then use an iterative process to allocate known system-wide solar curtailment across those solar plants facing the lowest LMPs, on the theory that these lowest-value projects will be the first to be curtailed. This process begins by focusing on projects that face negative LMPs; if those projects are unable to absorb the full amount of known system-wide curtailment, we then focus on projects facing very low positive LMPs (e.g., \$0–\$5/MWh) to allocate the remainder, progressively ratcheting up the LMP threshold as needed until all system-wide curtailment has been allocated to specific projects.

Finally, we establish the age,  $t$ , of each project in each calendar year by defining the first full calendar year following the project’s commercial operation date (COD) as “age one.” This effectively means that we are disregarding anywhere from zero to twelve months of initial operations (a period that we refer to as “age zero”), depending on whether the project achieved commercial operations late or early in the calendar year, respectively. On average, projects within our sample achieved commercial operations in September, which suggests that we exclude 3–4 months of age zero data on average. Although discarding up to a year of a project’s initial output is not ideal—particularly given the short track record of many of the projects in our sample—this data sacrifice is necessary due to our focus on calendar year data, which itself stems from inconsistencies in the quality of monthly generation data for some projects, as noted earlier. In addition, excluding this age zero period provides somewhat of a buffer against other potential data quality problems (e.g., the COD actually occurring a bit later than reported) and normal “teething issues” (e.g., power quality or interconnection issues that typically get ironed out in the first few months following the COD) inadvertently biasing our analysis.

Whereas our historical capacity factors are empirical, our ideal capacity factors [ $CF_{f,t}^{ideal}$  in Eq. (2)] are simulated and control for both the interannual variation in the solar resource (at each plant’s



coordinates) and known plant-specific characteristics (e.g., capacity, tracking vs fixed-tilt mount, tilt, azimuth, and DC:AC ratio). Specifically, drawing upon our extensive database of project characteristics described earlier, in conjunction with site-specific irradiance data from the NSRDB, we simulate hourly solar generation for each project using NREL’s System Advisory Model (NREL, 2019b), which, in turn, relies upon performance algorithms from NREL’s PVWatts (NREL, 2019c). Although Hansen *et al.* (2015) found that NSRDB overstates global horizontal irradiance (GHI) by 3%–5%, for present purposes, we are less-concerned about absolute bias and more-interested in NSRDB’s ability to capture year-to-year variability, which Habte *et al.* (2017) described as reasonably good. Modeling inputs for which we use known characteristics that vary by plant include capacity, DC:AC ratio, tilt, azimuth, module type, and mount type; all other parameters are left at prepopulated default values. We roll up the estimated hourly generation to an annual timescale and then calculate capacity factors as described above (without any adjustments for curtailment in this case).

The site-specific fixed effects [ $S_i$  in Eq. (2)] control for all remaining (but unknown) differences across projects and sites that are not already incorporated into and reflected by the ideal capacity factor. In practice, these site-specific fixed effects are implemented as a dummy variable and are expressed as an absolute deviation from a reference plant.

Finally, the age fixed effects [ $A_t$  in Eq. (2)] capture the average influence of age on the capacity factor in each year, implemented as a dummy variable and expressed as an absolute deviation from the average historical capacity factor at age one (i.e., the reference age). Adding the age fixed effects to the average historical capacity factor 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.

If not already clear from the preceding descriptions of the variables in Eq. (2), it is worth emphasizing 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. Differences between individual plants are, instead, captured by the simulated ideal capacity factor ( $CF_{fit}^{ideal}$ ) and the site-specific fixed effects ( $S_i$ ). The more-accurate the simulated ideal capacity factor, the fewer the remaining differences across plants that need to be explained by the site-specific fixed effects. 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 observations (plants) at each age, to yield a single, linear average degradation rate, with confidence intervals, for the entire sample.

**B. Degradation results**

Figure 4 plots the age fixed effects (i.e., the indexed capacity factor) for the full sample (blue circles), bounded by 95% confidence intervals (blue shaded area), along with a best-fit line (blue dashed line) that is weighted by the number of observations (i.e., plants) in each year. Over the full eleven-year period, the slope of the best-fit line is highly significant, at  $-1.3\%/year$  ( $\pm 0.2\%/year$  with 95% confidence). At first glance, the  $\pm 0.2\%$  uncertainty around the best-fit line might seem too narrow, in light of the high degree of uncertainty surrounding the individual age fixed effects for ages 8 through 11 (the blue-shaded area widens considerably as sample size dwindles). We reiterate, however, that the best-fit line is weighted by the number of plants for each age. As such, ages 8 through 11 receive very little weight, enabling a relatively tight confidence interval around the best-fit line (in fact, a best-fit line spanning just ages 1–7 yields an almost identical slope as for the full eleven-year period). To further check for robustness, we re-ran the model on just the central 80% of plants (in terms of each individual plant’s degradation rate)—i.e., excluding the top and bottom 10%—and found similar results as for the full sample, suggesting that our results are relatively insensitive to possible outliers. The lack of apparent degradation at age 2 is perhaps an indication that




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

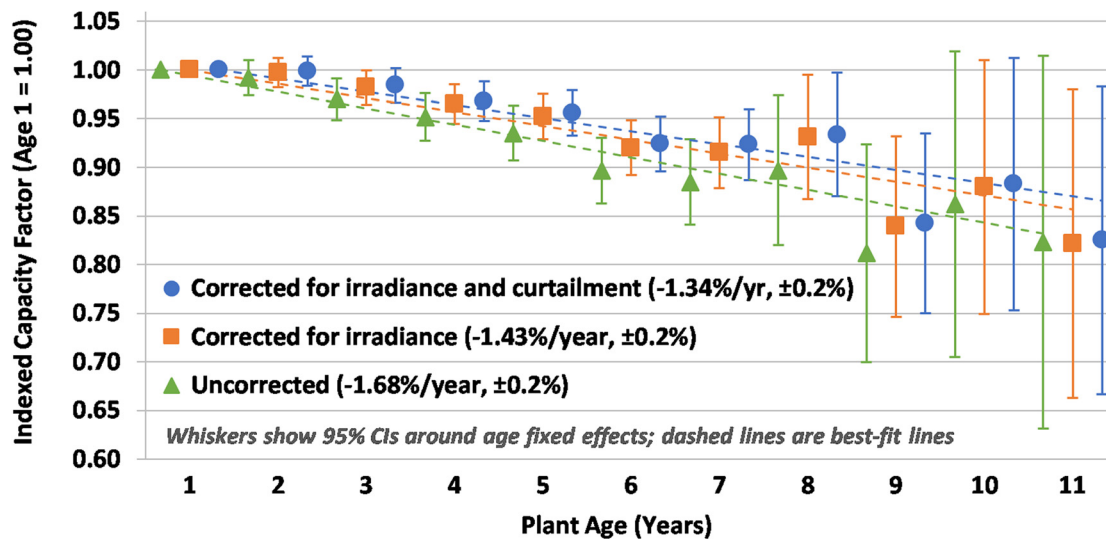


FIG. 5. Age fixed effects and best-fit lines for three different model specifications.

we are not fully excluding the initial “teething period” for all projects (despite defining age 1 as the first full calendar year following COD). If we exclude the first year (age 1) due to these potential ongoing teething issues and focus just on ages 2–7, the slope steepens somewhat to  $-1.6\%/year$  (orange dashed line).

To reiterate, Fig. 4 shows the output of our final fixed effects model specification—i.e., after correcting all projects for the interannual variation in irradiance and some projects in California and Texas for curtailment. Figure 5, meanwhile, shows the incremental effect of these two corrections—without which the implied degradation rate would be worse than shown in Fig. 4. Specifically, if we do not adjust the historical capacity factors in Eq. (2) for curtailment in California and Texas and remove the ideal capacity factor from the right-hand side of that equation, the slope from ages 1 to 11 is considerably worse, at  $-1.7\%/year$  (see the green triangles, 95% confidence intervals, and best-fit line in Fig. 5). Adding the ideal capacity factor to the right-hand side of the equation improves the slope to  $-1.4\%/year$  (orange squares), and then controlling for curtailment in California in Texas as well results in a modest further improvement to the  $-1.3\%/year$  ( $\pm 0.2\%$ ) shown in both Figs. 4 and 5 (blue circles). Not surprisingly, the irradiance correction—which affects all projects in the sample—has a larger impact than the curtailment correction, which only affects certain projects in California and Texas.

### C. Degradation drivers

The low temporal (i.e., annual) and spatial (i.e., project-level) resolution of our generation data prohibits identification of specific degradation pathways, but we did nevertheless look for statistically significant relationships between a range of project-level characteristics and degradation rates. The specific project characteristics that we tested include the commercial operation date, project capacity, DC:AC ratio, mount type (fixed-tilt, single-axis tracking, and dual-axis tracking), module manufacturer (SunPower, First Solar, and others), and

the long-term average irradiance and temperature at each project site (sourced from NSRDB). We analyzed these potential drivers in two ways.

First, for many of these project characteristics, we divided the full project sample into two sub-samples (e.g., for the commercial operation date, we divided the full sample into projects built pre-2013 vs post-2012) and then ran the fixed effects model on each sub-sample, assessing whether the slopes of the “age fixed effects” for each sub-sample, as well as the *difference between the slopes* of each sub-sample, are statistically significant. This exercise yielded five pairings of statistical significance:

- (1) post-2012 projects ( $-0.9\%/year$ ,  $\pm 0.3\%$ ) degrade less than pre-2013 projects ( $-1.5\%/year$ ,  $\pm 0.2\%$ );
- (2) projects of 25 MW<sub>AC</sub> or larger ( $-0.8\%/year$ ,  $\pm 0.3\%$ ) degrade less than projects less than 25 MW<sub>AC</sub> ( $-1.5\%/year$ ,  $\pm 0.2\%$ );
- (3) projects with an average site temperature of less than 15 °C ( $-0.9\%/year$ ,  $\pm 0.4\%$ ) degrade less than projects at warmer sites ( $-1.4\%/year$ ,  $\pm 0.2\%$ );
- (4) projects at sites with an average long-term global horizontal irradiance of less than 210 W/m<sup>2</sup> ( $-0.8\%/year$ ,  $\pm 0.3\%$ ) degrade less than projects at sunnier sites ( $-1.4\%/year$ ,  $\pm 0.2\%$ ); and
- (5) projects with a DC:AC ratio of 1.25 or greater ( $-1.2\%/year$ ,  $\pm 0.2\%$ ) degrade less than projects with lower ratios ( $-1.4\%/year$ ,  $\pm 0.2\%$ ).

Although the specific thresholds chosen for each variable (2013, 25 MW<sub>AC</sub>, 15°, 210 W/m<sup>2</sup>, and 1.25 DC:AC ratio) are somewhat arbitrary, these findings are intuitive. For example, photovoltaic technology should improve over time, larger projects presumably receive more attention than smaller projects in terms of maintenance and repair protocols, and higher temperatures and irradiance have been found to be positively correlated with performance degradation. In addition, higher DC:AC ratios will tend to mask some system-level performance degradation due to power clipping (i.e., if degraded DC

power output is further reduced or “clipped” by an inverter operating at its maximum capacity, then that degraded DC output will not negatively impact system-level performance).

The bivariate nature of the fixed effects sub-sample analysis described in the preceding paragraph, however, does not allow one to control for variables other than the single variable being directly examined (e.g., commercial operation date, project size, average site temperature and irradiance, and DC:AC ratio). To protect against the possibility of other variables driving the results (e.g., if newer projects are also generally larger and have a higher DC:AC ratio, then it is difficult to interpret the fixed effects results described in the previous paragraph), we also ran a multivariate regression on the same set of project characteristics. Two variables stood out as highly significant: commercial operation date and long-term average site temperature (long-term average site irradiance was also significant, but only at the 10% level, and it is highly correlated with the more-significant long-term average site temperature). Newer projects were found to degrade less, while hotter sites were found to degrade more; once again, both results are intuitive and are consistent with the fixed effects results in the preceding paragraph.

VI. DISCUSSION

Although the first-year performance of our sample seems roughly in line with expectations [particularly as represented by TMY-based modeling—see Fig. 3(a)], the  $-1.3\%/year$  average system-level degradation rate implied by Fig. 4 is generally larger in magnitude than found in previous studies. It is also significantly worse than *ex ante* expectations, as revealed by expected degradation rates published within power purchase agreements. Figure 6 shows the distribution of expected degradation rates found within a sample of PPAs for 93 utility-scale PV plants totaling  $7.36\text{ GW}_{AC}$ . Nearly half of these PPAs codify an expected degradation rate of  $-0.5\%/year$ , while nearly two-thirds expect either  $-0.5\%/year$  or better. Only one PPA in the sample expects degradation to be worse than  $-1.0\%/year$ .

While the significant misalignment of these expected degradation rates with our findings of  $-1.3\%/year$  is alarming, it is, once again, worth considering the source of these expectations—power purchase agreements—and whether that might have any bearing on the discrepancy. If, as discussed earlier, developers are somewhat conservative when establishing first-year generation expectations, then they can afford to be less-conservative when setting expectations about degradation—this could explain some of the discrepancy between our findings in Fig. 4 and the expectations shown in Fig. 6. In other words, performance guarantees within PPAs are based on an assessment of actual vs expected generation over time, with those expectations dependent on both the first-year generation estimate and the expected degradation rate. If the first-year generation estimate is conservative enough, then one could potentially assume no degradation at all and still meet contractual performance requirements. Of course, another possible explanation is that the PPA counterparties are simply confused about module- vs system-level degradation.

It is possible that the discrepancy between our  $-1.3\%/year$  system-level degradation estimate and the somewhat milder degradation rates found in earlier studies and memorialized in PPAs is methodological and/or related to data resolution. In particular, though a necessity due to the monthly data problems described earlier, our reliance on annual generation data could nevertheless be problematic, given that such low resolution restricts visibility into potentially important events occurring over finer timescales—e.g., seasonal soiling, maintenance events, and other downtime—and also complicates filtering of anomalous data. Such soiling and availability-related events can lead to apparent higher-than-typical degradation rates (Jordan et al., 2020). That said, our modeling of ideal capacity factors does occur on an hourly timescale and—once rolled up to the annual level—matches actual performance reasonably well [e.g., see Fig. 3(a), earlier].

In an attempt to validate our annual fixed effects regression model and its annual resolution, we analyzed a sub-sample of 16

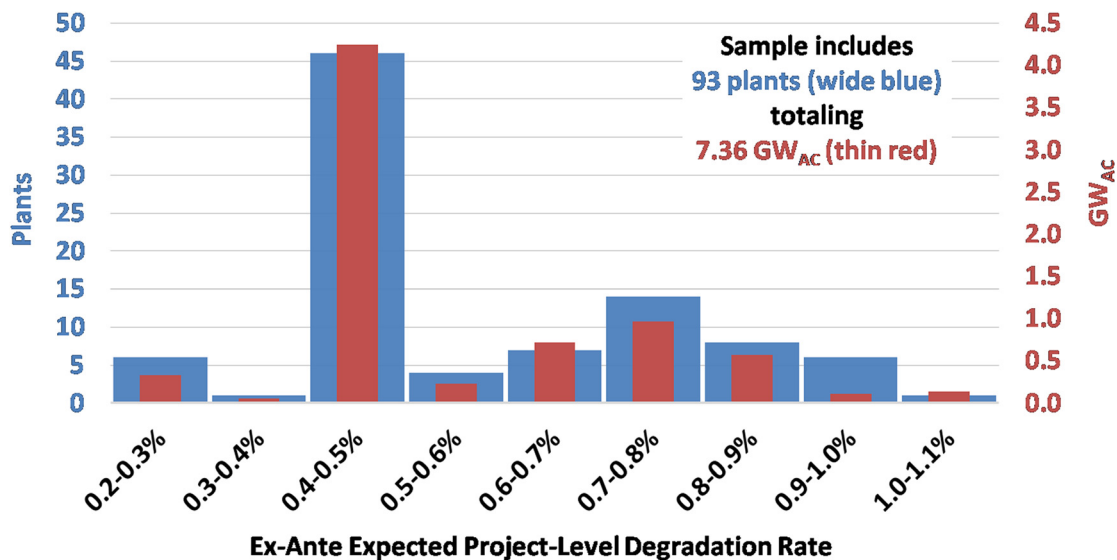


FIG. 6. Histogram of degradation rates contained in PPAs for utility-scale PV plants.

projects totaling 2.1 GW<sub>DC</sub> for which we have hourly generation data going back in time. Specifically, we ran our fixed effects model on the rolled-up calendar year generation data for all 16 of these projects and compared the results to an analysis of the hourly data conducted using NREL’s RdTools open-source platform (NREL, 2019d). Using annual data, the fixed effects model found an average degradation rate of  $-1.23\%/year$  for this sub-sample, while the average across RdTools’ estimates using hourly data were  $-0.99\%/year$ —a gap of  $0.24\%/year$ . Again, it is worth noting that for both datasets, we lack visibility into soiling and availability, which can exacerbate system-level degradation estimates if not properly controlled for. Here, though, we are most interested in the *relative difference* between the annual and hourly results, rather than the absolute numbers for each.

Figure 7 shows the individual plant-level comparisons for each plant-method pairing. Because the fixed effects model cannot be used for individual plants (i.e., it is only useful when analyzing a portfolio of projects), the annual results presented in Fig. 7 are instead based on a simple regression of the irradiance-normalized capacity factor against project age for each plant (and so are not directly comparable to the fixed effects approach or results presented earlier). The two approaches—i.e., the simple regression using annual data and the RdTools analysis of hourly data—are in good agreement on some plants but further apart on others. Where discrepancies exist, they could be caused by each approach analyzing slightly different time periods—e.g., in some cases, the hourly data did not go back as far as the annual data, and/or RdTools made use of partial-year hourly data that were excluded from the annual analysis.

For all 16 plants, the median absolute difference in degradation rates across the two methods is  $0.22\%/year$ , with the annual data showing greater degradation (consistent with the fixed effects model

results of  $-1.23\%$  compared to the hourly data average of  $-0.99\%/year$ ). Though we would not expect these two different methods (and different data resolutions) to yield identical results, this comparison nevertheless suggests that we cannot rule out the possibility that we find a greater degradation rate than past studies due, in part, to methodological differences and/or our use of low-resolution annual generation data. We also emphasize once again the fact that we estimate degradation at the system-level rather than at the module-level, which opens up many more degradation pathways. For example, based on recent news of tracker replacements after just six years of operations (Weaver, 2019), the high degradation rate of plant number 2 in Fig. 7 (i.e., the Alamo 1 project in Texas) is likely the result of malfunctioning trackers rather than of sub-par modules.

Whatever the cause of the apparent difference, we note that the gap between the  $-1.3\%/year$  system-level degradation rate that we find, the  $-0.8\%/year$  to  $-1.0\%/year$  range found in earlier studies, and the  $-0.5\%/year$  rate most-often stated in PPAs is significant from a financial perspective. Using an in-house financial pro forma model and assuming the following parameters— $\$1.2/W_{AC}$  CapEx,  $\$20/kW_{AC}\text{-year}$  OpEx, 30% net capacity factor (in AC terms), 4% debt interest rate, 2.5% inflation, 27% combined federal and state income tax rate, 30% ITC, and a capital structure (debt/equity ratio) that varies based on a debt service coverage ratio of 1.3—we find that a generic utility-scale PV project with a levelized 25-year PPA price of  $\$31.5/MWh$  would generate an internal rate of return of 10% assuming a degradation rate of  $-0.5\%/year$ , but only 5.1% if actual degradation turned out to be  $-1.0\%/year$  and only 2.6% if actual degradation turned out to be  $-1.3\%$ . In other words, worse-than-expected degradation of the magnitude found in this analysis can claw back most or all of an investor’s return in a utility-scale PV project.




FIG. 7. Degradation rate comparison: annual (regression) vs hourly (RdTools) data.

## VII. CONCLUSIONS

Our analysis of the performance and reliability of a large sample of utility-scale PV projects built from 2007 through 2016—and which, in 2017, generated more than 50% of all solar electricity in the United States—finds that first-year generation has generally matched *ex ante* expectations but that subsequent energy yield degradation has been greater than expected on average. Compared to past studies, which have generally found system-level energy yield degradation to be in the  $-0.8\%/year$  to  $-1.0\%/year$  range, and a sample of PPAs, two-thirds of which expect degradation rates of  $-0.5\%/year$  or better, we find an average fleet-wide degradation rate of  $-1.3\%/year$  ( $\pm 0.2\%$ ). This rate is an improvement over the  $-1.7\%/year$  found prior to correcting for interannual variance in irradiance and curtailment. A side analysis of a variety of project-level characteristics suggests that degradation rates tend to be lower in magnitude (i.e., closer to estimates from past studies but still worse than rates assumed in most PPAs) among newer projects and at sites with lower long-term average temperatures (both intuitive results).

Some of the difference between our results and those of previous studies could be methodological and/or related to our use of low-resolution annual data. Yet, for large-scale, *ex post* studies of fielded performance—which are of great interest to both potential and current solar project investors—the use of publicly available, low-resolution generation data may be par for the course. While our coarse data resolution prohibits us from properly attributing and allocating the estimated  $-1.3\%/year$  system-level degradation rate among its probable causes—e.g., module-level degradation, balance of plant degradation, soiling, and availability—from an investor's perspective, this breakdown is perhaps less important than the topline estimate of  $-1.3\%/year$  ( $\pm 0.2\%$ ) in terms of affecting the bottom line.

Even if our estimated degradation rate of  $-1.3\%/year$  is somewhat excessive—as potentially suggested by a comparison to RdTools analysis of a sub-sample of plants for which we have hourly data—the absolute magnitude of the difference to RdTools results ( $0.24\%/year$ ) would still put our adjusted system-level degradation rate at the high end of the range of past studies. Perhaps more importantly, our results—and the results of past studies—are significantly higher than the expected degradation rates most-often memorialized within PPAs. Basic financial analysis suggests that an actual degradation rate of  $-1.0\%$ —which is twice the rate expected within the majority of PPAs we sampled—can wipe out roughly half of an investor's return in a utility-scale PV project. For utility-scale PV to live up to its potential in helping to de-carbonize the electricity sector, investors and other stakeholders need a better understanding of the long-term reliability of these plants' performance over time. This study—which introduces a new approach based on publicly available, low-resolution data—represents an initial step in that direction.

## ACKNOWLEDGMENTS

This work was authored by Lawrence Berkeley National Laboratory, under Contract No. DE-AC02-05CH11231 with the U.S. Department of Energy. Funding provided by the U.S. Department of Energy Office of Energy Efficiency and Renewable Energy Solar Energy Technologies Office is acknowledged. The views expressed in the article do not necessarily represent the views of the DOE or the U.S. Government. 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 work, or allow others to do so, for U.S. Government purposes.

## DATA AVAILABILITY

The data that support the findings of this study are available from the corresponding author upon reasonable request.

## REFERENCES

- Anderson, M., Defreitas, Z., and Hasselbrink, E. F., Jr., "A system degradation study of 445 systems using year-over-year performance index analysis," in Photovoltaic Module Reliability Workshop 2013, SunPower Corporation, Golden, Colorado, February 26–27 (2013).
- Bolinger, M., Seel, J., and Robson, D., *Utility-Scale Solar: Empirical Trends in Project Technology, Cost, Performance, and PPA Pricing in the United States (2019 Edition)* (Lawrence Berkeley National Laboratory, Berkeley, CA, 2019).
- Clarion Energy, "Revolutionizing solar market transactions," *Global Power and Energy Elites 2020: Finance and Investment* (Clarion Energy, 2019), pp. 68–69.
- Deceglie, M., Jordan, D., Nag, A., Shinn, A., and Deline, C., "Fleet-scale photovoltaic energy-yield degradation analysis applied to hundreds of residential and non-residential PV systems," *IEEE J. Photovoltaics* **9**(2), 476 (2019).
- DNV GL, "Solar power performance white paper—United States," Report No. ENA-WP-21-A (2019).
- Germer, S. and Kleidon, A., "Have wind turbines in Germany generated electricity as would be expected from the prevailing wind conditions in 2000–2014?," *PLoS One* **14**(2), e0211028 (2019).
- Habte, A., Sengupta, M., and Lopez, A., "Evaluation of the National Solar Radiation Database (NSRDB): 1998–2015," Report No. NREL/TP-5D00-67722 (National Renewable Energy Laboratory, Golden, Colorado, 2017).
- Hamilton, S., Millstein, D., Bolinger, M., Wisner, R., and Jeong, S., "How does wind project performance change with age in the United States?," *Joule* **4**, 1004–1017 (2020).
- Hansen, C. W., Martin, C. E., and Guay, N., "Analysis of global horizontal irradiance in version 3 of the National Solar Radiation Database," Report No. SAND2015-8023 (2015).
- Hasselbrink, E., Anderson, M., Defreitas, Z., *et al.*, "Validation of the PVLIFE model using 3 million module years of live site data," in 2013 IEEE PVSC (2013), pp. 0007–0012.
- Hughes, G., *The Performance of Wind Farms in the United Kingdom and Denmark* (Renewable Energy Foundation, 2012).
- Jordan, D., Deceglie, M., and Kurtz, S., "PV degradation methodology comparison—A basis for a standard," in 2016 IEEE 43rd Photovoltaic Specialists Conference (PVSC), Portland, OR, June 5–10 (2016a).
- Jordan, D., Deline, C., Kurtz, S., Kimball, G., and Anderson, M., "Robust PV degradation methodology and application," *IEEE J. Photovoltaics* **8**(2), 525 (2018).
- Jordan, D. and Kurtz, S., "Field performance of 1.7 GW of photovoltaic systems," *IEEE J. Photovoltaics* **5**(1), 243 (2015).
- Jordan, D., Kurtz, S., VanSant, K., and Newmiller, J., "Compendium of photovoltaic degradation rates," *Prog. Photovoltaics* **24**, 978–989 (2016b).
- Jordan, D., Marion, B., Deline, C., Barnes, T., and Bolinger, M., "PV field reliability status: Analysis of 100,000 solar systems," *Prog. Photovoltaics* (2020).
- Köntges, M., Kurtz, S., Packard, C., Jahn, U., Berger, K. A., Kato, K., Friesen, T., Liu, H., and Van Iseghem, M., "Review of failures of photovoltaic modules," Report No. IEA-PVPS T13-01 (International Energy Agency, Photovoltaic Power Systems Programme (IEA-PVPS), Paris, 2014).
- Kraus, P., Massué, C., Heumann, S., and Schögl, R., "Reliable long-term performance assessment of commercial photovoltaic modules tested under field conditions over 5 years," *J. Renewable Sustainable Energy* **11**, 063501 (2019).
- kWh Analytics, see <https://www.kwhanalytics.com/solar-risk-assessment> for "Solar Risk Assessment 2019: Quantitative Insights from the Industry Experts" (last accessed December 12, 2019).

- National Renewable Energy Laboratory (NREL), *National Solar Radiation Database (NSRDB)* (National Renewable Energy Laboratory, Golden, CO, 2019a).
- National Renewable Energy Laboratory (NREL), *System Advisor Model (SAM)* (National Renewable Energy Laboratory, Golden, CO, 2019b).
- National Renewable Energy Laboratory (NREL), *PVWatts Calculator* (National Renewable Energy Laboratory, Golden, CO, 2019c).
- National Renewable Energy Laboratory (NREL), *Accurate Degradation Rate Calculation with RdTools* (National Renewable Energy Laboratory, Golden, CO, 2019d).
- Norton Rose Fulbright, “Solar finance outlook,” *Project Finance Newswire* (Norton Rose Fulbright, 2019), pp. 24–34.
- Olauson, J., Edström, P., and Rydén, J., “Wind turbine performance decline in Sweden,” *Wind Energy* **20**, 2049–2053 (2017).
- Staffell, I. and Green, R., “How does wind farm performance decline with age?,” *Renewable Energy* **66**, 775–786 (2014).
- U.S. Energy Information Administration, *Electric Power Monthly, with Data for December 2018* (U.S. Department of Energy, Washington, DC, 2019a).
- U.S. Energy Information Administration, *Form EIA-923 Data for 2018* (U.S. Department of Energy, Washington, DC, 2019b).
- Weaver, J., “Six-year-old solar power plant partially repowered with new trackers and bifacial solar modules,” *PV Magazine* (2019).