andlinger center
for energy+the environment

PRINCETON
UNIVERSITY

Sunlight to Electricity:

Navigating the Field

An Energy Technology Distillate of the
Andlinger Center for Energy and the Environment
Princeton University

Contributors

**Barry P. Rand, Forrest Meggers, William C. Witt,
Manali Gokhale, Samantha Walter, Robert Socolow**

August 2017

Table of Contents

Acknowledgements

Ted Borer was a valuable guide to the Princeton University case study. William Hackett, Anne Hoskins, Benjamin Scott Hunter, Ralph Izzo, Tom Kreutz, Tom Leyden, Andrew Mills, and Michael Winka provided insights into enabling policy and commercialization. We benefited from the summer research of three Princeton undergraduates: Angelo Campus, Lindsey Conlan, and Isabella Douglas.

An intermediate version of this report was critiqued by the eight graduate students in a Policy Workshop at Princeton University on community solar power taught by Anne Hoskins and Jeanne Fox. Stephen Lassiter, one of these students, judiciously combined his comments with the others – from Vivian Chang, Cara Goldenberg, Jack Hoskins, Zhongshu Li, Eri Nakatani, Sheree Oluwafemi, and Hannah Safford.

We also benefited from comments on early drafts from Emily Carter, Philip Chew, Phillip Hannam, David Kanter, Elena Krieger, Eric Larson, Nicolas Lefevre, and Sigurd Wagner. A late draft was helpfully reviewed by Joe Berry, Michael McGehee, Cliff Rechtschaffen, Lenny Tinker, and Jessika Trancik.

The production of this report was guided diligently and creatively by Nari Baughman, Darcy Cotten, Joel Oullette, and Jennifer Poacelli.

At every stage, our revisions were propelled by Lynn Loo’s commitment to clear and effective communication.

The Andlinger Center for Energy and the Environment is grateful to the High Meadows Foundation, the Nicholas Family, and an anonymous donor whose gifts are helping to advance public understanding of critical issues related to energy and the environment through this Energy Technology Distillate.

Article 1: Overview

The goal of this Distillate is to enable the reader to understand the state of solar energy today and to develop his or her own views of some of the key issues that loom over solar power's future.

Today, for the first time in human history, a commercially significant quantity of solar energy is being turned directly into electricity. Global capacity to produce solar electricity was about 50 times greater in 2016 than 10 years earlier. Solar power has grown rapidly in Europe, Asia, and North America.

No one knows how long future solar growth will resemble the past. It is conceivable but far from certain that solar power will dominate the global electricity system by mid-century. There is still a long way to go. In 2016 only about 1.5 percent of total global electricity came from solar power. In the U.S. the percentage was nearly the same, even though one million U.S. homes have solar panels on their roofs and several of the world's largest solar installations are in the deserts of the southwestern U.S.

This Distillate explores five open questions related to solar power's future:

1. *Will distributed and centralized deployment both flourish?* Solar cell technology is spectacularly modular: a solar cell will convert sunlight into electricity whether on a rooftop or in a multi-thousand-acre field, and assemblages of these cells are housed in panels that are essentially the same wherever they are used. Due to this modularity, the plummeting costs of solar cell technology have had a dramatic, positive impact on the growth of solar power at all scales. Future deployment could tilt toward very large projects because of economies of scale: large projects have substantially lower construction costs than small projects, for the same amount of electricity generated. However, distributed electricity generation, especially if accompanied by distributed electricity storage, may enable innovative grids that are more flexible and resilient than the centralized grids of the past. If deep penetration of distributed solar generation into electricity markets is achieved, political support for pro-solar policy will strengthen, to the likely benefit of centralized solar power as well. The path forward may well feature parallel development at large and small scale – with much geographical variation in the mix of the two scales.
2. *How much can balance-of-system costs be reduced?* The principal challenge of past decades was reducing the cost of the solar cell and the solar panel that houses the cell. Now, "balance-of-system" costs are emerging as the principal cost concerns. The balance of system, here, is all of a solar power project except for the solar panels: the land, the structure that holds the panel, any tracking hardware, the inverters that change the direct current (DC) produced by the cell into the alternating current (AC) required by the user, installation at the site, interconnection to the grid, and business costs such as financing, permitting, and insurance.
3. *Will crystalline silicon remain the workhorse of solar power?* Today, crystalline silicon has 90 percent of the solar cell market. Can any of the new thin-film technologies challenge silicon's dominance, now that the silicon cell industry has developed so much infrastructure and experience? The crystalline silicon solar cell has the limitation (thus far) of being available only as a rigid structure, which limits potential applications. Its competitors are thin films, whose versatility assures that there will be at least niche markets for some of them, even if they do not become significant producers of solar electricity. To enter the market, a thin film will need to convert sunlight to electricity at substantially higher efficiency than the crystalline solar cell, demonstrate stability and ease of manufacture, and avoid scarce or toxic materials.
4. *Will solar power subsidies disappear?* Government policies favorable to solar electricity are called "incentives" by their proponents and "subsidies" by their detractors. Subsidies have enabled solar energy to mature, and now they are shrinking, both for centralized and distributed generation, as solar power becomes increasingly competitive. The system costs of incentives for distributed generation were small when they were paid only to "early adopters," but as the fraction of beneficiaries in an eligible group grows, the non-adopters bear more noticeable costs and push back.

On the other side of the argument, those who favor incentives stress their direct environmental benefits, including cleaner air and less rapid climate change. Specific to distributed generation, they also note that producing electricity closer to the point of use, especially in combination with dispersed electricity storage, can reduce grid congestion and improve system performance.

5. *Will the intermittency of solar power soon throttle its expansion?* The full system in which the solar panel is embedded consists not just of the panel and the balance of the system at the project level, but also the electricity grid that must accommodate an intermittent and only partially predictable supply. Grid regulators and operators are able to accommodate the intermittency and unpredictability inherent in solar power when its market share is tiny. However, in an increasing number of places and increasingly often, solar power is now raising grid-management problems resulting from oversupply at mid-day, as well as rapid output gains in the morning and losses in the evening.

   Compensating responses are emerging. The grid can be made more accommodating by reducing the presence of nuclear and coal power plants, which function best when running at a constant rate, in favor of hydropower and gas-turbine power, whose output can be varied rapidly. The grid can make greater use of centralized and distributed electricity storage and can be extended geographically to integrate distant sources that have complementary time profiles. Time-of-use pricing can be more aggressively implemented to induce shifts in supply and demand by several hours. The choice of compass orientation for some stationary solar collectors in the northern hemisphere will then become southwest instead of due south, flattening the peak at noon. People will become more aware of whether a day is sunny or cloudy as they find themselves washing their dishes (and, perhaps, charging their electric-car battery) preferentially at mid-day on sunny days.

   Intermittency has been relatively invisible thus far, even though intermittency is arguably the Achilles heel of solar power, hobbling its path forward. Indeed, a widely cited objective for solar power, "grid parity," neglects intermittency. Solar energy achieves nominal grid parity when the cost of a kilowatt-hour of solar energy is the same as a kilowatt-hour from, say, natural gas or coal. But this takes no account of solar power's limited ability to produce power when desired, and therefore its higher grid-integration costs.

Roadmap

We have endeavored to treat technology and policy with equal seriousness. We have written for the reader who has little technical background but an appetite for scientific argument and curiosity about the policy domain. Discussions of technology and policy are in separate articles, however, so as to enable readers to read selectively if they come to the subject with a stronger interest in one than the other.

The four articles in this Distillate that follow this *Overview* (Article 1) address the five questions above, and they go considerably further. Article 2, "The Solar Panel: Key Concepts and Vocabulary," introduces the quantitative concepts widely used to discuss the deployment of solar power and to measure the performance of solar projects in physical and economic terms.

Article 3, "*From the Sun to the Solar Project*," deals with the first of the five questions above by comparing projects at various scales. It calls attention to the prominence of "mid-scale" projects on the rooftops of commercial buildings, in public parks, and in other settings, much larger than projects on the roofs of homes but similarly not owned by electric utilities. It also discusses the second question, balance-of-system issues. Prior to dealing with these questions, Article 3 describes the massive amount of energy arriving from the sun and then provides views of solar power's deployment at descending geographical scales: planet, country, state, and individual project.

Article 4, "*Solar Cell Technology*," illuminates the third question, the dominance of crystalline silicon. It describes many of the technologies used to convert sunlight into electricity, both currently commercialized and on the technological frontier, highlighting features that affect their competitiveness.

Article 5, "*Grid Integration and Policy*," takes on the fourth question (subsidies), with an emphasis on the U.S. and the state of New Jersey. Although New Jersey is atypical in the extent to which solar energy has been promoted by the state government, many other states have adopted similar policies and are confronting similar controversies. Article 5 also explores the critical fifth question (intermittency), describing its current emergence as a priority for the electric grid and the variety of partial responses in view.

The Distillate concludes with a brief appendix that presents some illustrative results from Princeton University's own solar project, which was the initial springboard for our report. Indeed, this Distillate generalizes what we learned from that project's technology choices, the many projects in New Jersey that it resembles, and its interactions with New Jersey and federal incentives.

Missing from our Distillate are several important issues. There is almost no discussion of the world's many giant solar projects (nearly all of them thousands of miles from New Jersey), and solar power at all scales in China, India, and elsewhere in the developing world.

Regarding technologies, this Distillate considers only flat-panel solar photovoltaic (PV) technology. It excludes "concentrating solar power,"[1] which is a second, currently competitive large-scale (but not rooftop) solar electricity technology. Also excluded are direct solar power for water heating and cooking, applications that are expanding rapidly in the developing world.

---

[1] Concentrating solar power (sometimes called "solar thermal" power) uses mirrors to focus sunlight and heat a fluid (liquid or gas) to a high temperature, whereupon the hot fluid powers an engine to produce electricity. In one version, sunlight is focused on long tubes running along the axis of parabolic troughs; in another, sunlight is focused onto a small spatial region at the top of a "power tower."

Article 2: Key Concepts and Vocabulary

In this article we introduce key concepts and specialized vocabulary for solar energy. We explain some quantitative characteristics of the individual solar panel, including electricity produced, cost, and carbon dioxide saved. We work out deliberately oversimplified numerical examples. Our objective is to demystify.

Watts and Watt-Hours

The Watt

Some electrical devices produce electricity and others consume it. The rate at which electricity is produced or consumed is measured in watts, and the amount is measured in watt-hours. Producing or consuming electricity at the rate of 1 watt for an hour results in the production or consumption of 1 watt-hour.

A 60-watt light bulb consumes electricity at the rate of 60 watts when turned on, a toaster making toast consumes power at a rate of about 1,000 watts, or 1 kilowatt, and the largest jet engines can produce power at a rate of about 100 million watts, or 100 megawatts.

Notably for this Distillate, the intensity of sunlight on a surface perpendicular to the Sun's rays when the Sun is high in the sky on a clear day (peak conditions) is approximately 1,000 watts for each square meter of surface. A typical solar panel has an area of 1.5 square meters. It therefore can receive sunlight at a rate of 1,500 watts under peak conditions.

The Watt-Hour

The dash (hyphen) in watt-hour means that a multiplication is involved. A 60-watt bulb will consume 60 watt-hours when it is turned on for one hour and 120 watt-hours when it is on for two hours.

The kilowatt-hour is the unit most commonly used to track electricity consumption and production, and it is the unit that appears on home electricity bills. Electricity is also often measured using the megawatt-hour, which is equivalent to 1,000 kilowatt-hours. In energy markets where solar energy certificates are bought and sold, one certificate represents 1 megawatt-hour of solar electricity production.

Watts and watt-hours are frequently confused, in part because the watt is one of the few rates with a name of its own.[2] Dividing watt-hours (a unit of energy) by hours (a unit of time) yields watt-hours per hour, or watts. If a home consumes 360 kilowatt-hours of electricity in a 30-day month, it consumes at an average rate of half a kilowatt (500 watts), since a 30-day month has 720 hours.

Conversion Efficiency

The most cited attribute of a solar cell and solar panel is its efficiency, which is electricity output divided by solar energy input. A "rated efficiency" is determined in the laboratory in a simulation of direct sunlight.

The panel efficiency is approaching 20 percent in projects being built today. Cells with efficiencies of 10 percent or less have special applications, and a conversion efficiency above 30 percent can be achieved today with some expensive composite ("multijunction") solar cells.

We return to our 1.5 square-meter panel that receives 1,500 watts of solar energy under peak conditions. If it has a conversion efficiency of 20 percent, it can produce electricity at a peak rate of 300 watts. It is called a "300-watt panel," and 300 watts is its rated output.

Capacity Factor

The "capacity factor" is a widely used index of performance, applicable to any power plant. It is the *actual* production of electricity produced at a power plant, divided by the maximum amount of electricity the plant could have produced if it had run at full rated capacity (over some common period such as a year). It is not unusual for a modern nuclear power plant to achieve a capacity factor of 90 percent, given that nuclear plants run at nearly their maximum capacity almost every

---

[2] Others units that describe rates include the ampere (a rate of flow of electric current) and the knot (a measure of nautical speed).

day of the year. Some power plants follow and respond quickly to the ups and downs of electricity demand in a region and have capacity factors near 50 percent. Still others are "peaking plants," designed to run only during the few times of the year when demand is particularly high (for example, on an extraordinarily hot summer afternoon); these have capacity factors in the single digits.

The capacity factor for a solar power plant is the electricity produced by the plant over some time interval, divided by the electricity the plant would have produced if all of its panels had produced electric power at their rated output throughout the same time interval. The capacity factor is affected by the sunniness of the location, how steeply the panels are tilted relative to a horizontal surface and their compass orientation, and whether the panels are stationary or track the sun. The capacity factor is reduced to the extent that the plant's panels at certain times are covered with snow or debris, or they are in the shadow of trees, nearby buildings, or other panels. The capacity factor is also reduced when a plant is shut down for maintenance, or if a plant is producing electricity but a manager of an electric grid forbids an operating plant from sending its electricity onto the grid because of some grid-management issue.

The capacity factor for the world's solar power (an average over all the solar power plants) in 2014 can be estimated from estimates that global installed capacity was 181 million kilowatts and global solar production was 211 billion kilowatt-hours.[3] Global production, therefore, was equivalent to production at full capacity for 1,160 hours and no electricity production during the rest of the year. Rounding up to 1,200 hours and dividing by the 8,766 hours in an average year gives a capacity factor of 14 percent.

Combining the Capacity Factor and the Conversion Efficiency

The capacity factor and the conversion efficiency are entirely different concepts, but they combine multiplicatively to determine the output of a solar power plant. The capacity factor measures how much sunlight falls on the panels. The conversion efficiency measures how much electricity is produced by that sunlight.

Quantitatively, the capacity factor and the conversion efficiency are of comparable importance. A representative value for both is 20 percent: a power plant located in a favorable location has a capacity factor of 20 percent or more, and the conversion efficiency of most commercial solar panels is close to 20 percent.[4] Moreover, in both cases most values for real projects fall between 10 to 30 percent.[5]

To be sure, for a specific solar facility, the actual scores within these two ranges are critically important determinants of its attractiveness as an investment. A facility with two scores of 30 percent produces roughly nine times as much power as an identical facility where both scores are 10 percent.

We return again to our 1.5 square-meter, 20-percent-efficienty panel with a rated capacity of 300 watts. If its capacity factor is also 20 percent, it will produce electricity at an average rate of 60 watts. Over a year, it will produce (rounding off) about 500 kilowatt-hours of electricity (60 watts, multiplied by 8766 hours, equals 526 kilowatt-hours).

Panel Economics: Balance of System and Payback Period

The "payback period" is the amount of time required for an investment to break even. To find the payback period for a residential solar project, we require the cost of residential electricity and the cost of the residential project.

Representative costs for electricity in the U.S. are 5 cents per kilowatt-hour for wholesale electricity (the cost to the utility of producing the power) and 15 cents per kilowatt-hour for retail electricity (the cost of electricity provided to a household by the utility). The difference is attributable to the capital and operating costs of the transmission and distribution system and overhead (maintenance, billing, profit, etc.).

The average cost of a panel in the U.S. has recently dropped below $1 per peak-watt and is still falling. Non-panel costs, referred to as "balance of system" costs, make up the majority of project costs today, and their costs are falling too. Representative (conservative) total project costs are $2 per peak-watt for a utility-scale system and $4 per peak-watt for a residential rooftop system. With these cost assumptions, a single 300-watt panel installed at a utility-scale project will cost its owner $600, and the same panel installed on a residential roof will cost $1,200.

---

[3] https://www.worldenergy.org/wp-content/uploads/2016/09/Variable-Renewable-Energy-Sources-Integration-in-Electricity-Systems-2016-How-to-get-it-right-Executive-Summary.pdf, table on p. 2.

[4] In several examples in this Distillate, we use 20 percent for both.

[5] A 30-percent capacity factor can even be exceeded if a panel is located in a desert and is mounted on a motor-driven support that tracks the sun.

You can walk past a house with a solar panel array and estimate its cost by counting the number of panels. The average capacity of the solar collection system in a U.S. home is approximately 5,000 peak-watts, which corresponds to a home with about 16 panels and a cost, at today's prices, of about $20,000. A large solar power plant in the desert in the southwestern U.S. rated at 300 million peak-watts has about one million panels; if built today, at $2 per peak-watt, it would cost $600 million.

We can work out the payback period for this 300 peak-watt, $1,200 residential panel, knowing that it produces 500 kilowatt-hours of electricity annually. Valuing the 500 kilowatt-hours at the retail rate above, the panel saves the residential customer $75 of purchased electricity each year. If a homeowner spends $1,200 to save $75 per year, her payback period (the time to break even) is 16 years.

Here, we have not included any state or federal incentives. In Article 5 this calculation is redone with specific New Jersey and federal incentives included, and the payback period is found to be three times shorter, or about five years.

Value of Improved Efficiency

Improvements in solar cell efficiency translate into reduced costs for the balance of the system, per unit of electricity produced, because more electricity is produced for the same balance of system cost. We work out an example, starting from the 300 peak-watt, $1,200 rooftop panel, above, where the panel costs $300 and the balance of system costs $900. We assume that a homeowner decides to install six of these panels to meet her budget and provide the solar electricity that she wants. She spends a total of $7,200: $1,800 for the six panels and $5,400 for the balance of system. We further assume that the available panel is 20 percent efficient in converting sunlight to electricity.

Now, a new panel becomes available which costs exactly the same but is one-fifth more efficient (24 percent efficient), so she can buy five panels instead of six panels and get the same amount of solar electricity. We make the rough approximation that that the cost of the balance of system depends only on the number of panels, and is now five-sixths as much, or $4,500, because there are now five panels instead of six. (We neglect costs, like permitting, which might not come down when there are fewer panels.) The more efficient panel has reduced the balance of system cost by $900. The homeowner should be willing to pay up to $900 more for the five panels, or $180 more per panel, and still come out ahead. Since the original panel costs $300, the homeowner should be willing to pay as much as $480 per panel for the more efficient panel, 60 percent more. This example thus illustrates the trade-off, where paying more for increased efficiency results in paying less for the balance of system.

Levelized Cost of Electricity

The levelized cost of electricity is the cost of building, operating, and maintaining a facility over its lifetime, divided by the amount of electricity it produces in its lifetime. If we make the assumption that the residential panel above, which produces 500 kilowatt-hours of electricity each year, will have a lifetime of 20 years, then it will produce 10,000 kilowatt-hours over its lifetime. If we further make the simplifying assumption that the only significant cost for the panel is the $1,200 installation cost at the beginning (for example, we neglect maintenance costs), then the levelized cost of electricity is 12 cents per kilowatt-hour. This is higher than the levelized cost of new natural gas power today, but lower than the levelized cost of new nuclear power. The levelized cost would be much lower for a panel used at a large utility project. The levelized cost is a problematic concept for solar power because complications due to its intermittency are ignored.

Cost of Avoided Emissions of Carbon Dioxide

How cost-effectively does the residential panel, above, reduce carbon dioxide emissions to the atmosphere? Our panel produces 10,000 kilowatt-hours of electricity over its 20-year lifetime, and so, presumably, some mix of other power plants that serve the same region produce 10,000 kilowatt-hours less. Thus, the answer depends on the carbon dioxide emissions of the other power plants: the displaced electricity could be assignable to either coal plants or nuclear plants, for example. Let's assume that what is displaced is an average U.S. power plant, which emits a ton of carbon dioxide for each 2,000 kilowatt-hours of power produced. In that case, about five tons of carbon dioxide is not emitted into the atmosphere thanks to our residential panel. Since the cost of the panel is $1,200 (ignoring all costs after the panel is installed), it costs $240 to prevent one ton of carbon dioxide from entering the atmosphere. The corresponding estimate could be several times less for a panel at a large utility installation in a favorable location, and after costs have fallen further. This calculation neglects the carbon dioxide emissions associated with manufacturing the panel in the first place; including manufacturing emissions will decrease the net emissions reduction achieved by the panel and increase the cost of avoided emissions.

Article 3. From the Sun to the Solar Project

In this article we first describe sunlight and how it falls on the earth. We then provide a high-level view (global, national, and by U.S. state) of the rate of production of solar electricity and its remarkable growth in recent years. We conclude with observations about distributed generation and other issues at the project level.

A discussion of solar projects follows, where we develop information relevant to the intriguing question of whether the mix of large and small projects characterizing current solar power will be sustained in the future. We introduce a three-part categorization of solar projects (“utility,” “mid-scale,” and “residential”). We emphasize the importance of the little-discussed mid-scale: the projects that are built on the rooftops of commercial buildings and on land owned by public and private institutions, not owned by utilities but much larger than projects on the roofs of homes. We also discuss “distributed generation” (both mid-scale and residential projects), which may conceivably become the basis of a restructuring of the current centralized utility.

We conclude Article 3 with a discussion of the “balance of system,” which is every aspect of a project other than the high-tech panel. Costs for typical projects are disaggregated to highlight the balance of system, whose cost is now at least as important as the cost of the panel. The underlying question is the extent to which “balance-of-system” costs can continue to fall in the future. The article concludes with a description of some imaginative uses of solar collectors in buildings, where the production of electricity is a side objective.

A. The Sun and the Earth

Sunlight Above the Earth

The solar energy that can be made useful to people in the form of electricity is a tiny fraction of the solar energy that the Sun produces and radiates to space. The Sun emits energy at a rate of 400 billion quadrillion (4×10^{26}) watts, uniformly and in all directions. This rate varies by about one-tenth of one percent from year to year, depending upon the number of sunspots on the Sun’s surface.

The Earth intercepts a tiny fraction of this energy: 170 quadrillion (1.7×10^{17}) watts, or about one half of one billionth of the energy emitted by the Sun. The power for space satellites exploring distant parts of our solar system is produced with solar panels that intercept sunlight which would not have hit the Earth, but with this exception only the solar energy that the Earth intercepts is available today to power our civilization. Perhaps, someday, human beings will build structures in the solar system to harvest sunlight emitted in other directions.

It should be noted that Earth’s orbit around the Sun is not a perfect circle, so the amount of usable solar energy varies over the course of the year as the distance between the Earth and the Sun grows and shrinks. Sunlight is about 7 percent stronger when the Earth is closest to the Sun (at the beginning of January) than when it is furthest from the Sun (at the beginning of July).

Sunlight on the Earth’s Surface

The intensity of the Sun’s energy is about 30 percent greater at the top of the atmosphere, but various gases and aerosols reduce the intensity by absorbing sunlight as it travels through the atmosphere toward the Earth’s surface. Where the Earth’s surface is flat on a clear day at sea level, with the Sun directly overhead, the average intensity of direct sunlight is about 1,000 watts per square meter.

As shown in Figure 3.1, the average annual intensity of sunlight varies by location. The highest intensities are found in most of Africa, the Middle East, and Australia, as well as the southwestern U.S. and Mexico. The values mapped in Figure 3.1 include both sunlight that comes from the direction of the Sun and sunlight coming from other directions, known as diffuse sunlight. On a sunny day, levels of pollution, dust, and humidity determine the ratio of the direct to the diffuse components of incident solar radiation; averaged over a sunny day in a low-pollution environment with the Sun high in the sky, the energy arriving at a flat panel from diffuse light is about 20 percent of the total. On a fully cloudy day, a horizontal solar panel will collect two to five times less energy over a day than on a sunny day. On a partly cloudy day, the amount of sunlight incident on a solar panel can fluctuate by a factor of five or more over the course of minutes as a cloud passes between the Sun and the panel. This short-term variability is one of the key challenges to scaling up the deployment of solar power.

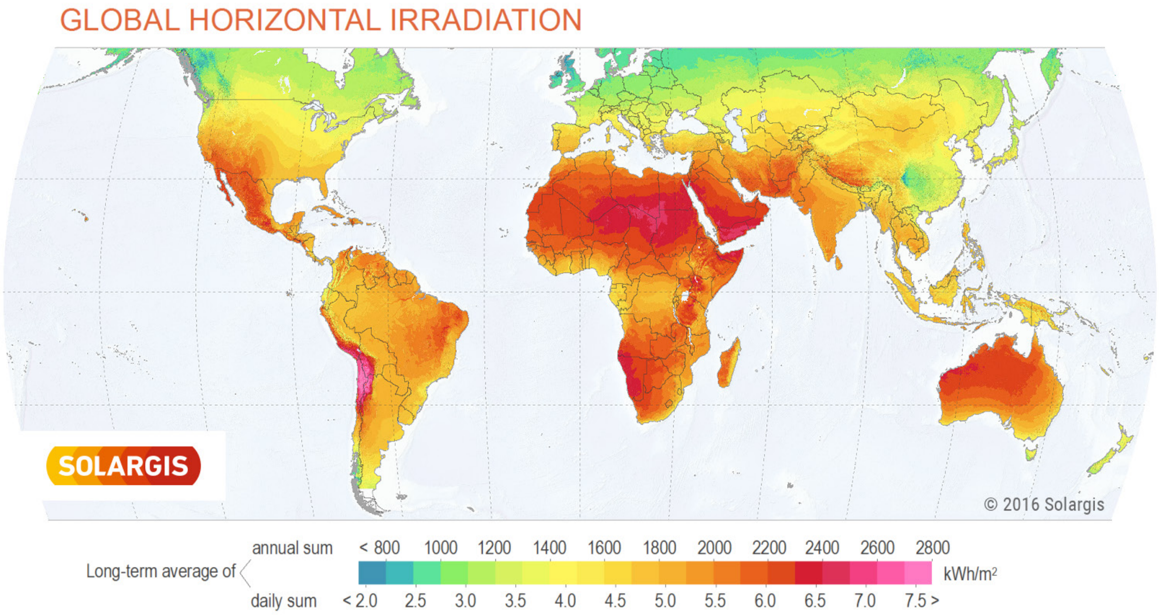

Figure 3.1: Annually averaged "irradiation" incident on a horizontal surface (the sum of direct and diffuse sunlight arriving over a period of time, presented here for an average day). The upper and lower horizontal scales are the annual and daily sums, respectively, in units of kilowatt-hours per square meter (kWh/M²) of surface area. Source: https://www.solargis.com

In Figure 3.1, the strength of average incident sunlight is measured in kilowatt-hours per square meter of horizontal surface per year (upper scale) and per day (lower scale). In particular, the lower scale runs from just below 2.4 kilowatt-hours per square meter to just above 7.2 kilowatt-hours per square meter. This scale is the one that the solar industry uses most frequently in quantifying the solar resource.

Dividing the numbers on the lower scale by 24 produces average power measured in kilowatt-hours per square meter per hour, which is the same as kilowatts per square meter. Thus, the average strength of sunlight ranges from just below 100 watts per square meter to just above 300 watts per square meter (square meter of *horizontal surface*). This range in the strength of incident sunlight can also be expressed as 10 to 30 percent of the peak rate, 1,000 watts per square meter, at which sunlight can be collected at the Earth's surface (clear day at mid-day, with the collector aligned perpendicular to the Sun's rays).

Yet another way to express the amount of sunlight that falls on a horizontal surface over a year at some location is in terms of the number of hours required to collect that much energy at that location from hypothetical panels collecting sunlight at its peak incident rate. Where the strength of incident sunlight is 20 percent of the peak rate, these panels would need to operate 20 percent of the year, and since there are 8,766 hours in an average year, they would need to operate about 1,750 hours per year. In these units, annual incident sunlight on a horizontal surface at specific locations on the globe varies from less than 1,000 hours per year to as much as 2,500 hours per year of peak sunlight.

The Solar Spectrum

Sunlight is a mixture of light of many colors; the mixture forms a spectrum. Figure 3.2 shows the spectrum of incident sunlight both at the top of the atmosphere and at the Earth's surface. The spectrum is conventionally divided into three regions, with "visible" light in the middle (here, violet on the left and red on the right), ultraviolet (more violet than the eye can see) on one side and infrared (more red than the eye can see) on the other.

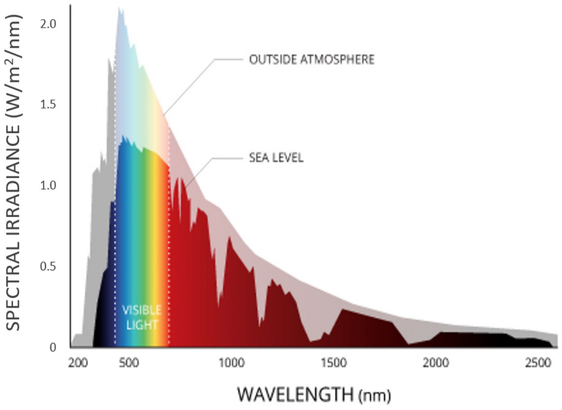

Figure 3.2: Distribution of incoming solar energy across the spectrum at the top of the atmosphere and at ground level. "nm" is nanometer, one billionth of a meter. Source: http://www.fondriest.com/environmental-measurements/parameters/weather/photosynthetically-active-radiation/.

The strength of incident sunlight is reduced throughout the spectrum, but unevenly. On its way through the atmosphere, much of the ultraviolet radiation is absorbed (by ozone and oxygen). Similarly, much of the infrared radiation is absorbed (especially by water vapor), resulting in most of the notches in the curve for ground-level solar energy at the right in Figure 3.2. At the Earth's surface, about 42 percent of direct sunlight is visible light, approximately 4 percent is ultraviolet, and the remaining 54 percent is in the infrared region.

Sunlight can be thought of as a collection of individual particles (photons), each carrying a specific amount of energy. The energy of a photon depends on the color of the light. Ultraviolet light is the most energetic, then visible, then infrared; within the visible spectrum, blue is more energetic than red. In the world of solar cells, a photochemical process typically requires some minimum amount of photon energy; thus, blue light can drive some processes that red light cannot.

The Path of the Sun through the Sky over a Year

The angle between the Sun and a solar panel determines how much power the panel can generate. The angle is easiest to understand at solar noon, when, every day of the year, the Sun is either directly to the south or directly to the north. The noon positions of the Sun are shown schematically in Figure 3.3 for a summer day and a winter day at a latitude typical of China and the U.S.

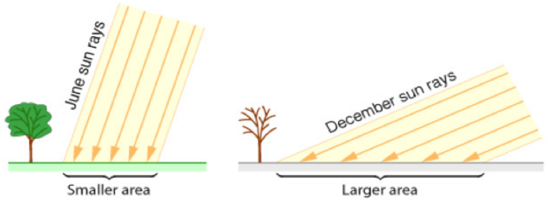

Figure 3.3: The angle of the Sun with the vertical at solar noon is displayed for a mid-latitude location in the northern hemisphere. Source: http://physics.weber.edu/schroeder/ua/SunAndSeasons.html

Figure 3.4 augments Figure 3.3 by showing four moments along the trajectory of the Earth around the Sun. At solar noon on March 21 and September 21, the Sun is to the south everywhere in the northern hemisphere. The angle between a line to the Sun and a vertical line is the same as the latitude at that location. For example, at the equator, where the latitude is zero, the Sun is straight overhead.

Relative to its position at solar noon on March 21 and September 21, the Sun at solar noon is further north throughout the period between March 21 and September 21 and further south between September 21 and March 21. On June 21 (the summer solstice and the longest day of the year) at solar noon, it is furthest north, 23.5 degrees further north than its location on September 21 and March 21. A person's shadow is shorter at solar noon on June 21 than at any other time of the year. On December 21 (the winter solstice and shortest day of the year) at solar noon, it is furthest south, again by the same 23.5 degrees relative to its position on September 21 and March 21. The 23.5 degrees angle is the tilt of the axis of the Earth's rotation, relative to the plane that contains the Earth's path around the Sun.

B. The Scale of Current Solar Power

The sunlight that strikes Earth in one hour carries more energy than is required to power human civilization for an entire year. This frequently encountered statement accounts for the energy consumed by power plants, vehicles, furnaces, boilers, and other facilities (in aggregate, "primary energy"), but excludes the sunlight required to grow food, to evaporate water so that we receive rain, and to enable other "ecosystem services."

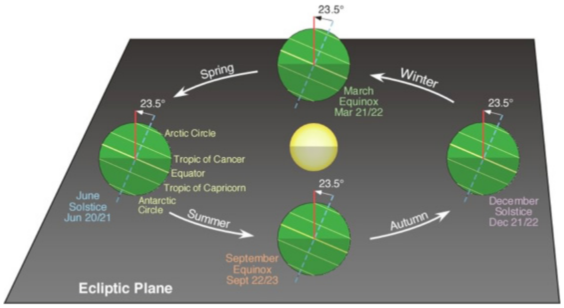

Figure 3.4: The Earth's position relative to the Sun on four key days of the year. Source: http://www.physicalgeography.net/fundamentals/6h.html

Electricity Production from All Sources

Of the total primary energy used by humans, about 40 percent is used to produce electricity. The rest is used directly by industry, vehicles, and buildings. Currently, the total capacity of the world's electric power plants of all kinds is approximately six billion kilowatts, and the world's annual electricity consumption is approximately 25,000 billion kilowatt-hours. Since there are 8,766 hours in an average year, the world's power plants would have produced approximately 50,000 billion kilowatt hours (6 times 8,766, rounding off) if the plants had run steadily at full capacity all year. We conclude that the world's power plants produce, on average, about half of the output that they could produce if they ran continuously at peak capacity. For any single power plant or group of plants, the "capacity factor" is the ratio of the actual production divided by the hypothetical production at peak capacity. Thus, the capacity factor of the world's power plants is currently about 50 percent (25,000 divided by 50,000).

Global Solar Electricity Production

At the end of 2016, the amount of solar photovoltaic (PV) power installed worldwide was 300 million peak kilowatts, 5 percent of the total capacity of the world's power plants of all kinds. Solar output is not as well documented as solar capacity, but if the capacity factor for global solar electricity production was 14 percent in 2016, as it was in 2014,[6] global solar electricity consumption in 2016 would have been about 360 billion kilowatt-hours of electricity, or about 1.5 percent of that year's total electricity from all sources.

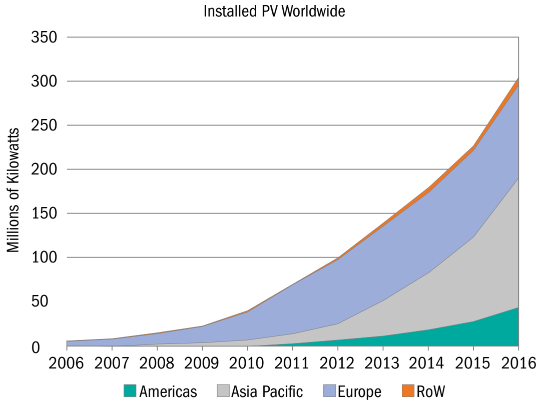

Figure 3.5: Installed generation capacity of solar photovoltaic (PV) production facilities, by world region, 2006- 2016. 1 gigawatt = 1000 megawatts = 1,000,000 kilowatts. RoW is the rest of the world. Source: International Energy Agency, Photovoltaic Power Systems Program, IEA PVPS Snapshot 2017: http://iea-pvps.org/index.php?id=trends0.

Figure 3.5 shows the growth of global solar power plant capacity from 2006 to 2016 and its distribution over broad geographical regions. Deployment in Europe dominated global expansion initially: since 2010 the annual growth rate in the Asia-Pacific region has been larger than in Europe, and the absolute increment over the previous year has been larger in the Asia-Pacific region since 2013. In 2016 the Asia Pacific region accounted for two-thirds of the growth in global capacity. Relatively, the Americas have been small players.

Deployment by Country

Figure 3.6 shows the solar capacity in place, by country, in 2016. More than half of the capacity is located in just four countries: China, Germany, Japan, and the U.S. During the year 2016, China installed about half of the world's total added capacity, as total global capacity grew by one third. In 2016 China and the U.S. added about 80 percent and about 60 percent to their 2015 solar capacity, respectively, with the result that China ended 2016 with about twice as much installed solar capacity as the U.S., which has about the same total capacity as Germany and Japan. Germany has the most installed capacity per capita of the nations with large deployment: about 0.5 kilowatt per capita.

The U.S. was estimated to produce about 56 billion kilowatt-hours of solar electricity in 2016, out of roughly 4,000 billion kilowatt hours of electricity from all sources. For Greece, Italy, and Germany, solar electricity production accounted for about 7 percent of national electricity production from all sources.

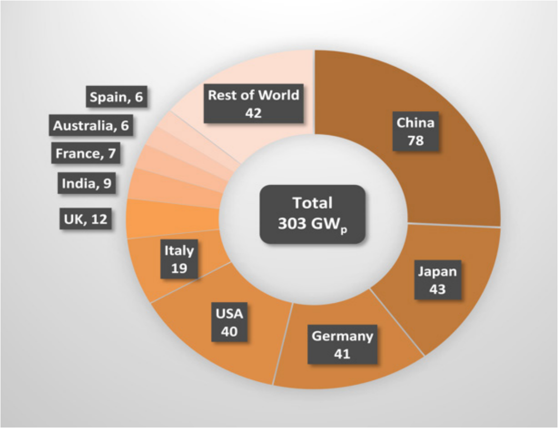

Figure 3.6: Cumulative installed solar PV capacity, in peak-gigawatts, for the ten countries having more than five peak-gigawatts (GW$_p$) of capacity by the end of 2016. 1 gigawatt = 1000 megawatts = 1,000,000 kilowatts. Data: International Energy Agency, Photovoltaic Power Systems Program, IEA PVPS Snapshot 2017: http://iea-pvps.org/index.php?id=trends0.

---

[6] https://www.worldenergy.org/wp-content/uploads/2016/09/Variable-Renewable-Energy-Sources-Integration-in-Electricity-Systems-2016-How-to-get-it-right-Executive-Summary.pdf, table on p. 2, cited also in Footnote 3.

Deployment by U.S. State

Figure 3.7 breaks down the installed solar PV capacity in the U.S. by the state in which it is installed. California is responsible for nearly half of current installed capacity. That North Carolina is in second position, New Jersey in fifth, and Massachusetts in seventh – despite being neither especially large nor especially sunny – is a reflection of consistent state-level policy support.

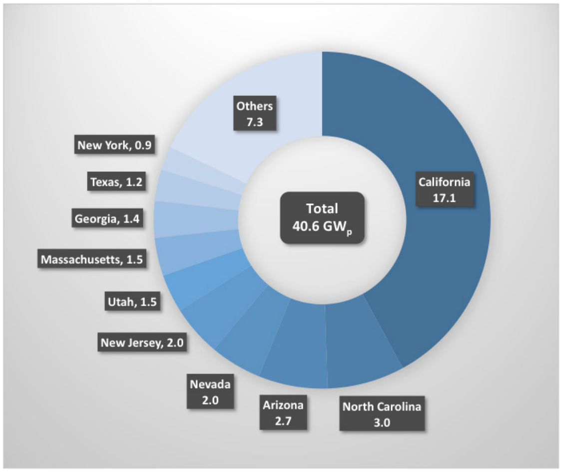

Figure. 3.7: Cumulative installed solar PV capacity at the end of 2016 by U.S. state, in peak-gigawatts (GW_p). 1 gigawatt = 1,000 megawatts = 1,000,000 kilowatts. Data: Solar Energy Industry Association, http://www.seia.org.

Land Required to Produce Electricity from Sunlight

The route from sunlight to electricity using solar cells can be compared to another route, the "biopower" route, where sunlight enables the growth of vegetation (crops, grasses, trees), which is then harvested and converted into electricity. The land demand to convert sunlight to electricity directly with solar cells is far less than the land demand for the "biopower" route. On the other hand, competition for land is often fierce, and solar power requires dedicated land, while biopower is compatible with simultaneous use for other purposes. Dedicated land for solar power can conflict with urban green space and, on a larger scale, with demand for national parks and wilderness.

Solar power requires less land than biopower because the efficiency of conversion of sunlight to commercial energy is so much higher for solar power. A reference efficiency for solar panels today is 20 percent. A conversion efficiency of even 1 percent represents an extremely high yield for biomass, relative to actual yields in crops and forests. The two conversion efficiencies – 20 percent for a representative solar panel and less than 1 percent for biomass – mean that biomass requires at least 20 times more land as solar panels to produce the same amount of energy. (The comparison is simplistic, to be sure, since biomass requires further processing to be useful, but on the other hand biomass not only collects solar energy but also stores it for use at a later time.) The significantly smaller land requirements for solar energy production are a fundamental reason why solar electricity has the potential to transform the global energy system.

It is instructive to calculate how much land fully devoted to PV solar power would be required to meet the entire electricity demand of a specific geographical region. For simplicity, we ignore solar collection on the roofs of residential and commercial buildings and work out the amount of land required to meet total U.S. electricity demand from horizontal stationary solar panels sited near Phoenix, Arizona. The amount of electricity consumed in the U.S. in 2015 was about 4,000 billion kilowatt-hours. Solar energy falls on Phoenix at an average rate of approximately 6.5 kilowatt-hours per square meter of land per day, or 2,400 kilowatt-hours per square meter of land per year (see Figure 3.1). Thus, 480 kilowatt-hours would be produced each year from each square meter of stationary horizontal power in Phoenix, assuming 20-percent efficiency panels. Dividing 4,000 billion kilowatt hours by 480 kilowatt-hours per square meter, 8.3 billion square meters (8,300 square kilometers, or 3,200 square miles) of panels near Phoenix could collect this much energy.

We could double this area to take into account gaps between the rows and to include supporting infrastructure beyond the site. The result, 6,400 square miles (about 1/600th of the area of the U.S.), is roughly the size of metropolitan Phoenix and is compared with

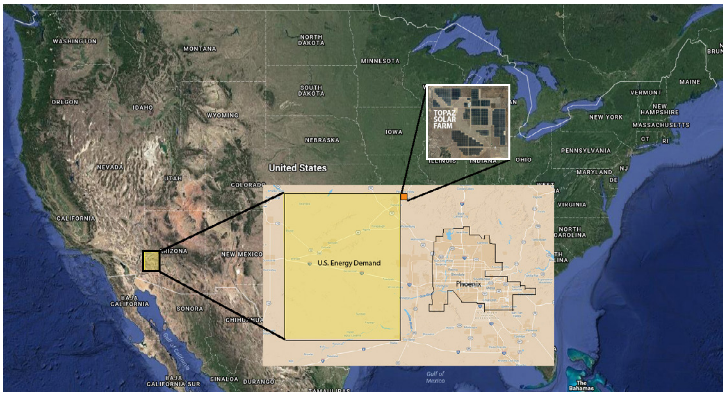

Figure 3.8: The area of land outside Phoenix, AZ, about 6,400 square miles, required to generate the entire U.S. electricity demand if fully devoted to solar power, is shown in position on a map of the U.S. and in an inset as a rectangle adjacent to the city boundaries of Phoenix. (For assumptions, see text.) The red rectangle, shown to scale in this inset, is expanded in a second inset to reveal the land required for the Topaz Solar Farm in San Luis Obispo County, California.

Phoenix on a map in Figure 3.8. Additional dedicated land would be required for energy storage facilities and transmission corridors. Note that if the solar cell efficiency were 25 percent instead of the assumed 20 percent, all of these area calculations would be reduced by one-fifth. For example, our estimate of 6,400 square miles would become 5,100 square miles.

For comparison, Figure 3.8 also shows the 550-megawatt Topaz Solar Farm in San Luis Obispo County, California, one of the largest solar farms in the U.S., which went on line in 2014.

The calculated land area meets the *current* demand for electricity, but not the additional demand that would be required if the U.S. economy were completely electrified – where cars run on batteries, houses are electrically heated, and all industrial processes are powered by electricity. Currently, about 40 percent of U.S. primary energy is used for electricity; thus, as a very rough estimate, the required land area to power a totally electrified U.S. economy might be 2.5 times the area calculated in Figure 3.8. This figure would be approximately 16,000 square miles, which is roughly the size of Maryland.

C. Solar Energy Projects

Utility, Mid-scale, and Residential Projects

Commercial solar power is arriving at all sizes at once. The usual distinction for solar projects is between 1) *utility* projects that deliver power directly to a utility (sometimes called projects "in front of the meter"), and 2) *distributed generation* projects ("behind the meter"), where a portion of the produced electricity is consumed on site.

We have found it useful to divide distributed generation group into residential projects (a billing category widely used by the industry) and *mid-scale* projects, which are all distributed-generation projects that are not residential projects. Mid-scale projects are almost always larger than residential installations but smaller than utility arrays. *Commercial* projects (another billing category) are included in the mid-scale category: these are the projects on rooftops of warehouses and on other private property. Also in the mid-scale category are the many installations on public land, including those on or around schools, hospitals, parks, municipal centers, and parking structures.

Residential and utility projects have recognizable archetypes, seen in Figure 3.9: a residential installation of rooftop panels (left) and a project comprising fields of panels delivering power directly to utilities (right). Mid-scale projects, like Princeton University's project (bottom), by contrast, are rarely included in the visual imagery of solar power.

Photo by Tom Grimes.

Figure 3.9: A representative residential PV installation (upper left, 10 kilowatts, estimated), Solarpark Meuro, the largest installation in Germany (upper right), more than 150,000 kilowatts, not all shown), and the Princeton University mid-scale project (bottom, 5,400 kilowatts). Source: https://www.habdank-pv.com/en/portfolio-item/soft-soil.

Solar PV Projects in New Jersey

Mid-scale projects dominate the deployment of solar energy in New Jersey. They account for 58 percent of New Jersey's solar capacity, even though they account for only about one tenth of all projects. Utility projects account for 23 percent of capacity, and residential projects account for the remaining 19 percent (even though residential projects constitute almost nine tenths of all projects). These findings come from a database of nearly all of New Jersey's solar PV projects, maintained by the New Jersey Board of Public Utilities.[7] The data are displayed in Figure 3.10 and further reported in Table 3.1. Projects smaller than 10 kilowatts contribute roughly one-fifth of the total capacity, those between 10 kilowatts and 1,000 kilowatts add another two-fifths, and those larger than 1,000 kilowatts contribute the remaining two-fifths. Half of New Jersey's solar PV projects have a capacity below 8 kilowatts.

Even New Jersey's largest projects are far smaller than the largest utility solar projects found in the southwestern U.S. The largest single project in the New Jersey database is a 19.9-megawatt utility project, whereas in 2016 there were six solar projects in the U.S. whose capacity exceeded 300 megawatts – three in California, two in Nevada, and one in Arizona.[8]

---

[7] The Board of Public Utilities database catalogs all projects eligible to receive New Jersey's solar renewable energy credits (SRECs). As of February 29, 2016, the database included more than 40,000 projects totaling more than 1,600 megawatts of generating capacity.

[8] Globally, there are 14 projects whose capacity exceeded 300 megawatts. Two of the world's three largest solar projects are in India and the third is in China. See https://en.wikipedia.org/wiki/List_of_photovoltaic_power_stations.

	Residential	Mid-Scale	Utility	Total
Number of Projects	38,859	4,827	136	43,822
Percent of all Projects	88.7	11.0	0.3	100
Capacity (Megawatts)	314	952	377	1,643
Percent of All Projects	19	58	23	100
Median (Kilowatts)	7	50	1,246	8
Mean (Kilowatts)	8	197	2,772	37

Table 3.1: Summary statistics for all solar projects in New Jersey. "Residential" and "Utility" are categories used in the database; utility projects provide power directly to a utility. "Mid-scale" groups together all other categories. Source: New Jersey Board of Public Utilities, http://www.njcleanenergy.com/renewable-energy/project-activity-reports/project-activity-reports.

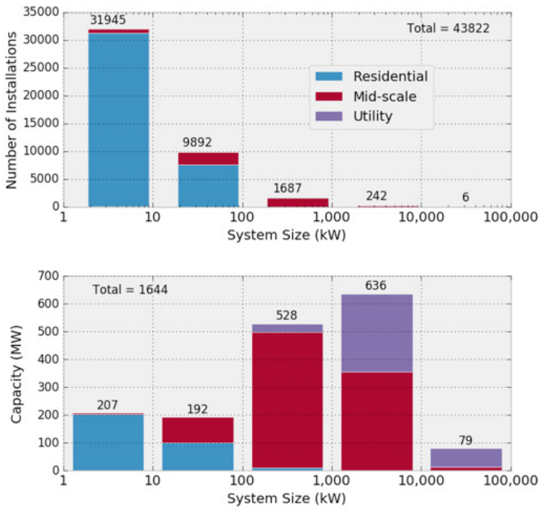

Figure 3.10. Contributions to the total solar generating capacity of New Jersey, as of February 29, 2016, binned by size (capacity). The more than 40,000 installations (top panel) overwhelmingly have a capacity less than 10 kilowatts, but the roughly 1.6 million kilowatts of total capacity (bottom panel) is dominated by large facilities. Numbers above each bar are totals for that bar; totals are in thousands of kilowatts (megawatts) in the bottom panel. Bars and segments of bars are colored to sort projects into our three categories. *Source:* New Jersey Board of Public Utilities, http://www.njcleanenergy.com/renewable-energy/project-activity-reports/project-activity-reports. Data as of February 29, 2016.

D. Distributed Generation

Large solar plants, like those in the deserts of the southwestern U.S., fit nicely with century-long trends: the size of the individual power plant of all kinds has increased steadily, as has the distance between the site of power production and the site of electricity use. By contrast, residential and mid-scale solar power production reverses historical patterns. A household can meet its annual power requirements from a collector on the roof and trade power with its utility, buying or selling depending on whether household demand exceeds or is less than the collector's supply. Several households can link themselves together and locate their collector in a nearby field, creating a solar power system with its own microgrid. Private companies and public institutions of all kinds can do the same. In each case, a specialized business can own the collectors and rent them to the households and companies, achieving economies of scale. And in each case, the project can be augmented by electricity storage: add enough batteries and any of these entities can disconnect from the grid entirely. This is the new world of "distributed generation" of solar power.

"Distributed generation" is a general concept. It describes not only dispersed solar production facilities but also dispersed electricity production from other energy sources, notably dispersed production of electricity from natural-gas. In principle, distributed generation can take over the entire electricity system, displacing central station power entirely. More credible is a future grid that combines large amounts of both distributed power and centralized power. Such a grid can be more resilient and flexible than a grid consisting only of large power plants, especially if the sites for distributed components are chosen so as to reduce congestion and relieve bottlenecks in transmission of bulk power. Distributed generation also provides back-up power when natural disasters or hacking produce widespread outages at centralized facilities.

A major constraint on the expansion of such a mixed system becomes the grid itself, which must be developed in new ways. The grid must continue to provide reliable electricity service; electric utilities often affirm that reliability is their most important objective. The entire infrastructure needs to remain reliable, including the distribution system of power lines running down every street, even as new sources of electricity are introduced at the outermost branches of the distribution system, leading to two-way flows of electric current on lines that were designed for one-way flow. When a decentralized generator fails, the grid must provide an alternative.

Distributed electricity storage is key to the future of distributed energy. The first solar power projects in homes and on farms came with banks of batteries, enabling a user to become completely independent of any grid, but these early systems were largely supplanted when grid connection was offered on favorable terms. Now, once again, distributed solar electricity storage is being offered in combination with distributed power generation, and the two are being tied to each other and to the grid by “smart” information sharing. Down this road, decentralized solar power becomes dispatchable, back-up *by* the grid becomes less demanding, and back-up *of* the grid becomes more credible.

Community Solar Power

Constituting a new class of mid-scale projects are “community solar” projects (also called “shared solar,” “solar gardens,” and “community distributed generation”). The objective of a community solar project is to expand solar energy access to renters, homeowners with unsuitable roofs, low-income and moderate-income consumers, and others who cannot otherwise “go solar.” A community solar project could be organized by a solar company, a local organization, or some other entity; its participants are “subscribers” who purchase fractions of the project’s installed capacity or fractions of its electricity production. The project’s solar power need not be produced on the premises of any of the subscribers, and it need not be delivered to the subscribers. States with community solar programs require utilities to credit all participants for the solar power their portion of the project produces, lowering their monthly utility bills. Unused credits typically roll over to the following month, but in some states credits expire at the end of the calendar year. Colorado, Massachusetts, Minnesota, and New York have been pioneers in the development of community solar projects. They and eleven other states, as well as Washington, D.C., have enacted policies authorizing community solar programs.[9]

E. Balance of System

A PV system is much more than just solar panels. We use the term “balance of system” to refer to everything related to a solar project *other* than the panels – both non-panel hardware and so-called “soft costs.”[10] Non-panel hardware includes panel mounts, transformers, wiring, enclosures, and the inverters which convert electricity from direct current (DC) to alternating current (AC). Soft costs include the costs associated with land, customer acquisition, financing, permitting, property taxes, installation labor, and installer profit. Balance of system costs do not include costs for integrating a project into an electricity grid, such as associated electricity storage or back-up power; grid-integration costs are treated extensively in Article 5.

Both the PV panel and the balance of systems have become steadily less expensive, as seen in Figure 3.11, which shows representative costs for 2009 through 2016 for residential, commercial, and utility installations, for projects modeled by the National Renewable Energy Laboratory. (The “commercial” category is roughly equivalent to our Distillate’s “mid-scale” category.) Panel costs, which are presumed to be the same for the three kinds of projects, are about three times more expensive at the beginning of the period than at the end. Balance-of-system costs also fall, although not as dramatically.

Also seen in the most recent bars in Figure 3.11, balance-of-system costs now dominate total costs for residential and commercial projects and account for about half of total costs for utility projects. And within balance-of-system costs, soft costs have become the major component of balance-of-system costs, especially for smaller projects. Figure 3.12 elaborates this argument with an independent estimate of balance-of-system costs, where 36 percent are hardware costs and 64 percent are soft costs, and the soft costs are distributed into nine categories, none of them

---

[9] http://www.communitysolaraccess.org/wp-content/uploads/2016/03/CCSA-Policy-Decision-Matrix-Final-11-15-2016.pdf.

[10] An alternative use of the phrase, “balance of system,” restricts its meaning to non-panel hardware.

dominant. One of the reasons that the hardware component of the balance-of-system cost has fallen, when measured in dollars per peak-watt of capacity (the unit used in Figure 3.11), is that solar cells have become more efficient. Less balance-of-system hardware is required for the same amount of electricity produced, even when the exactly the same hardware is used to mount and connect the panel.

Soft costs are being steadily reduced. Strategies internal to the solar industry to reduce these costs include standardization of hardware, workforce training, and financial risk management. Local governments are also contributing, to the extent that they modify local land use and zoning policies to encourage (or at least not inhibit) solar projects and simplify the acquisition of construction permits.

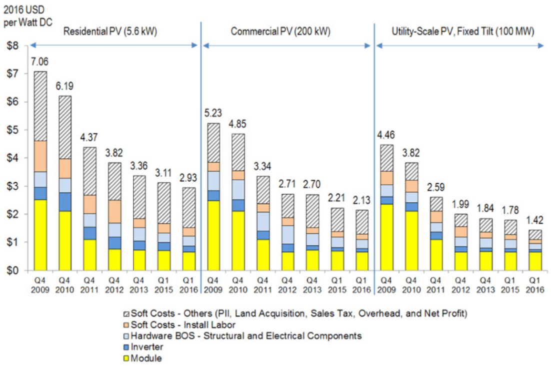

Figure 3.11: Costs for representative residential, commercial, and utility solar projects modeled by the National Renewable Energy Lab (NREL). Q1 and Q4 are a year's first and fourth quarters, respectively. PII is "Permitting, Inspection, and Interconnection." BOS is "Balance of System." Source: NREL, "NREL report shows U.S. solar photovoltaic costs continuing to fall in 2016." http://www.nrel.gov/news/press/2016/37745.

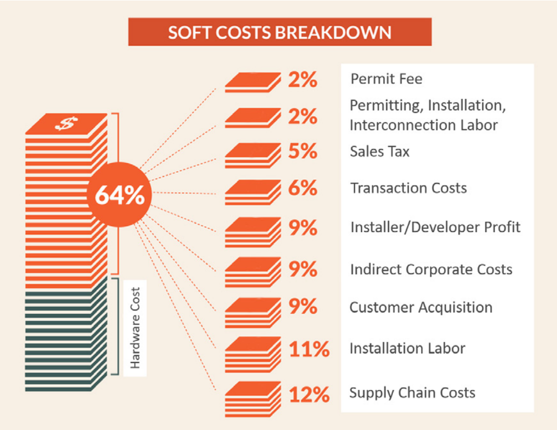

Figure 3.12: A representative distribution of the "soft-cost" component of balance-of-system costs. Source: U.S. Department of Energy, Office of Energy Efficiency & Renewable Energy: http://energy.gov/eere/sunshot/soft-costs

Fixed Panels versus Tracking Panels

There are two strategies for collecting sunlight on a flat panel. The panel can be placed on a rigid mount, or it can be placed on a movable tilted frame. Many fixed panels lie flat on the roofs of buildings, their orientation and tilt dictated by the roof's orientation. Other fixed panels are mounted on the ground, in which case the orientation and tilt can usually be freely chosen. In the northern hemisphere a typical ground-mounted fixed panel will be tilted so that its north edge is higher than its south edge and will be oriented due south, thereby benefiting more from the path of the Sun through the sky than if the panel were lying flat.

The strategy of moving a panel during the day is called "tracking," because the panel tracks the Sun's path through the sky. Tracking adds initial costs and maintenance costs, but tracking results in greater amounts of solar energy striking the panel. The most expensive tracking, "double-axis tracking," maximizes solar collection by keeping panels perpendicular to the Sun throughout the day, every day of the year. This strategy requires the mount to be able to rotate around two axes, so as to change both its east-west orientation and its tilt relative to the horizon.

More common is "single-axis tracking," where the panel rotates around a fixed axis that has a single orientation throughout the year. The axis of rotation for single-axis tracking is usually horizontal, resulting in a panel that moves like a seesaw and is horizontal at noon. The axis can also be vertical, resulting in a panel that is vertical and (in the northern hemisphere) faces due south at noon. Still a third option is for the axis to be oriented at an angle between horizontal and vertical.

Moving clockwise from the top-left, the three photos in Figure 3.13 show panels mounted with a fixed tilt, two-axis tracking, and one-axis tracking. The orientation of the axis of the single-axis tracking system is north-south at a small angle relative to horizontal, resulting in panels that at noon face south at that same angle.

The cost of land can be a determining factor in choosing between fixed panels and tracking panels. In general, tracking panels require extra land (for the same amount of solar power capacity) relative to fixed-axis panels, because tracking panels cast larger shadows. Expensive land can drive the choice toward fixed panels over tracking panels or toward tracking panels placed closer together (accepting more shadowing).

Costs Related to Voltage and Current

Even when residential, mid-scale, and utility installations utilize the same PV panels, the optimal designs for the management of the electricity output can be very different. Panels on the roof of a residence are easily

Figure 3.13 Panels in a fixed array at Eastern Mennonite University, Harrisonburg, Virginia (top left); with double-axis tracking in Toledo, Spain (top right); and with single-axis tracking in Xitieshan, China (bottom). Sources:
Top left: https://upload.wikimedia.org/wikipedia/commons/c/c5/Eastern_Mennonite_University_Solar_Array.jpg
Top right: https://upload.wikimedia.org/wikipedia/commons/8/8f/Seguidor2ejes.jpg
Bottom: By Vinaykumar8687 - Own work, CC BY-SA 4.0, https://commons.wikimedia.org/w/index.php?curid=35401850.

linked together at low voltage to feed power either to the building, or, when the panels produce excess power, back through the residential meter to the utility's low-voltage distribution system. By contrast, utility-scale solar arrays are stepped up to high grid voltages in order use the utility's mid-voltage and high-voltage transmission lines.

As for mid-scale projects, grid connection presents more individualized challenges. One general observation is that any project exceeding 100 kilowatts of capacity requires a significant investment in inverters to convert the DC power produced by the modules to the AC power required by the grid. The cost of these inverters has not fallen as quickly as the cost of modules.

F. Building-integrated Photovoltaics

Balance-of-system costs can become opportunities for systems design. In the building sector, roof and façade not only can support attached solar panels, but can actually be *constructed* of solar panels—an approach known as "building-integrated PV." While a number of companies have integrated solar modules into roof shingles, Tesla's "solar roof" recently popularized the technology (see Figure 3.14, where two other examples of structural PV are also shown).

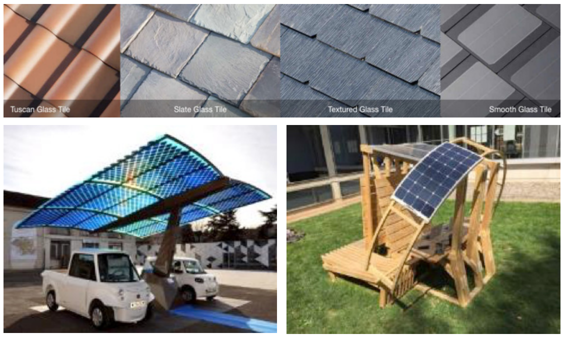

Figure 3.14: Examples of building-integrated solar PV. Top: Tesla's "solar roof" offerings. Lower left: A car shade made of PV collectors. Lower right: A solar umbrella that tracks the Sun so that the table's surface is always shaded. Sources: www.tesla.com/solar (top); https://en.wikipedia.org/wiki/Photovoltaic_system#/media/File:Ombri%C3%A8re_SUDI_-_Sustainable_Urban_Design_%26_Innovation.jpg (lower left); Meggers CHAOS lab (lower right).

Another class of novel specialty applications of PV cells features lightweight, colorful, and semi-transparent photoactive materials and devices to enhance aesthetic value while also generating electricity. An example is shown in Figure 3.15: the installation of dye-sensitized transparent solar cells in the façade of the SwissTech Convention Center in Lausanne, Switzerland. These cells have a conversion efficiency of only a few percent, but

Figure 3.15: SwissTech Convention Center installation of dye-sensitized solar cells in a large glazed wall. The cells help prevent overheating in the afternoon while simultaneously generating electricity. Source: © FG+SG fotografie de architectura, http://www.archdaily.com/519434/epfl-quartier-nord-swisstech-convention-center-retail-and-student-housing-richter-dahl-rocha-and-associes/53a84e94c07a80c112000101-epfl-quartier-nord-swisstech-convention-center-retail-and-student-housing-richter-dahl-rocha-and-associes-photo.

they serve additional functions as shades and filters. In addition, the cells maintain a relatively high efficiency in diffuse light, making them well suited for vertical surfaces.

Similarly, at Princeton University, the panels containing monocrystalline silicon solar cells on top of the Frick Chemistry Laboratory (Figure 3.16) provide both shade and energy. In designing the building, architects and engineers recognized that shading surfaces function best when aligned perpendicular to the Sun's rays, as is also true of solar panels. Hence, glass-mounted solar panels were used to shade the building's central atrium, intercepting the majority of glare-inducing intense sunlight while effectively letting light through between the cells. The principal justification for the panels is aesthetic interest, not the electricity they generate, which is only about one percent of the building's electricity, less than the panels save by reducing the need for cooling. But the incremental cost may also have been minimal, perhaps not exceeding what would have been spent for an internally integrated shading system.

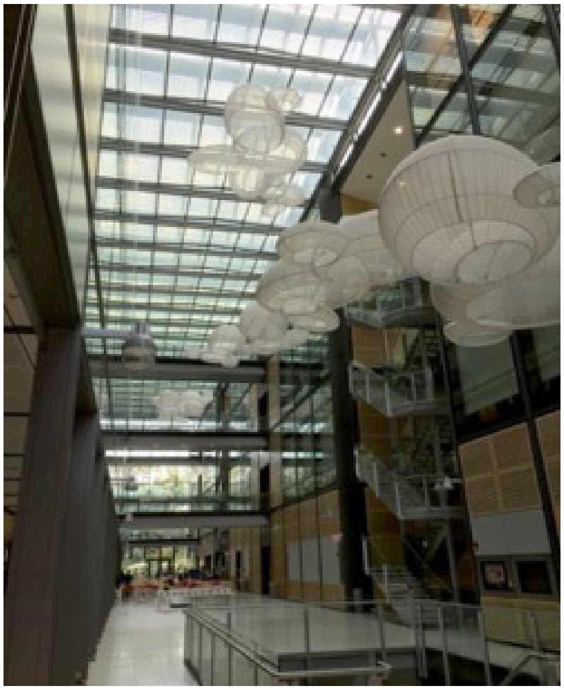

Figure 3.16: Frick Chemistry Laboratory at Princeton University, where solar panels above the atrium are used both for shading and for electricity generation. Source: Forrest Meggers.

Article 4: Solar Cell Technology

Article 4 is a survey of solar cell technologies. Eleven solar technologies are reviewed, five of them currently available and six of them still in the laboratory. A scoring system is introduced that highlights many of the issues that drive solar cell development. An underlying question is whether the current dominance of the crystalline silicon solar cell will be a permanent feature of the solar cell market for the indefinite future.

A. The Solar Cell

Light consists of discrete particles, called photons, each carrying a tiny amount of energy. A photovoltaic (PV) cell converts incident photons into electricity. If a PV cell is powered by incident sunlight, it is termed a solar cell.

At the core of most common solar cells is an interface between two semiconductors with different electronic properties. On one side of the interface is an “n-type” semiconductor with excess electrons (electrons carry *negative* charge, hence “n-type”). On the other side is a “p-type” semiconductor, which has a deficiency of electrons (equivalent to an excess of *positively* charged counterparts to the electron, called “holes”). The two materials create an internal electric field at the interface, which is called the “p-n junction.”

The device has a “band gap,” a specific amount of energy. That energy is the minimum needed to excite an electron into a state in which it can move through the device in the presence of an electric field. Upon the absorption of a photon whose energy exceeds the band gap, an electron is promoted across the band gap and makes what is normally a forbidden transition from a lower energy band (the valence band or ground state) to a higher energy band (the conduction band or excited state). At the same time, a “hole” (in effect, a missing electron) is created in the valence band. The effectiveness of a solar cell arises from the fact that the energy of the photon can be converted to electricity when the internal electric field separates the electrons and holes from one another and directs them to the two contacts of the device (through the cathode and anode, respectively).

In general, electricity is produced only when a photon with more energy than the band gap strikes a solar cell. Thus, there are materials where a blue photon is energetic enough to drive an electron across the band gap but a red photon is not. The excess energy carried by the blue photon, relative to the amount of energy that is sufficient to cross the band gap, becomes heat.

Solar Cell Efficiency

The efficiency of a solar cell is defined as the percentage of incident solar power that is converted to electric power. The efficiency is measured under laboratory conditions that mimic peak conditions, where the Sun is directly above the solar cell and high in the sky, and the day is clear. Efficiency is a solar cell’s most important attribute, because higher efficiency translates into smaller facilities on less land.

The electric power output is the product of the photocurrent and photovoltage of the solar cell. The photocurrent is directly proportional to the number of solar photons that an absorber is able to collect, while the photovoltage is determined by the semiconductor’s band gap. A material with a larger band gap provides a greater voltage but delivers less current because it absorbs less of the solar spectrum. Accordingly, there is an optimal band gap where the maximum output of solar electricity can be achieved, determined by the specific distribution of energies in the photons of sunlight. At that band gap, a solar cell with a single junction (the most common type), has the maximum possible efficiency. That efficiency is about 33 percent. In practice, the highest efficiencies achieved for single-junction cells are close to this limit: 28.8 percent for gallium arsenide and 26.6 percent for crystalline silicon.

Considerably higher efficiencies can be reached with a multijunction solar cell, where different solar cells are integrated together. A typical multijunction cell has two to five absorbers, each having a band-gap with a different amount of energy, so that complementary portions of the solar spectrum can be harvested. Multijunction solar cells are more expensive to fabricate than single-junction cells. As a result it is often worth enhancing their efficiency still further by using concentrators that intensify the strength of the sunlight that falls on these cells. The record efficiency to date is 38.8 percent for a multijunction cell without concentration, and 46 percent for a multijunction cell receiving a solar input concentrated more than 100 times.

	Efficiency	Abundance	Compatibility	Stability	Manufacturability	Versatility
Today's Technologies						
mono-Si	+2	+2	+1	+2	+2	-2
poly-Si	+1	+2	+1	+2	+2	-2
a-Si	-1	+2	+2	-1	+2	+1
CdTe	+1	-2	-2	+2	+2	+1
CIGS	+1	-1	-1	+2	-1	+1
Technologies on the Frontier						
GaAs	+2	+1	+1	+2	-2	+1
CZTS	-1	+2	-1	?	-1	+1
OPV	-1	+1	+2	?	+1	+2
DSSC	-1	+1	+1	+1	-2	+1
QD	-1	+1	-1	?	?	+1
Perovskite	+1	+1	-2	?	?	?

Table 4.1: Scores (on a four-point scale) of five current and six frontier solar technologies with respect to six attributes. The row labels are names of cells. mono-Si: monocrystalline silicon. poly-Si: polycrystalline silicon. a-Si: amorphous silicon. CdTe: cadmium telluride. CIGS: either copper indium gallium diselenide or copper indium gallium disulfide. GaAs: gallium arsenide. CZTS: either copper-zinc-tin-sulfur or copper-zinc-tin-selenium. OPV: organic photovoltaic. DSSC: dye-sensitized solar cell. QD: quantum dot. Perovskite is not abbreviated.

Will crystalline silicon ever lose its dominance?

Monocrystalline silicon and polycrystalline silicon, the two main crystalline silicon technologies, together account for about 90 percent of today's global installed solar power capacity. Will another solar cell ever beat crystalline silicon in the PV market?

Table 4.1 presents our attempt to benchmark eleven other solar technologies, five which we consider "today's technologies," and six which we place on the research frontier. We consider only single-junction cells. We compare these eleven cells across six metrics: efficiency, element abundance, compatibility with public health and the environment, stability, manufacturability, and versatility in deployment options. We use a four-point scale: +2, +1, -1, and -2 (approximating very poor, poor, fair, and good). We opt for question marks in a few instances. Below, we discuss first the six metrics and then the eleven cells.

Efficiency

The most heralded performance index of a solar cell is its efficiency. Raising the efficiency of a solar cell, other things being equal, lowers the cost of a project. Fewer structural supports, less installation labor, and less outlay in many other areas can produce the same amount of electricity when the cell efficiency increases.

Timelines of the highest efficiency achieved by each of the eleven technologies are plotted in Figure 4.1, which is a simplification of a widely cited figure prepared by the National Renewable Energy Laboratory and regularly updated on its website.[11] In Table 4.1, we assign the +2 score only to the gallium arsenide and the monocrystalline silicon cells. (The gallium arsenide cell, as discussed further below, is used on spacecraft but has been too costly for wide use elsewhere.) The other nine cells are scored either +1 or -1 (-2 is not used).

---

[11] Another excellent resource for following progress in the performance of the various solar technologies is the journal, *Progress in Photovoltaics*, which periodically publishes "Solar-cell efficiency tables" for cells and modules.

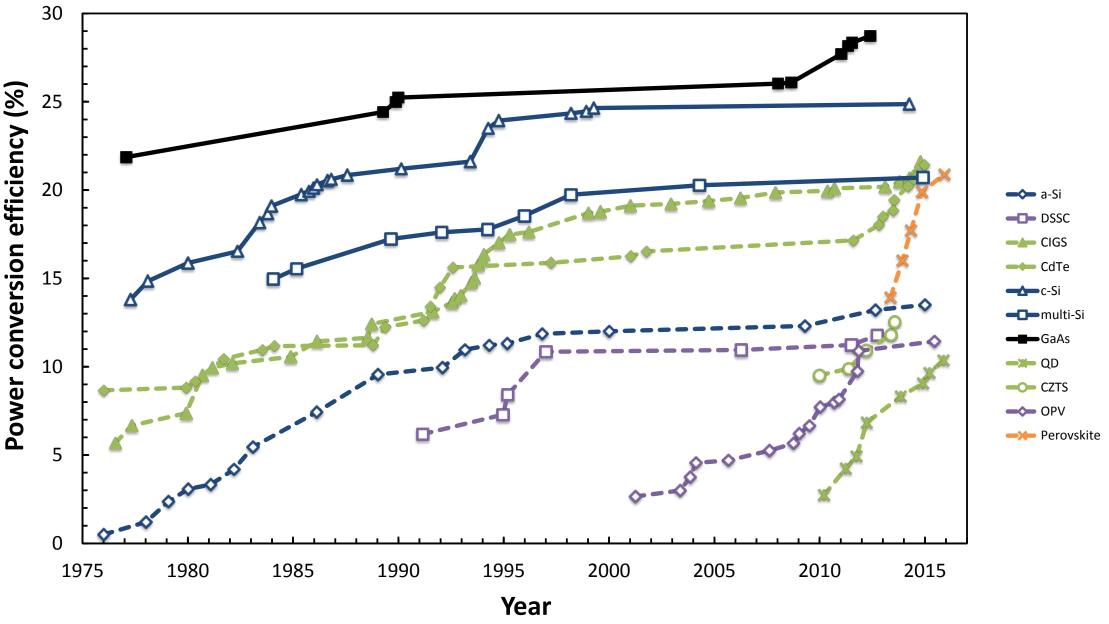

Figure 4.1: Timeline of the certified efficiency values of eleven solar-cell technologies. Source: "Best research-cell efficiencies," a chart from the National Renewable Energy Laboratory, http://www.nrel.gov

Abundance

Solar energy systems face resource limitations, especially when they rely on relatively rare elements. A point of reference for the scoring of scarcity in Table 4.1 is the abundance of an element in the Earth's crust. Consider the seven elements in the two thin-film non-silicon inorganic solar cells discussed under "today's technologies" above: the CdTe cell contains cadmium and tellurium and the CIGS thin-film cell contains copper, indium, gallium, selenium, and sulfur. Add an eighth element, ruthenium, which currently is used in DSSCs. The crustal abundance of three of these eight elements is substantially greater than one part per million: the abundances are about 400, 60, and 20 parts per million for sulfur, copper, and gallium, respectively. Of the other five, tellurium and ruthenium are substantially more scarce than cadmium, indium, and selenium: the abundances of both tellurium and ruthenium are one part per billion (about the same as the abundance of platinum), while the other three, roughly equally abundant, are between 50 and 250 times more abundant than tellurium and ruthenium and about 1,000 times more scarce than sulfur, copper, and gallium.[12] The scarcity of tellurium is a substantial obstacle to the expansion of CdTe cells, even though the amount of photo-active material required for a thin film is very small. The scarcity of indium is one of the reasons why alternatives to the CIGS cell have been sought.

To be sure, crustal abundance by itself is only a weak guide to scarcity, because mineral distribution is of course not uniform over the Earth's crust. Moreover, scarce elements are often produced not directly but as byproducts of the mining of more common elements with which they are associated (cadmium and indium are extracted primarily with zinc; selenium and tellurium with copper), resulting in much lower costs for these elements than the costs of production from dedicated mines.

Compatibility

The *compatibility* index is intended to reflect hazards both to public health and the natural environment. Our compatibility score takes note of only lead and cadmium, two heavy metals that are particularly toxic. Indeed, the replacement of nickel-cadmium batteries by other kinds of nickel batteries was in part driven by the desire to avoid cadmium's toxicity. Similar pressure drove lead compounds out of paint and tetra-ethyl lead out of gasoline. A concern for toxicity arises throughout a material's life-cycle, starting with the health of miners and workers in fabrication facilities; then during use when, for example, winds carrying sand can ablate and disperse the material; and, finally, during the disposal process. A prominent use of lead solder in a device affects our compatibility score slightly (it leads to a score of +1 instead of +2). Extensive use of cadmium or lead in the cell itself is given greater importance and results in a score of either -1 or -2, depending on how much metal is involved.

---

[12] https://en.wikipedia.org/wiki/Abundance_of_elements_in_Earth%27s_crust.

Toxicity has not inhibited extensive use of lead, notably in the lead acid battery, which accounts for a large fraction of the lead in industrial products.[13] The need to manage battery lead is partially responsible for the "secondary" lead acid battery industry, which rebuilds lead acid batteries in dedicated facilities that use lead exclusively taken from batteries collected from users, rather than newly mined lead. Moreover, the environmental mobility of lead in the three chemical forms found in a battery (metallic lead, lead oxide, and lead sulfate) is small. As a result, the lead battery system, in principle, can produce relatively little toxic environmental impact.[14] Comparable solutions may enable the use of lead in commercial solar cells.

Stability

Stability reflects the average loss of performance over time for real-world installations. For several of the frontier technologies conclusive data are not available, and for most of the other technologies is not an issue.

Manufacturability

Manufacturability refers to the current cell fabrication technology. Reasons for the scoring vary and are presented for each cell below.

Versatility

In scoring "versatility" our dominant criterion is whether the cell is crystalline or a thin film. Thin films are inherently more versatile because, relative to crystalline cells, they can be shaped and can weigh less, permitting use in a wider variety of applications. Eight of the eleven technologies in Table 4.1 are based on thin films.

B. Today's Technologies

Five technologies in Figure 4.1 and Table 4.1 are classified as "today's technologies." Three of these (shown in blue in Figure 4.1) are silicon-based: monocrystalline silicon (also called single-crystalline silicon), polycrystalline silicon (also, multi-crystalline), and amorphous silicon. The other two are based on cadmium telluride and copper indium gallium diselenide cells (shown in green).

Monocrystalline and Polycrystalline Silicon Cells

In electronic devices, the silicon is extremely pure: the silicon is called "nine nines" silicon (99.9999999 percent silicon), because less than one atom in a billion in the crystal is not a silicon atom. In crystalline silicon solar cells, the silicon can be either as pure or a little less pure.

The n-type and the p-type silicon semiconductors are formed by "doping" the silicon, that is, by introducing very small amounts of impurity atoms into silicon's crystal lattice. Silicon is made n-type by introducing an impurity atom, such as phosphorus; the one additional valence electron of phosphorus, relative to silicon, contributes an electron to the solid, making the material rich in electrons, which are negatively charged. Silicon is made p-type by introducing boron, which has one less valence electron than silicon, creating a positive charge (a hole) in the lattice.

The result is a silicon solar cell with a band gap corresponding to the photon energy of near-infrared light (light that is slightly less energetic than red light), which enables the cell to absorb the whole spectrum of visible light as well as all light more energetic than visible light. Silicon solar cells operate close to their maximum power point (maximum product of voltage and current), which for a typical crystalline silicon solar cell is a voltage of about 0.6 volts, about half the voltage of a AA battery. To reach higher voltages for practical implementation, cells are strung in series and encapsulated into what are known as modules or panels.

The monocrystalline silicon cell is based on a single silicon crystal, whereas the polycrystalline silicon cell contains numerous crystalline grains, each a few centimeters in size. The crystalline solar cell is typically a square, 15 centimeters by 15 centimeters – about the size of a compact disk case – and the module contains dozens (typical values are 60, 72, or 96) of individual cells. These modules are then connected in series and parallel to other modules to form an array, and power conditioning elements are incorporated. The two main power conditioning elements are 1) transformers, which change the voltage, and 2) inverters, which change the direct current (DC) produced by the cell into the alternating current (AC) required by the user.

Monocrystalline silicon is more costly to produce than polycrystalline silicon, but the monocrystalline cell's efficiency is higher, so there is a trade-off, and both are widely produced. The loss of efficiency in polycrystalline silicon results from defects that promote the recombination of electrons and holes. The efficiency of polycrystalline silicon solar cells is less than the efficiency of monocrystalline silicon solar cells by a few percentage points: roughly 20 percent for polycrystalline silicon versus 25 percent for monocrystalline silicon (Figure 4.1). Compared with other cell technologies, these efficiencies are relatively high; in fact, monocrystalline silicon is one of two technologies earning the top score (+2) for efficiency in Table 4.1 (the other is gallium arsenide).

[13] http://www.ila-lead.org/lead-facts/lead-uses--statistics

[14] Lead contamination at homes near a closed secondary lead acid battery plant in California is a counterexample. http://www.latimes.com/local/lanow/la-me-ln-exide-cleanup-expedite-20170112-story.html.

We do not differentiate the two forms of crystalline silicon on the five attributes other than efficiency in Table 4.1. We assign +2 scores for stability and manufacturability to both cells, because the processes that produce durable cells are by now well developed. To fabricate either a monocrystalline or a polycrystalline silicon cell, a saw cuts wafers out of a large silicon crystal ingot. High temperatures are required to purify and crystallize silicon, and the energy to provide these temperatures is a major component of the total energy used to create a silicon solar panel. The polycrystalline cell can be formed into a rectangular shape in a ceramic crucible, meaning that wafers can fill the module space more efficiently than the slightly rounded monocrystalline wafers that are cut from a single cylindrical silicon ingot (Figure 4.2).

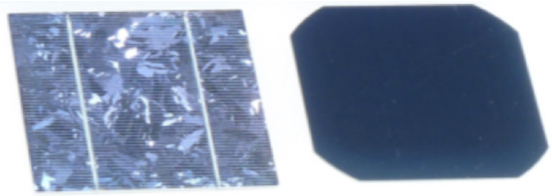

Figure 4.2: A polycrystalline silicon cell (left) and a monocrystalline silicon cell (right) Source: https://upload.wikimedia.org/wikipedia/commons/7/71/Comparison_solar_cell_poly-Si_vs_mono-Si.png

In both cases, the wafer is typically 150-200 microns thick, thick enough to support itself when undergoing the further processing required to make it into a solar cell. (A micron is one millionth of a meter.) A thick wafer is necessary, because silicon is an indirect band gap semiconductor, meaning that photons are not absorbed very strongly, so a thick layer is needed to ensure nearly complete absorption of the incident sunlight. The need for thick layers does not pose a long-term supply issue for silicon, because silicon is the second-most-abundant element (after oxygen) in the Earth's crust.

To increase the efficiency of the silicon cell, its surface is deliberately modified. An example is shown in Figure 4.3, where a complex microstructure has been created on the surface of the silicon. The result is that when a photon is reflected off of the surface it has a higher likelihood of hitting another part of the surface and being absorbed. To the eye, it looks blacker. Another strategy to improve its efficiency is to cool the cell actively (or, alternatively, to harvest its waste heat), because the performance of a silicon solar cell is degraded when hot.

We give the two forms of crystalline silicon a -2 score for versatility because, in addition to not having the versatility of a thin film, the crystalline silicon cell must be thick as a result of its poor absorption of incident sunlight; by contrast, the one other crystalline cell on our list, gallium arsenide, can be thin because it is a strong light absorber.

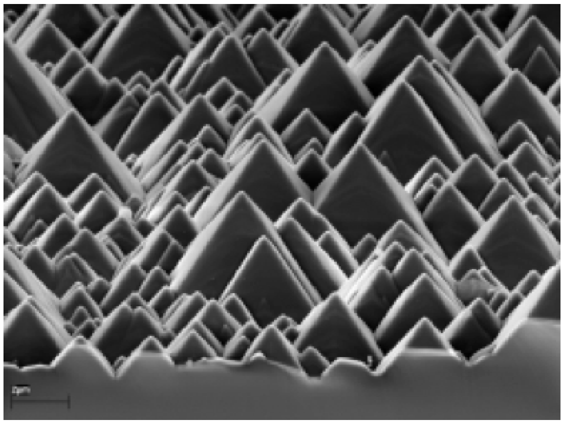

Figure 4.3: Detailed microstructure of a silicon cell whose surface has been modified to trap light and increase the probability of photon absorption. Source: https://www.flickr.com/photos/zeissmicro/10995781963

As for "compatibility" with public health and the environment, we stretch the definition to include not only the cell but the module in which it is contained. We score both types of crystalline silicon cells +1 rather than +2, because lead solder is used for the many metallic interconnects that join crystalline silicon cells to one another in a module.

Amorphous Silicon Cells

Amorphous silicon cells commanded more than 30 percent of the (albeit rather small) solar cell market in the late 1980s. While less expensive than the crystalline forms, amorphous silicon cells have much lower conversion efficiencies – roughly 10 percent, or around half of the efficiency of the two crystalline forms (Figure 4.1). Today, the market for amorphous silicon is largely confined to consumer products.

Amorphous silicon does not have a crystalline structure and cannot be fabricated into wafers. Rather, it is fabricated into thin films. Indeed, amorphous silicon was the first major commercial technology for thin-film solar cells. Thin-film technologies utilize films that are much thinner than wafers: they are at most a few microns thick and therefore need to be supported by a substrate – typically glass, metal, or plastic. The main benefits of thin-film technologies over those that are wafer-based are that less specialty material is used and fabrication throughput is higher – potentially lowering costs as well as the amount of energy required to manufacture the cells. These cells are also easier to integrate into building materials.

The performance of amorphous silicon solar cells is improved by adding hydrogen to the thin films. The low-temperature processing that leads to amorphous silicon results in silicon atoms with dangling chemical bonds, because the atoms are not positioned to form a crystal. (In crystalline silicon, by contrast, each silicon atom bonds to four other silicon atoms, and there are no dangling bonds.) Hydrogen, with a single valence electron, has the capacity to terminate (passivate) the dangling bonds, improving the electronic performance of amorphous silicon.

Unless an amorphous silicon thin film is very thin, it cannot be passivated effectively. If it is made thicker, the passivation isn't stable, and the cell loses efficiency. However, a thin film can be too thin to manufacture. Therefore, to obtain sufficient thickness without sacrificing internal efficiency, several passivated amorphous silicon thin films are often placed on top of each other in a tandem architecture. To express this inherent complexity on Table 4.1, we grade amorphous silicon as +2 for manufacturability but -1 for stability. We also give amorphous silicon a +2 score for compatibility with public health and the environment, because, compared to crystalline silicon, lead solder is used in a far more limited way.

Cadmium Telluride and Copper Indium Gallium Diselenide Cells

Cadmium telluride (CdTe) solar cells are found at nearly all of the world's solar power plants that do not use crystalline silicon cells. The CdTe plants account for roughly 10 percent of global power capacity. The 550-megawatt Topaz Solar Farm shown in Figure 3.8 uses CdTe solar cells. The copper indium gallium diselenide (CIGS) solar cell accounts for about one percent of global capacity. Variants of the CIGS cell replace all or some of the selenium with sulfur. The S in CIGS, therefore, identifies both sulfur and selenium.

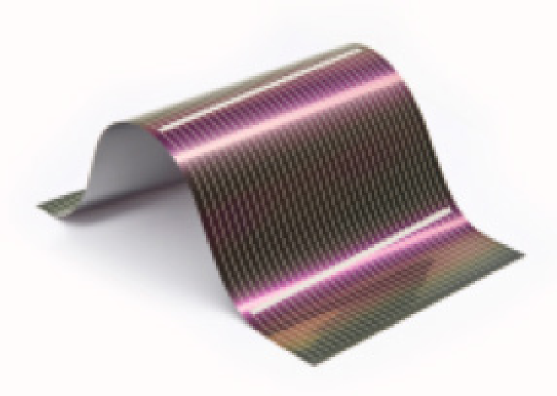

Figure 4.4: A flexible copper indium gallium diselenide (CIGS) panel. The cells form long stripes – the preferred arrangement to collect the charge built on a thin film. Source: http://materia.nl/article/innovation-thin-film-solar-cells-at-mx2016/innovation-thin-film-solar-cells-at-mx2016-1/

Like the amorphous silicon cell, the CdTe and the CIGS cells are thin-film cells. Panels made from thin-film solar cells can be flexed and deformed, as seen in Figure 4.4, which shows a flexible ribbon made of CIGS cells. Deformability makes thin-film panels attractive for installations where curved surfaces are encountered (such as on vehicles) or where the objective is to integrate the cell into architectural surfaces.

Both the CdTe and the CIGS cells are named after its p-type absorber, which in both cases is a chalcogenide semiconductor – a semiconductor formed with an atom in the group of the periodic table that includes sulfur, selenium, and tellurium, known as chalcogens. Having already achieved record cell efficiencies of just below 22 percent (Figure 4.1), and still improving, chalcogenide solar cells offer a module performance that competes effectively with polycrystalline silicon.

A useful distinction among solar cells is whether the interface is a homojunction or a heterojunction. The p-type and n-type semiconductors on the two sides of the interface of a homojunction are the same, while the two materials are different for heterojunctions. The silicon cells described above have silicon on both sides of the interface (doped in different ways), so they are homojunction cells. By contrast, both the CdTe cell and the CIGS cell have heterojunctions: the n-type material is most often cadmium sulfide (CdS).

One reason why the CdTe cell has been a strong competitor is because both its p-type semiconductor and its n-type semiconductor are binary compounds (compounds with only two elements); binary compounds can be produced industrially with better reproducibility than compounds made from three or more elements.

Both the CdTe cell and the CIGS cell are very stable. However, the extreme scarcity of tellurium and the relative scarcity of indium lead to the -2 and -1 abundance scores for the CdTe and CIGS cells, respectively. Also, the toxicity of cadmium in combination with its relative prevalence in the cell leads to compatibility scores of -2 for CdTe and -1 for CIGS. CdTe cells (made from two elements) are easier to manufacture than CIGS cells (made from four or even five elements). Both cells can be made to be rigid or flexible, enhancing their versatility.

The Learning Curve

Historically, many industries realize lower costs and therefore lower prices over time. The solar industry is no exception, as shown in Figure 4.5, which shows the "learning curve" for the crystalline-silicon solar panel (module). The average module price is plotted as a function of cumulative module shipments, and data are shown for the 40 years from 1976 to December 2016. Among the contributing factors to "learning" are steady improvements in efficient production processes and throughout the supply chain, as well as benefits from research and development and from the spillover of positive results achieved by other industries.

The fit to the data seen in Figure 4.5 (which has a logarithmic scale on both axes) corresponds to a 22.5 percent reduction in module price for each doubling in cumulative production. Between the first and last data point, cumulative production increased by a factor of about one million, from 0.3 megawatts to 300,000 megawatts: a million is approximately 20 doublings. With 0.225 as the learning curve parameter, the price should have fallen by a factor of 0.775 (1.0 minus the learning curve parameter) for each doubling, or by a factor of approximately 160 for 20 doublings. Actually, the 2016 data point lies considerably below the learning-curve line, meaning that the price fall has been even faster.

Figure 4.5 also shows a vertical line that represents the solar capacity that would be required to meet *all* of the 2016 global electricity use, which was approximately 24,000 billion kilowatt-hours. We again assume that, on average, one kilowatt of installed solar capacity will produce 1,200 kilowatt-hours of electricity each year. Then, the required solar capacity would be 20 million megawatts, which is about 70 times more than current capacity (about six more doublings would be required). This estimate implicitly assumes the existence of abundant electricity storage, so that sufficient solar electricity can delivered to the user at all times. It also assumes that the losses of electricity associated with storage (charging and discharging a battery, for example) are negligible.

There is a price bubble in Figure 4.5 corresponding to the years 2001-2010, when the average sales price was above the long-term learning curve. The high module price was due to a sharply rising price for pure silicon, the raw-material precursor to the silicon wafers that are used by both the solar cell and microelectronics industries. Silicon's price rise was the result of a surging solar market (a growth that began in the early to mid-1990s) that caused demand to increase well beyond supply. The price of pure-silicon feedstock increased from less than $30 per kilogram in the early 2000's to more than $400 per kilogram in 2008. This price increase created incentives for new suppliers to enter the market and for current providers to increase capacity, which led to an oversupply. By 2010, when much of the added silicon production capacity had come online, the price of silicon feedstock had dropped to about $50 per kilogram, and today it is relatively stable at about $20 per kilogram.

During the decade-long price bubble, non-silicon competitors, most prominently cadmium telluride, had a chance to prosper. In the absence of another price bubble, it is not clear how a new technology would be able to enter the market.

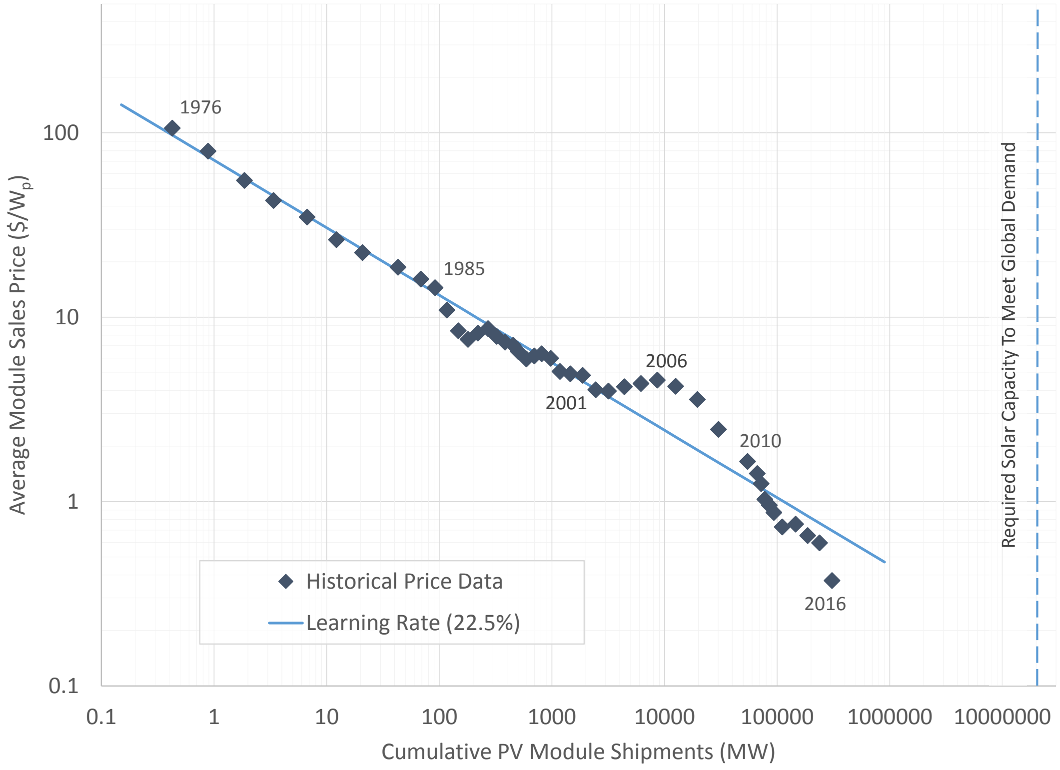

Figure 4.5: Learning-curve analysis of the solar panel with crystalline silicon cells. The average panel price (in units of 2016 U.S. dollars per peak-watt, $/Wp) is plotted against the global cumulative capacity of panel shipments (in units of thousands of kilowatts, or megawatts, MW), 1976-2016. The straight line represents a learning rate (LR in the figure) where the price falls by 22.5 percent for each doubling of cumulative production. The data extend across 20 doublings of cumulative production. The vertical line identifies the required global solar capacity (approximately 20 million megawatts) to satisfy the entire global demand for electricity, assuming perfect storage. Source: International Technology Roadmap for Photovoltaic, Eighth Edition, March 2017, Figure 3: http://www.itrpv.net/Reports/Downloads/.

Today, monocrystalline and polycrystalline silicon modules are being sold at shrinking profit margins, and some high-cost producers and small-volume producers are unable to operate profitably. The result is a widespread consolidation in the PV industry. It is expected that in the next decade only a limited number of major suppliers will remain – a consolidation that is not without precedent, as recent history shows similar consolidation phases for both the microelectronics and electronic display industries.

C. The Photovoltaic Frontier

Recognizing that technologies based upon crystalline silicon will continue to rule the PV marketplace for years to come, research is proceeding in two directions. In one direction, research focuses on further improving today's commercial technologies. For example, in 2014, the 25 percent efficiency record for crystalline silicon that had stood for 15 years was broken with improvements to an architecture based on a heterojunction between the crystalline silicon absorber and an amorphous silicon charge-collection layer.

In the other direction, the goal is to find new technologies and materials that can achieve some combination of dramatically higher performance and dramatically lower costs. Because the cost of crystalline silicon devices is only a small fraction of total system costs, a new kind of solar cell might become competitive, for example, if it lowered balance-of-system costs or served significant niche markets. A successful entrant into the PV cell market need not outcompete silicon, at least initially, as long as it can offer features not possible with silicon. For example, if it can be lower in weight, a system built around a thin-film cell rather than a crystalline silicon cell might be able to be integrated in roof shingles or building façades, and it might have mobile applications.

Here, we take note of six emerging options: gallium arsenide cells, copper-zinc-tin-selenium (or sulfur) cells, organic photovoltaic cells, dye-sensitized solar cells, quantum-dot cells, and perovskite cells. Their efficiency trajectories are plotted in Figure 4.1 (quite a few pages back), alongside today's cells.

Gallium Arsenide Cells

The gallium arsenide (GaAs) solar cell has a specialized use today: powering space satellites. It is the solar cell of choice for this application because of its high efficiency and ability to withstand the radiation in space with limited damage. It possesses a nearly optimal direct band gap, and in fact the record for the efficiency of a single junction cell (28.8 percent) was achieved with this material (see Figure 4.1, shown in black).

The GaAs cell has limited uses on Earth, because high-throughput, low-cost production has not been achieved. Accordingly, we place this cell on the technological frontier, and we assign this cell (and only one other) a score of -2 for manufacturability. The high-efficiency GaAs cell produced today requires a monocrystalline wafer 100 to 200 microns thick, which must be fabricated by slow deposition to realize the proper p-n junction formation. However, given its direct band gap, the cell could be as thin as a few microns. Efforts are therefore underway to develop fabrication processes in which the costly substrate on which the cell is grown is used multiple times. For each use, the top layer containing this solar cell can be peeled off, and the substrate is cleaned and reused. If successful, this approach could propel GaAs technology forward.

Some of the potential of the GaAs cell comes from the abundance of the elements of which it is composed, the low toxicity of the material (arsenic is toxic, but in the cell it is tightly bound to gallium and not easily mobilized), and its stability. The GaAs cell is also versatile: thin versions of the cell are lightweight and efficient, and the cell's performance degrades relatively little over time when subjected to the radiation beyond the atmosphere, which adds to its attractiveness for space applications.

Copper-zinc-tin-sulfur and Copper-zinc-tin-selenium Cells

To overcome the scarcity of tellurium and indium, a chalcogenide cell that might replace the CdTe cell or the CIGS cell is being investigated which uses copper, zinc, and tin instead. This cell is the CZTS cell (C for copper, Z for zinc, T for tin, and S – as with the CIGS cell above – for either selenium or sulfur or both). As with the CdTe and CIGS cells, cadmium sulfide (CdS) is the n-type semiconductor in the highest efficiency version of CZTS cell. The efficiency of the CZTS cell has reached 12.6 percent (see Figure 4.1, where, like the other chalcogenide semiconductors, the CZTS cell is shown in green). The cell can be processed with high throughput via solution-based coating techniques, but precise manufacturing is difficult. The major challenges for further efficiency gains lie in better controlling the ratios of the constituent elements while the polycrystalline material grows and in constructing a heterojunction that increases the photovoltage. The similarity of cadmium use in CIGS and CZTS cells leads us to assign the CZTS cell the same -1 score for compatibility. The stability of the CZTS cell is a current concern and may ultimately remove this cell from contention as a commercial product. Even if this happens, the CZTS cell's role in solar cell research will have been important, because it is inspiring extensive efforts to explore chemical element substitution, with the objective of replacing specific scarce or toxic elements with more abundant or less hazardous ones.

Organic Photovoltaic Cells and Dye-sensitized Solar Cells

Scientists are also exploring solar cells based on thin-film organic (carbon-containing) molecular absorbers. Particularly effective at absorbing solar photons are "conjugated" organic molecules (molecules that contain alternating carbon-carbon single and double bonds). Two promising versions of this technology are organic photovoltaic cells (OPV cells) and dye-sensitized solar cells (DSSCs), both of which have achieved efficiencies of approximately 12 percent (see Figure 4.1, data points in purple). In both cases, the conversion of sunlight to electricity is a two-step process. When light falls on a molecular absorber, electrons and holes are not produced directly; rather, the molecular absorber's internal energy is increased (the absorber is in an "excited" state). Much of the absorbed energy can then be transferred to a second, adjacent electronic material in which electrons and holes are produced. In combination, the two materials make the solar cell.

In OPV cells, both the organic absorber and the adjacent electronic material are organic, whereas in DSSCs the adjacent material is a metal oxide, often titanium dioxide, usually chemically bonded with the organic absorber. In OPV cells, the molecular absorbers are either polymers (containing many repeated chemical units, called monomers) or specific small organic molecules. In DSSCs, the molecular absorbers are more complicated organic molecules, incorporating a metal atom. The "abundance" score for OPV cells in Table 4.1 is +1 rather than +2, because the OPV cell uses a small amount of indium in one of its electrodes. As for the +1 score for the DSSC, it reflects the fact that currently the best DSSCs have the very rare element, ruthenium, at their core, but only a small amount of ruthenium is present (there is only one ruthenium atom per molecule). Moreover, it is likely that similar DSSCs without ruthenium will be developed, either fully organic dyes or ones that use fewer rare metal atoms.

In both OPV cells and DSSCs the band gap can be sensitively tuned by small changes in the chemical structure of the absorber, enabling devices built from these cells to absorb and emit only a small fraction of the incoming solar spectrum, thereby producing a specific color; the nearly infinite range of colors of flowers is evidence of the variety of selective organic molecular absorbers found in nature. A combination of such devices has created the colorful curtain wall at the SwissTech Convention Center (Figure 3.15). Other cells absorb incoming solar radiation only at ultraviolet wavelengths, or only at infrared wavelengths, or at both ultraviolet and infrared wavelengths but still not where light is visible to the eye. Such absorbers open up the possibility of use in window coatings that at the same time are totally transparent and a source of electrical energy; the harvested electricity could be used, for example, to change some property of the window glass, like its ability to transmit heat in summer vs. winter.

OPV cells can be made lightweight and flexible relatively easily. The combination of flexibility and color variability lead us to assign the highest score (+2) for versatility uniquely to the OPV cell in Table 4.1. DSSCs score only +1 on versatility, because the most efficient thin films of DSSCs must be kept rigid in order to encapsulate its liquid electrolyte. As for manufacturability, DSSCs get the lowest score, -2, for two reasons related to the requirement of rigidity. First, the preparation of the titanium dioxide layer requires a high-temperature process, which makes the DSSC incompatible with any lightweight and flexible substrate. Second, the encapsulation of the liquid electrolyte requires a pair of glass substrates.

Key issues that remain for both the OPV cell and the DSSC are long-term stability and fabrication at low cost with high throughput. For improved stability, the cells must be sealed to prevent air from contacting the organic absorbers, which are sensitive to photo-oxidation. Low-cost encapsulation will be critical for commercialization.

Quantum-dot Cells and Perovskite Cells

Two technologies are new entrants to frontier research on solar cells: quantum dot (QD) and perovskite cells. Both involve new materials, and as seen in Figure 4.1, neither has a data point before 2010. Quantum dots are nanometer-scale inorganic crystals, fabricated by deposition of inks, often using lead sulfide. QD cells have efficiencies of just over 10 percent (see Figure 4.1, where its data points are in green because a sulfide is a chalcogenide, and the other three solar cells shown in green also involve chalcogenides). We assign a score of -1 to the QD cell for public health compatibility because of lead's toxicity; current research seeks a substitute for lead that does not compromise performance. The QD cell's versatility results from the ability of the band gap to be tuned by varying the physical dimensions of the dot. We assign a score of +1 rather than +2 for abundance to both the QD and perovskite cells, because, like the OPV cell, indium is used in one electrode.

Perovskites are crystal structures, the most studied of which is methylammonium lead iodide ($CH_3NH_3PbI_3$), which contains both an organic molecule and an inorganic metal halide and which can be processed at low temperatures from solutions to form highly crystalline layers. This perovskite has a direct band gap near the theoretical optimum for single-junction cells and can be produced with relatively few defects, even though deposition is at a relatively low temperature. Benefiting from the knowledge of thin-film solar cells gained over the last decades, the efficiency of perovskite solar cells has had a meteoric rise, from 3 percent in 2013 to 22 percent in 2016 (see Figure 4.1, where the perovskite data points are the only ones in orange, reflecting the distinctiveness of these materials).

The properties of perovskites addressed in Table 4.1 are still emerging, leading us to assign question marks to perovskites under "stability," "manufacturability," and "versatility." We assign a -2 score for "compatibility," because of the lead in current perovskites. Lead accounts for a large fraction of the perovskite's weight, and the most efficient perovskites today, unlike the quantum dots, are soluble in water. Thus, the leaching of lead from the cell into the local environment is a real possibility, raising public health concerns and complicating the management of the cell over its lifecycle, from manufacture through disposal. Research is under way to render perovskites insoluble in water and to find alternative perovskites that do not use lead.

Multijunction cells

The eleven solar cells we have discussed here should not be thought of as alternatives because they can be used in combinations, creating multijunction cells. Indeed, the cells we have grouped as belonging to the "photovoltaic frontier" may turn out to be particularly useful when layered with "today's technologies" cells. For example, much work is underway to create a "tandem" solar cell (a multijunction cell with just two components) that adds a thin-film perovskite cell to a crystalline solar cell. The result is a multijunction cell whose efficiency exceeds the efficiency of a crystalline silicon solar cell on its own: the perovskite layer and the crystalline silicon layer have different band gaps and thus function in combination to absorb more of the solar spectrum. The result is less wasted photon energy and less generation of heat. If such an enhanced crystalline silicon cell becomes commercially competitive, a future for solar electricity based on solar cells made without crystalline silicon would become even less likely.

Earlier, we reported that multijunction cells have been developed principally to enhance the efficiency of high-cost cells for applications where extra efficiency is worth a high premium. New kinds of tandem solar cells will alter this perspective, if it turns out that the performance (efficiency and durability) of the low-cost cells required for commercial electric power can be enhanced with tandem cells produced from inexpensive materials with inexpensive manufacturing.

D. Energy and Greenhouse Gas Performance Indices, End of Useful Life

In this section we consider three indices that are used to evaluate solar panels: 1) the energy to make a panel versus the energy it produces, 2) the analogous question for greenhouse gases instead of energy, and 3) the cost of avoided greenhouse gas emissions. We conclude with a brief discussion of alternative strategies for managing solar power systems when they are no longer useful.

Energy and Greenhouse Gas Performance Indices

Energy Payback Period and Energy Return on Energy Invested

Consider the energy required to make a renewable energy system. Two common metrics are frequently used: *the energy payback period* and *the energy return on energy invested*. The energy payback period is the amount of time needed for the system to generate the amount of energy expended to make it. The energy return on energy invested, a closely related concept, incorporates an estimate of the expected operating life of the system; specifically, it is the energy produced throughout the life of the facility divided by the energy invested to make it. For example, if it takes 1,000 energy units to produce a solar-panel system, and the system produces 200 energy units each year for 20 years, *the payback period is five years and the energy return on energy invested is four.* Here, the energy to make something includes the energy to mine the elements and to transport components at various stages of production from place to place; in all, this is called the "embodied energy" of the device.

A recent meta-analysis of research published on the energy payback and energy returns for solar panels showed that the energy payback period for monocrystalline and polycrystalline panels is about four years.[15] With the further assumption of a panel lifetime of 20 years, *the energy return on energy invested would be about five*. The energy-intensive step of creating crystalline silicon wafers from pure silicon feedstock dominates the front-end energy investment. By contrast, some new thin-film cells require far less energy to build and have *payback periods as short as one year.*[16]

If the energy return on energy invested is less than one (equivalently, if the energy payback period is longer than the lifetime of the system), the system never breaks even. Although panels are designed to last 20-30 years and to lose less than one percent of their conversion efficiency each year over that time period, the performance of a solar panel may fall short. For example, solar panels located in deserts can be degraded by windstorms carrying sand. Panel surfaces can also be fouled by dust and bird droppings. To limit the resulting damage, panels are routinely wired together in ways that assure the shading or fouling of one cell will not degrade the performance of the entire panel.

---

[15] New estimates suggest even shorter payback periods are now being achieved. http://www.sciencedirect.com/science/article/pii/S1364032116306906.

[16] http://www.sciencedirect.com/science/article/pii/S136403211500146X.

Greenhouse-gas Payback Period

Similar calculations can be done for the *greenhouse-gas payback period.* The greenhouse gas emissions required to make and install a solar collector system can be compared to the reductions in greenhouse gas emissions each year achieved by making electricity from the solar collector system instead of some other electricity source. These calculations are highly site-specific for two reasons. (For simplicity, we consider only carbon dioxide – the most important greenhouse gas, but not the only one.) First, the calculation depends on the greenhouse gas emissions associated with the energy expended to produce the device: the electrical heating of the silicon ingot (prior to the cutting of the wafers) may have used coal-based electricity or windpower, for example. Second, there can be big differences in the amount of greenhouse gas not produced because the solar energy source produced that energy instead: the solar source could have resulted in a reduction of electricity production from hydropower (with almost no greenhouse gas emissions) or from natural gas (with considerable greenhouse gas emissions).

Thus, greenhouse gas emissions appear in two ways in a calculation of the greenhouse gas payback period: emissions associated with 1) the energy required to produce the device, and 2) the energy displaced from the grid each year when the device is producing electricity. If the two emissions are equal (for the same amount of energy produced), the energy payback period and the greenhouse-gas payback period are identical; for our example, *the greenhouse-gas payback period would be four years.* But if a device is produced in China, where the grid is dominated by coal, and then used in California where the grid is much less carbon-intensive, the greenhouse payback period might be twice as long as the energy payback period.[17]

A full estimate of greenhouse-gas intensity must include not only emissions related to energy flows but also emissions relative to chemical treatment. For example, nitrogen trifluoride (NF_3), a particularly potent greenhouse gas, is used currently to etch openings into the coatings of the silicon wafer to enable electrical contact.

Cost of Avoided Emissions of Greenhouse Gasses

A related index, for solar power and other technologies that reduce greenhouse gas emissions, is the cost of avoided emissions per ton of reduced greenhouse gasses. To illustrate this calculation, consider a large utility facility whose construction cost is $1 per peak watt. Also assume that the capacity factor is 20 percent (about 1,750 hours per year), so that 1.75 kilowatt-hours of electricity are produced each year for each dollar spent. One must also decide how many years the facility will operate; If it operates for 20 years (clearly, a critical assumption), it will produce 35 kilowatt-hours of electricity. The carbon intensity of the power that it displaces must be specified: assume 500 grams of carbon dioxide emissions for each kilowatt-hour produced (an average value, and about twice as much where electricity is produced from coal), so 17.5 kilograms of carbon dioxide emissions will be avoided. In the usual units for this topic, the cost of emissions reduction is then *about $60 per ton of carbon dioxide.* (Among the issues ignored here are costs associated with operation and maintenance during the life of the plant, as well as emissions associated with construction, as discussed above.) These costs are somewhat higher than the "social cost of carbon" introduced by the Obama Administration and currently being set aside by the Trump Administration.

End of Useful Life

The infrastructure for the management of solar cells at the end of their useful life scarcely exists, but it will become important. Reduce-reuse-recycle is a well-known hierarchy. Reducing the amount of material requiring handling comes along with improvements in a solar cell's conversion efficiency, since much of the bulk of a solar device is in its balance of system, and less balance of system is required for the same power output when the system is more efficient.

Reusing the balance of system after treatment is easy to imagine, but even the cell may be designed for reuse.

As for recycling, whether the valuable materials in a solar cell will be recycled is not clear. They often represent only a small fraction of the total weight of the device. Yet platinum is so valuable that it is often retrieved from the catalytic converters of junked cars. The recycling of cadmium and tellurium retrieved from CdTe cells, which would be analogous, got off to a halting start when the largest manufacturer of CdTe panels first announced that it was embedding the cost of recycling in the panel cost, but then decided to defer the recycling cost.

Finally, disposal. The two principal managed destinations today are the landfill and the incinerator, each of which can be state-of-the-art, with hazardous materials well contained.

---

[17] A careful definition of the greenhouse gas emissions associated with electricity produced in some particular political region requires taking into account not only production within the region, but also imports of electricity into the region and exports from it. Strictly speaking, for an intermittent resource like solar power, one should also identify the hours of the year when the system displaces other power sources and consider greenhouse gas emissions only for those hours. Moreover, one should identify the marginal, rather than the average, power source that will be added to the grid in order to manufacture the panels and the specific source that will be taken off the grid when the panels produce power.

Article 5: Grid Integration and Policy

A full accounting for any solar power project must consider not only the panel and the balance of the system at the project level (discussed in Article 3), but also the project's impact on the grid. Article 5 focuses on solar power's intermittency and only partial predictability, which are creating problems for grid management that threaten to restrict future growth of solar power. Article 5 also discusses the variety of technological and policy responses that the intermittency problem is eliciting, including the promotion of natural gas and electricity storage, the enhancement of electricity transmission in order to access a diversity of sources, and the preferential use of electricity at times of the day when electricity is available in excess.

A second focus of Article 5 is the policies that have enabled the rapid growth of solar energy, with a focus on the U.S. and, within the U.S., the State of New Jersey. New Jersey, relative to many other states, has been particularly determined to create incentives for solar power projects that provide electricity directly to users, not only to the grid. Worldwide, incentives are diminishing, and a major open question is the extent to which the growth of solar power capacity will be adversely affected. A question within that question concerns the reduced incentives specifically for small-scale and dispersed electricity production.

Article 5 concludes with a discussion of "grid parity," an awkward metric widely used by the solar industry to measure its progress against conventional energy sources. The problem with "grid parity" is that it ignores the costs of grid integration.

A. Grid Integration and Supply Variability

Electricity supply is managed today in large systems, called grids. The largest grids coordinate the provision of power to millions of customers. Grid operators working at a central location inform the operators of various power plants that output from their plant will be required, with various notice periods from less than a second to days. In this way, the grid's variable demand is accommodated. Demand variability arises from predictable behaviors (for instance, most people sleep at night, or electricity consumption rises as viewers separate themselves from their televisions at half-time during the Super Bowl) and unpredictable ones (a large motor in a factory shuts down). Variability and unpredictability are no strangers to the grid.

With solar power's arrival, however, an electric grid now needs to respond not only to predictable and unpredictable demand but also to predictable and unpredictable supply – with minimal help (at least today) from electricity storage. Solar power's intermittency over days and seasons is largely predictable. However, solar power can be unpredictably intermittent on the scale of minutes (as clouds block the Sun) and days (from bad weather). Balancing supply and demand in the presence of unpredictable intermittency is a challenge to grid management that grows in importance as solar power gains market share.

The general challenge here is "dispatchability." A dispatchable source of electricity provides power when power is required. Solar power on its own is not dispatchable. To be embedded in a dispatchable system, it must be augmented by some combination of other power sources, electricity storage, and demand management.

The Duck Curve

Figure 5.1 illustrates these issues. It shows a recently popularized curve, the Duck Curve, which highlights the complications for grid management that accompany a rising fraction of solar power on contemporary grids. The curve was developed by the California Independent System Operator (CAISO), the organization responsible for the performance of the electricity grid that provides electricity to nearly the entire state of California. Figure 5.1 shows two curves of real data: actual hourly "net load" for Saturday, March 31, 2012 (light blue) and Sunday, March 31, 2013 (dark blue). "Net load" is CAISO's total production of electricity minus its production of electricity from solar and wind energy at "utility facilities" that directly supply its grid.

A comparison of the two blue curves shows that the net load during the day shrank between 2012 and 2013. This is because the combined solar and wind contribution to total supply grew faster than the total load. As a result, there was less production of electricity during the day, in aggregate, from all of the other in-state sources (fossil fuels, nuclear power, hydropower, bioenergy, and geothermal energy) and the out-of-state plants whose electricity CAISO imported. In the evening, when the solar load was absent, the net load was substantially larger in 2013 than in 2012.

The 2012 and 2013 curves look nothing like a duck. But Figure 5.1 also shows modeled data for the same March day for several future years (at the time of the preparation of the figure). As was the case between 2012 and 2013, the production of solar and wind energy during the day increases year after year and results in an ever smaller mid day net load. Also, the future net load in the evening increases. By 2020, the net-load curve outlines the underside of a duck – with a belly that is closest to the ground not long after noon, a long neck stretching upward in the evening hours, and even a tail during the first hours after midnight.

This visual metaphor, it seems, has injected exactly the amount of levity to enable candid discussion of the challenges that are beginning to arise as intermittent resources achieve deeper penetration on the grid. There are two separate concerns in the figure, the first at midday and the second in the evening. At midday, the combined output of solar and wind energy could drive down the need for other power sources to such an extent that there is no longer any need for some current sources that would normally run continuously (baseload power plants). Reducing the power output of a baseload plant and then raising it again, if it can be done at all, is likely to degrade the plant's long-term performance. The grid operator wishing to sustain constant output from the baseload plants has an alternative, which is to require the grid's solar and wind facilities to "curtail," or "spill" some of the power they produce at midday. These renewable power sources will then sell less electricity to the grid and lose revenue. Either way, at some high level of penetration of intermittent renewable energy, system costs become formidable.

The second challenge occurs in the early evening and may become even more daunting and costly. From 4

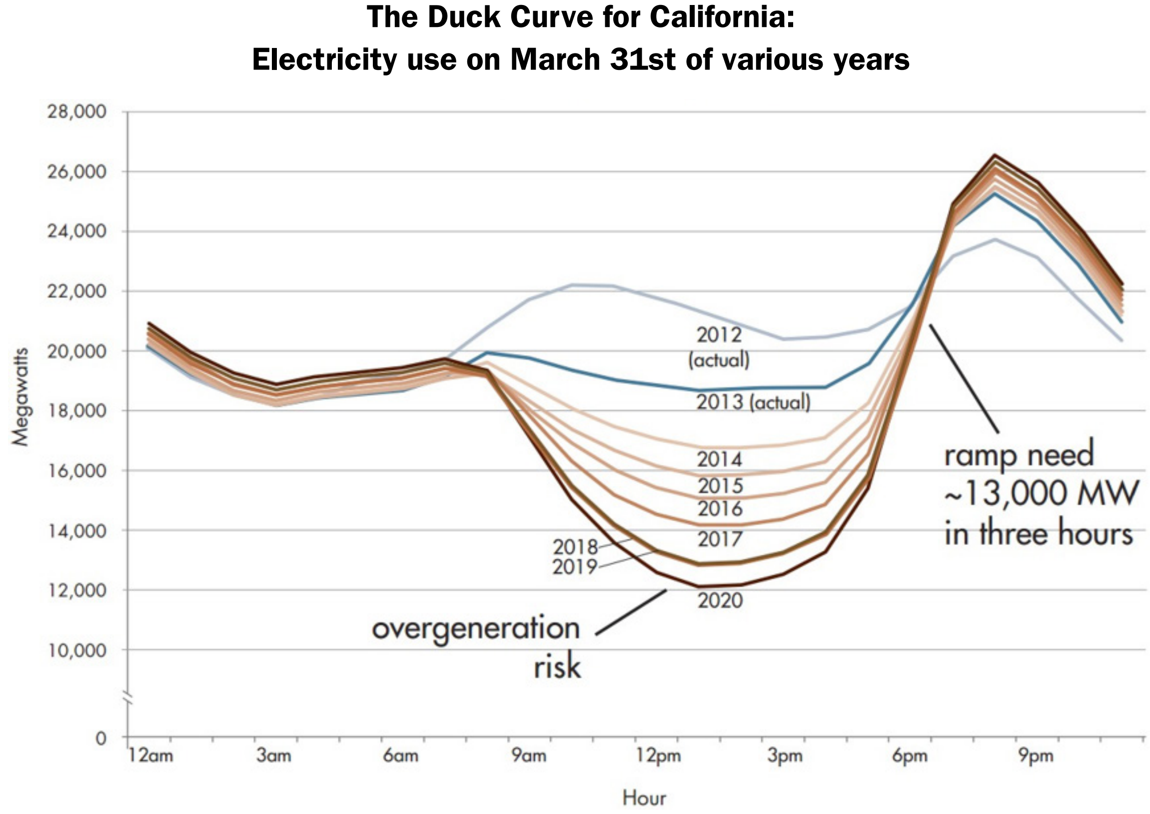

Figure 5.1: The original Duck Curve. The hourly net load (total electricity consumption minus electricity produced from utility wind and solar sources) on the CAISO grid for March 31 of successive years. Actual data for 2012 and 2013, modeled data for later years. Source: CAISO, the California Independent System Operator: https://www.caiso.com/Documents/FlexibleResourcesHelpRenewables_FastFacts.pdf

p.m. to 7 p.m., as the Sun descends, solar power's contribution to the grid falls, while the demand for electricity rises (people are returning home and running appliances while stores remain open). The extra demand for power at 7 p.m. relative to 4 p.m. (most of the length of the duck's neck), as noted in Figure 5.1, is projected to reach 13,000 megawatts in 2020, as total net power approximately doubles. In the CAISO system, gas turbines have been playing the primary role, ramping up their power as needed; in the CAISO models this role continues to dominate through 2020.[18]

Figure 5.1 accounts only for solar projects where all of the solar electricity is sold directly to utilities. It does not include solar power from customer-owned projects (also called "behind-the-meter" projects and "non-utility" projects) where some of the solar electricity is not sold directly to the utilities, a category that includes solar electricity produced on residential and commercial rooftops. Figure 5.2 is an instructive augmentation of the CAISO Duck Curve that repairs this omission by including an estimate of "non-utility" solar electricity. The data shown are for August 7, 2016, when at midday about 500,000 distributed solar energy sources in California were contributing an estimated 4,000 megawatts of non-utility solar power – at the same time as utility solar projects were contributing about 8,000 megawatts.

The gray region of Figure 5.2 represents electricity provided to customers in California from all sources except wind and solar sources. The production shown comes from fossil energy sources (natural gas and coal) as well as from nuclear fuels and several renewable energy sources other than wind and solar energy ("small" hydropower, geothermal power, and electricity from biomass). Figure 5.2 also shows, as three separate regions, three other contributions to California's electricity production that day: electricity from utility wind turbines (blue), utility solar facilities (orange), and non-utility solar facilities (yellow). Note that wind power was strongest at night and weakest in

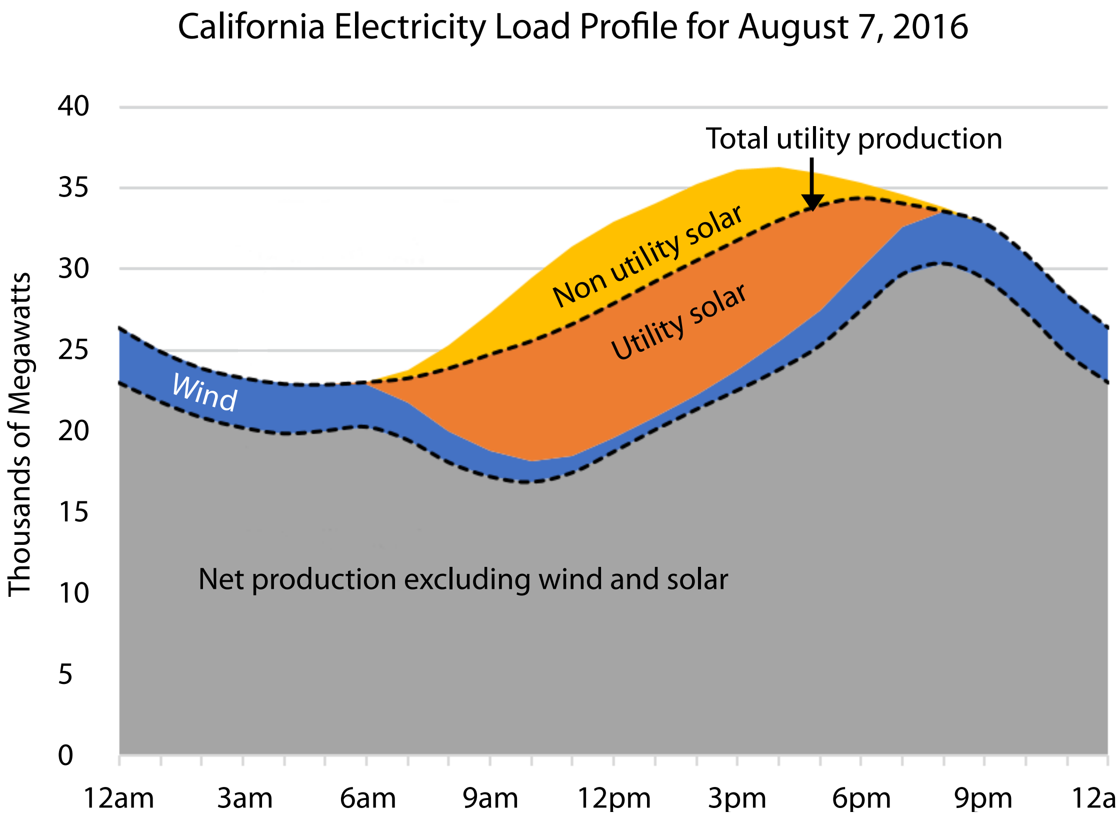

Figure 5.2: A Duck Curve for the same CAISO grid as Figure 5.1, but for August 7, 2016, and with the addition of solar electricity production from distributed sources. From bottom to top, the lowest (gray) region represents production from all utility sources aside from wind and solar. The next (blue) region is wind generation, and the region above it (orange) is utility solar generation. The top-most (yellow) region is an estimate for customer-owned ("non-utility") solar power that is produced "behind the meter." Unlike Figure 5.1, the vertical scale is a continuous linear scale from zero, without a gap. Adapted from: Paulos, Bentham, 2016 "California has more solar power than you think – a lot more." Greentech Media, http://www.greentechmedia.com/articles/read/california-has-more-solar-than-you-think.

---

[18] For a sketch of a duck superimposed on the data, see http://insideenergy.org/2016/10/25/learning-how-to-adapt-to-more-renewables-as-duck-curve-deepens/. For an updated discussion from CAISO, November 3, 2016, see https://www.greentechmedia.com/articles/read/the-california-duck-curve-is-real-and-bigger-than-expected.

the middle of the day, when the Sun was strongest; this beneficial anti-correlation is observed in many locales at many times, but of course not everywhere nor all the time. As for "utility solar" power, about one tenth of the electricity in this category was electricity from solar thermal power plants rather than solar photovoltaic power plants. "Non-utility solar" can only be estimated roughly, because much of this production is used at the site of the producer without ever being sent to the grid. During that particular day, distributed solar electricity production (yellow region) was about half as large as centralized solar electricity production (orange region). Wind electricity was roughly half as large as centralized plus distributed solar electricity.

Many states have a renewable electricity target that is a percent of total electricity. Various choices for this fraction can be formulated. The data behind Figure 5.2 reveal that between them, wind and utility electricity accounted for 20 percent of the day's electricity load recorded by utilities; including non-utility solar, total electricity production from solar of both categories and wind accounted for 25 percent of total electricity production from all sources. Including, as well, the other electricity production that day that in California counts as "renewable" (from small hydropower, geothermal, and biomass sources), utility production of renewable energy from all sources was 27 percent of total utility production, and total renewable energy production including distributed solar electricity was 31 percent of total electricity production from utility and non-utility sources. Other renewables percentages could take into account renewable energy embedded in imported electricity.

The lower black dashed line in Figure 5.2 corresponds to the Duck Curve lines in Figure 5.1. Looking ahead, California can expect growth in both utility and non-utility solar power. Both will affect the grid in the same way, further suppressing daytime net load and further steepening the evening ramp. It is clear from Figure 5.2 (whose vertical scale, unlike Figure 5.1, has no "suppressed zero") that another doubling of solar capacity, keeping the total load fixed, would cut deeply into the gray baseload region, significantly increasing solar power's disruption of the grid as a whole.

Fattening and Flattening the Duck

A recent report from the U.S. National Renewable Energy Laboratory distinguishes two approaches: reduce the cost of a fat duck, and flatten the duck so that it is less fat.[19] The cost is reduced if the grid can be made more flexible, notably, by reducing the importance of sources of electricity that are hard to scale back; "must-dispatch" nuclear power plants and coal power plants are relatively inflexible, while hydropower and gas-turbine power are relatively flexible.

The duck can be flattened both by eliminating some of the sources of peak load and by shifting the load away from the peak. Some portion of peak load can be eliminated in buildings. One way is to improve the efficiencies of the electric appliances that contribute significantly to electricity demand (air conditioners, refrigerators, water heaters, lights, and electronic equipment). Another way is to build buildings with better insulated roof, walls, windows, and with façades that allow sunlight to enter the interior in winter but not in summer.

As for shifting the load, there are many alternatives that involve energy storage. Power can be used at midday to pump water uphill from a lower to a higher reservoir, and the flow can be reversed in the evening, retrieving nearly as much power as was used earlier – this strategy is called "pumped storage hydroelectric," or either "pumped hydro" or "pumped storage," for short. Storage in buildings can be in heated water, thermally charged at mid-day and discharged several hours later. Air conditioners, water heaters, and refrigerators can be made to run mostly at hours of peak electricity supply, and the batteries in electric vehicles can be charged preferentially then too. Distributed electricity storage (batteries in homes and larger buildings) and smart communication can help too.

Still another strategy is to orient solar panels southwest instead of south. This shifts the peak output of solar panels from noon to the afternoon, toward the peak in demand. There is currently a subsidy in California for new solar homes that have their panels on a roof oriented within 11 degrees of due west.

Load shifting can be incentivized by pricing. A common price incentive is the "time-of-use" rate, where electricity is valued at a higher price when demand is highest, such as on a summer afternoon. California is already offering time-of-use rates for customers installing distributed solar power, and the time-of-use rate will be the default rate for all customers in 2019.

A further strategy to reduce the stress on the grid from solar power is to extend the grid geographically to integrate loads and supplies that have complementary time profiles. For example, planning for greater integration of power sources in the western U.S. is underway, driven in this case especially by the desire

---

[19] Denholm, Paul, Matthew O'Connell, Gregory Brinkman, and Jennie Jorgenson, 2015. "Overgeneration from Solar Energy in California: A Field Guide to the Duck Chart," National Renewable Energy Laboratory Technical Report, NREL/TP-6A20-65023, November 2015.

to allow greater access to windpower produced beyond CAISO's territory.

A duck curve emerged from CAISO data in Figure 5.1 only because of several choices. First, on the vertical axis, the value "10,000 megawatts" was located close to the horizontal axis, truncating the vertical space; otherwise, the duck would not stand on the ground. Second, a March weekend date was chosen, when total demand is particularly low (there is minimal heating and cooling). In summer, demand is much higher. Third, in California, windpower has a smooth profile (Figure 5.2); in Texas and the Mid-West today, wind variability looks entirely different. Nonetheless the Duck Curve quickly became iconic. Alluding to the curve is a convenient shortcut for identifying the challenges of integrating intermittent renewable energy into a grid – and the solutions.

B. Enabling Policy

Government incentives have allowed solar electricity to grow rapidly, and costs to fall. A generic justification for incentives to new industries is that they accelerate the arrival of a desirable commercial product, especially when the alternative of raising the taxes on its competitors is politically infeasible. Solar power's desirability arises from its much lower emissions of traditional air pollutants and greenhouse gases, relative to fossil-fuel power sources. Traditional fossil-fuel-derived air pollutants (including soot, sulfur and nitrogen oxides, and mercury) adversely affect public health and agriculture. Greenhouse gases cause climate change. Pollution-control investments at fossil fuel power plants are gradually lowering these emissions, but usually not to the levels that solar power achieves, even when the full life cycle is taken into account.

It is useful to distinguish two kinds of solar incentives. One class of incentives lowers the cost of a solar project, independent of how large it is; in the U.S. these are largely incentives provided by the federal government through tax deductions. The other class specifically encourages distributed generation of solar power, and in the U.S. these are largely state-level policies.

We discuss two examples of the first and two examples of the second. The federal incentives are the Investment Tax Credit and accelerated depreciation. The state incentives are established through the Renewable Portfolio Standard (and the solar carve-out from that standard) and "net metering." We also mention the feed-in tariff, a stimulant of distributed generation widely used in Europe.

These five policies are not the only significant ways by which governments foster solar energy. Others include the funding of research and development and targeted aid to manufacturing companies. A carbon tax or a cap-and-trade regime for carbon dioxide emissions also improves the competitiveness of solar power, relative to many alternative sources.

Investment Tax Credit

The principal subsidy from the federal government that affects the cost of a U.S. solar project is the Investment Tax Credit. It applies to solar power projects at all scales. The recipient of an Investment Tax Credit may subtract the value of the credit from the tax that he or she (or it, in the case of a corporation) owes. Currently, the Investment Tax Credit equals 30 percent of the capital cost of a solar power project. In effect, the Investment Tax Credit allows the government to share in the cost of construction.

The size of an investment tax credit does not depend on how much power the system owner produces, only on the amount spent to bring the unit online. As a result, this kind of credit rewards investment-intensive projects. In the case of solar power, a residential project is usually more investment-intensive than a mid-scale and utility project, measured in dollars invested per kilowatt of capacity. As a result, the Investment Tax Credit may treat residential projects preferentially.[20]

In December 2015, the U.S. government renewed the Investment Tax Credit with a schedule of stepwise reductions. Projects where construction begins on or before 2019 are eligible for a 30 percent credit, those beginning in 2020 are eligible for a 26 percent credit, and those beginning in 2021 through 2023 are eligible for a 22 percent credit. After 2023, the tax credit is permanently zero for residential projects and 10 percent for mid-scale and utility projects.

Accelerated Depreciation

A fixed asset like a solar panel loses value over time, mostly due to wear and tear. The tax code in the U.S. allows businesses to recover this depreciation in the value of a fixed asset as a tax deduction spread over a specific number of years. An individual is not allowed to take the depreciation deduction for items of personal use, but a company that leases an individual's roof and puts a solar collector there can claim the deduction. According to the tax code, the solar panel is considered a "five-year asset" subject to "accelerated depreciation," and the initial cost basis for depreciation deductions is 85 percent of the original cost. The "five-year" classification is supportive of solar power. If the useful life of a solar panel for depreciation purposes were more reflective of its expected useful life – 20 to 30 years – the recovery of its initial cost basis would occur much more slowly. Also, the specific rules for accelerated depreciation allow more than half of the total depreciation to be deductible in the first and second year. The depreciation deduction is typically

---

[20] The Future of Solar Energy: An Interdisciplinary MIT Study, http://energy.mit.edu/research/future-solar-energy.

roughly as large as the Investment Tax Credit, where it can be claimed.[21]

The Renewable Portfolio Standard

Twenty-nine states have adopted a Renewable Portfolio Standard (RPS), which typically requires each of the state's retail suppliers of electricity either 1) to produce some minimum fraction of its electricity from prescribed renewable energy sources, or 2) to buy from others what it cannot produce itself, or 3) to pay an Alternative Compliance Payment. The list of allowed renewable sources varies from state to state but typically includes solar power, wind power, landfill gas, and small hydropower facilities.

The minimum-fraction requirement creates a market where, either under bilateral agreements or at an auction, each retail supplier meets a portion of its requirement by buying renewable electricity from other market participants, including (via brokers and aggregators) producers of distributed solar electricity. The currency in the RPS market is the Renewable Energy Certificate (REC), which is nominally equivalent to the environmental attributes of one megawatt-hour of renewable electricity. The Alternative Compliance Payment puts a cap on the REC price, because, when the supply of RECs is small relative to the required purchases, the retail electricity supplier will pay the Alternative Compliance Payment rather than pay for RECs at a higher price.

The Solar Carve-Out

Six U.S. states and the District of Columbia go beyond the RPS to incentivize solar power more directly. They have created a solar "carve-out," which requires each retail provider in the state to produce a minimum fraction of its total electricity from solar power sources. A separate market for solar power emerges, whose currency is the Solar Renewable Energy Certificate (SREC), equivalent to one megawatt-hour of solar electricity, and whose market price cap is the Solar Alternative Compliance Payment. In New Jersey (one of the six states) the authorized producers of SRECs must be connected to the distribution system serving New Jersey, whereas the authorized producers of RECs face weaker restrictions: they are required only to be connected to the much larger north-east U.S. grid ("PJM"), of which New Jersey is a part. The SREC market, therefore, directly stimulates New Jersey's in-state solar electricity production.[22]

In states with a solar carve-out, the markets for RECs and SRECs are distinct. In New Jersey, for example, the SREC market has been dwarfing the REC market, measured by the value of the certificates bought by the retail producers to meet their requirements. In 2016 the total value of the SREC market in New Jersey was 460 million dollars, and the total value of the REC market (excluding the SREC market) was 120 million dollars. The SREC price (preliminary data)[23] was 15 times higher than the REC price ($225 versus $15 per megawatt-hour).[24]

Future RECs and SREC prices are unpredictable, even when required percentages and compliance payments are announced far in advance.[25] From the perspective of a potential investor in a distributed solar energy project, the future SRECs price is one of the major uncertainties, alongside other uncertainties such as future project costs and government incentives.

[21] The initial cost basis that can be depreciated is the full value of the project, minus half of the Investment Tax Credit, thus 85 percent of the original cost. For a business with an assumed 35 percent marginal tax rate, therefore, the value of the deduction is 29.75 percent (35 percent of 85 percent) of the project value, almost exactly the same as the deduction for the Investment Tax Credit, 30 percent of the project value. As a result of the two deductions, about 60 percent of the cost of the system is recoverable through tax benefits. Governed by the Modified Accelerated Cost Recovery System (MACRS), the five-year assumed useful life leads to a six-year schedule of deductions; as percentages of the initial cost basis, they are 20, 32, 19.2, 11.52, 11.52, and 5.76, for years one through six (adding up to 100 percent of the cost basis).

[22] New Jersey's SREC market, which became operational in 2004, has had a complex interaction with its in-state solar industry. In 2010, New Jersey stimulated its in-state solar electricity industry by establishing a high value for the Solar Alternative Compliance Payment (above $600 per megawatt-hour) when the SREC supply was small, resulting in a spot-market price for SRECs at roughly the price of the compliance payment. The very high SREC price generated an abundance of new solar power projects and a plummeting SREC price. To prop up the price, in 2012 New Jersey more than doubled the effective percentage targets, starting in 2014, to above two percent, and the market stabilized.

[23] www.njcleanenergy.com/files/file/rps/EY%202015%20RPS%20Summary%20Result%20Tables%20Final%20082416.pdf

[24] In 2016, for New Jersey's RECs and SRECs markets, respectively, the required percentages of total electricity production for each retail supplier were 14.9 percent and 2.75 percent, and the compliance payments were $50 and $323 per megawatt-hour.

[25] New Jersey has announced a schedule for the solar carve-out and the Solar Alternative Compliance Payment through 2028. The solar carve-out in 2028 is set at 4.1 percent and the Solar Alternative Compliance Payment at $239 per megawatt-hour. See https://www.pjm-eis.com/program-information/new-jersey.aspx.

A Numerical Example: Subsidies Shorten the Payback Time

A homeowner who is eligible for the federal Investment Tax Credit and the Solar Renewable Energy Certificate finds a solar project on her roof to be much more attractive financially than a homeowner who can access neither of these incentives. In Article 2 we worked out the payback period (the number of years required for a homeowner to recoup an initial investment through a stream of savings) for a solar panel that costs $1,200 and produces 500 kilowatt-hours of electricity per year, with no incentives. We assumed the homeowner would otherwise have purchased that electricity at 15 cents per kilowatt-hour (a representative cost for retail electricity), so that she saved $75 per year. The payback period is then 16 years.

But if the homeowner actually pays only 70 percent of the cost, or $840, thanks to the Investment Tax Credit, the payback period drops to 11 years. And if the homeowner, because she lives in New Jersey, also receives Solar Renewable Energy Certificates for the 500 kilowatt-hours her solar collector produces each year, and the going rate for these certificates is (conservatively) also 15 cents per kilowatt-hour, so she receives a payment of $75 per year, as well. Each year she saves $75 by not buying 500 kilowatt-hours of electricity, and she earns a second $75 for producing that electricity with solar energy, so each year she saves $150. The new payback for the panel, with both incentives in place, is 5.6 years ($840 of one-time capital outlay, divided by $150 per year of benefit). The payback in this example is now three times shorter.

Similar calculations apply to mid-scale projects, like Princeton University's. In New Jersey, early in the SREC program, projects were eligible for SRECs only if their capacity did not exceed two megawatts. Then the cap was eliminated, and the Princeton University 5.4 megawatt project became more attractive financially. As seen in Figure 3.10, many qualifying projects larger than two megawatts have now been built in New Jersey.

Net Metering

"Net metering" is another important policy that many states have implemented to encourage residential and mid-scale solar projects, much as Solar Renewable Energy Certificates do. Net metering, in its simplest form, requires an electric utility to accept all of the electricity sent to the grid by every customer who is an approved owner of solar power systems and to value that electricity at the retail price for electricity. When the utility buys power from other sources, it pays a lower, wholesale price. Because net metering policy assigns the same price to the electricity transmitted from the customer to the electric utility and from the utility to the customer, the customer's bill can be determined by a single meter that runs forwards and backwards – hence the word, "net." Forty states, Washington, D.C., and three US territories have adopted some form of net-metering policy.[26] However, currently, electric distribution companies in several U.S. states are seeking revisions to regulations so that the solar power delivered to them from decentralized sources costs them less. They argue that paying retail prices for this power creates uncompensated costs. Yes, for some peak hours in the summer the customer's solar power may be worth more to the utility than its retail price. But for most of the hours the power a distributed solar generator sells to the utility is less valuable than the other forms of power that the utility introduces onto the grid, because the solar power is intermittent and unpredictable. All of the arguments for and against distributed generation of electricity, discussed in the previous section, come into play.

A Side Rule Prevents the Customer from being a Net Exporter

Some states with net metering have a side rule that treats the solar electricity that a customer sells to a utility differently, once its amount exceeds the customer's own electricity purchases from the utility (typically averaged over a year). If a homeowner produces less power over a year than she consumes, all of the power her panels produce is valued at the price of retail power. But if she produces more than she consumes, the extra power is valued at the price of wholesale power. Thereby, the net-metering incentive is capped. The larger the customer's demand, the larger the available subsidy. In effect, this side rule allocates the pool of net-metering subsidies in a way that favors the large consumer.

Let's continue our numerical example. Suppose that our homeowner uses 6,000 kilowatt-hours of electricity over a year and that she confronts a retail price of 15 cents per kilowatt-hour; without any solar panels, therefore, her electricity would cost her $900 per year. Now suppose she installs an eight-panel collector on her roof. As above, each panel produces 500 kilowatt-hours each year, so her panels eight produce 4,000 kilowatt-hours of electricity and save her $600 each year. She buys the remaining 2,000 kilowatt-hours from the utility each year, at a cost of $300.

Now suppose that she decides to double her project and install another eight panels, for a total of 16, thereby producing 8,000 kilowatt-hours each year from her panels. She now is producing 2,000 kilowatt-hours *more* than she is using, and she has become a net seller to the utility rather than a net buyer. Here's where the wholesale versus retail reimbursement rule applies. In states with this restricted form of net metering, the utility pays the householder not the retail price, but the much lower wholesale price – say, 5 cents per kilowatt-hour. So the first four of her new panels earns her $300 per year (because she saves that amount by not buying retail power from the utility), but the second four of her

---

[26] See map at https://www.seia.org/initiatives/net-metering

new panels earns her only $100 per year in actual reimbursement from the utility. The final four of the 16 panels may not be worth their investment, since the power they produce is worth three times less. The homeowner may settle for 12 panels, or (if she is allowed) she may opt for the full 16 panels, if the cost of these extra panels is small enough.

This asymmetry in the treatment of a net seller and a net buyer is designed to discourage distributed solar producers from becoming solar power exporters – for example, to prohibit a farmer with limited need for electricity from installing panels on several parcels of land and connecting the panels to the grid. However, what often happens with such mid-scale projects is that a third party who already buys a large amount of electricity from the grid rents the land from the farmer, buys the panels, and offsets its own purchases from the utility with the power it is credited with producing on the farmer's land.

The Feed-in Tariff

The feed-in tariff has been the backbone of the expansion of Europe's residential solar power, led by Germany. The tariff is a constant price per unit of solar electricity that a government guarantees a homeowner for a specific number of years for all of the solar electricity that the homeowner produces. The feed-in tariff provides greater certainty about the price that will be paid for a project's future electricity, relative to the Solar Renewable Energy Certificate, because that price is determined by the government in advance, not by the market of the day.

Government reimbursement per kilowatt-hour of solar production was generally much higher when feed-in tariff programs were launched than later. In the United Kingdom, for example, the feed-in tariff was first available in March 2010, and for the first two years the very high price of 43.3 pence (about 70 U.S. cents at the time) per kilowatt-hour was guaranteed for 25 years. Prices for installations authorized in the fourth quarter of 2016 are much lower. The nominal price is about 10 times less (4.18 pence, or about six U.S. cents, per kilowatt) for small projects (those with capacities below 10 kilowatts); even after adjusting the nominal price upward to take into account a modest credit for unmeasured but assumed "exports" to the grid, the effective tariff is still dramatically lower than it first was.

Third-Party Ownership

The deployment of distributed solar energy has been accelerated by the wide use of third parties, who are able to access financial incentives that are unavailable to the host individual or host institution. The general mechanism is the "power purchase agreement," a financial arrangement where a company owns solar panels located on a property that it does not own. In one version, simplified here, a specialized company, in effect, rents the roof of a home for a fixed number of years. It installs solar panels on the roof and agrees to maintain them. The company receives three subsidies: the Investment Tax Credit, a portion of the depreciation allowance, and the RECs or SRECs. The homeowner pays no money up-front. She pays the company for the electricity produced by her panels, but the company charges her a rate that is less than the rate that she had been paying to the electricity utility, so she saves money. The company makes money too, if its project cost (panels, installation, and maintenance), reduced by the Investment Tax Credit and the depreciation allowance, is less than its revenue from the homeowner and the project's SRECs. The company may lease thousands of roofs, lowering its per-household costs by streamlining the permitting and using its labor force efficiently.[27]

At the mid-scale, the institution that hosts the project may not pay federal taxes – the project may be at a municipal government facility or a school, for example. In these cases, a third party that does pay taxes and thus can benefit from the tax credit often owns the project. The third party leases the facility to the host institution and claims the tax credit. This is the legal arrangement in place for Princeton University's field, where the third party is a financial services company.

Pressures to Reduce Incentives

The societal impact of policy incentives for solar power was modest when there were only a few beneficiaries – the early adopters. But as solar power increases its share of electricity production, some utilities are pushing back, arguing for reductions in these incentives (which they call "subsidies"). These utilities emphasize the consequences for the non-adopters, in their twin roles as taxpayers and ratepayers: subsidies that reduce the taxes of the early adopters shift the cost of paying for government services onto other taxpayers, and subsidies that reduce the electric bills of the early adopters shift utility system costs onto the rest of the utility's customers (ratepayers), whose electric bills increase.

The utility system costs most cited include the costs of maintaining reliable infrastructure, assuring back-up, incorporating new grid-related technology as it becomes available, and providing universal access. Utilities maintain that these costs account for much of the difference between the retail and the wholesale price, and therefore that sellers of distributed solar electricity to the grid must be required to accept less than the retail price. Toward that end, regulators in some U.S. states are considering a "connect charge" that every residential and mid-scale solar power producer would pay for the option of selling any of its electricity to a utility.

---

[27] In New Jersey 84 percent of recorded residential projects involve "third-party ownership."

Advocates for smaller incentives have become a strong political force in several European countries. Often, they align with advocates for fairness in the distribution of government benefits across income levels, who observe that current programs mostly benefit wealthy people, because they reward those who are more willing to take risks, who have stronger credit ratings, and for whom a tax deduction is worth more.

On the other side of this argument, pressing for a continuation of the incentives for distributed solar power throughout the world, are the manufacturers, distributors, and installers of distributed solar power. They emphasize that quite soon distributed electricity storage may be twinned with distributed solar power, at which point distributed energy will be able to relieve bottlenecks and provide resilience. They note that every national energy system is replete with incentives of many kinds, and thus the incentives for solar power primarily offset the incentives given to its competitors. They have allies among those who give priority to environmental objectives and maintain that solar incentives are a proven mechanism for achieving cleaner air and less rapid climate change.

Utility Ownership of Distributed Generation versus Ownership by Others

Two alternative ownership patterns for distributed energy are in contention: ownership by utilities and ownership by others. Advocates for ownership exclusively by utilities point to the efficiencies achievable when a single owner optimizes the entire system. North Carolina is one state that opted for utility ownership of distributed energy production, for example. Advocates for diverse ownership emphasize that the system encourages competition and can be expected to lead to lower costs and greater innovation.

In most states, utilities have not made a priority of owning decentralized electricity. Instead, they have urged regulators, legislators, and the public to pay attention to the risk of financial collapse of the grid, unless subsidies for dispersed ownership are reduced. They point to a “death spiral”: demand for utility power falls, the costs of maintaining the grid remain constant, prices rise for the remaining participants, and demand falls further. Demand falls as some customers leave the grid-connected system entirely and others produce substantial amounts of power on their own while remaining on the grid. Prices rise as the grid covers its total costs from the sale of fewer units of electricity. Eventually, the grid crumbles. Some argue that the death spiral is already underway.

Intrinsic Value in Distributed and Centralized Generation

An argument about intrinsic value runs beneath the surface of policy debates about distributed versus centralized energy. Proponents of distributed energy affirm that it enhances the positive values of self-reliance and self-sufficiency, whether at the level of individual households or small communities. Proponents of centralized generation see a well-maintained grid as a social structure that enhances the positive value of broad-based mutual dependency. They also see virtue in the specialization that enables the few with special skills to free the many to pay attention to other things. Arguments for and against centralization are far from unique to solar power. They appear in similar form in debates over the structures of grids for food, water, wastes, and the communication of information.

Grid Parity

There is much talk today of a “breakeven price” or “grid parity” for solar power, where a kilowatt-hour of solar power costs no more than a kilowatt-hour of power from, say, natural gas or coal. What is usually being compared is the “levelized cost of electricity,” which is the total cost for building and operating a power-generating facility, divided by the amount of electricity the facility produces over its operating life. Comparing the levelized costs for a solar power plant and a fossil fuel power plant, therefore, requires assumptions not only about the two capital costs (where the dramatic cost reductions for solar power enter the comparison), but also about the price of fuel, the number of years that the plants operate, and – crucially – the capacity factors of the two plants (where the limited availability of the solar plant enters the comparison).

This definition of “parity” is inadequate, because the levelized cost ignores the grid as a whole. Levelized costs take into account the number of hours that a power plant operates over a year, but not the characteristics of those hours. A more compelling comparison would account for the costs associated with grid integration, which will generally be larger for solar power than for fossil fuel power, given solar power’s intermittency and unpredictability. Adding complication, grid integration costs can be fully evaluated only when the complementary providers of power to the grid are also specified.

The levelized cost also ignores costs associated with today’s electricity generation that are not fully priced (the system’s “externalities”), such as damage to public health and the environment. Solar power generally reduces these costs substantially. A fascinating question is whether solar power will dominate the world’s electric power system by mid-century. That will depend on whether the positive environmental benefits of solar power can more than offset the costs of compensating for its variability.

Appendix: The Princeton University Solar Project

Many of the general issues discussed in the five articles of this Distillate are illustrated here with specific reference to one project that the authors know intimately—the Princeton University solar project. The project produces 5.4 megawatts of peak power and occupies 25 acres (10 hectares) of university land. It has been operational since September 2012 and has been meeting approximately 5 percent of the university's annual electricity consumption. Figure A-1 shows an aerial view of the project.

Photo by Tom Grimes.

Figure A.1 Aerial view of the Princeton University photovoltaic field. The northern-most and western-most solar panels that appear darker are fixed-tilt panels; all of the others are tracking panels.

Princeton University is allowed to sell to the grid the solar electricity it generates at its mid-scale field. However, such sales practically never happen. The electricity used by the buildings served jointly by the university's solar electricity system and the utility nearly always exceeds the output of the university's system. On rare weekends in the spring, demand is low enough and the university's solar supply is high enough for the university to be able to export electricity, but in the first 2.5 years after the university's project started to produce electricity, only about one three-thousandth of the university's solar electricity was sold to the grid.

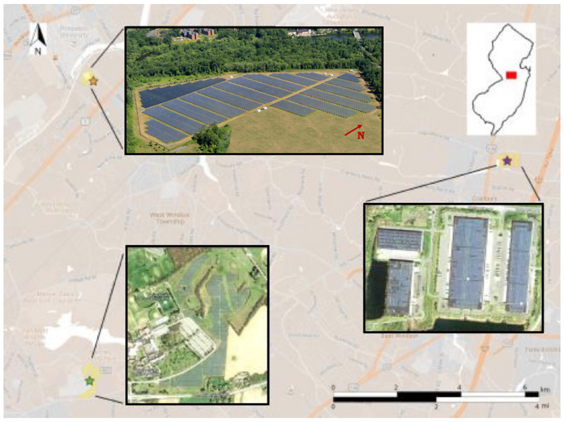

Figure A.2: Three mid-scale projects in central New Jersey: a ground-mounted system at Princeton University (top left, 5.4 MW), a ground-mounted system at Mercer County Community College in West Windsor (bottom left, 8 MW), and a system on the roofs of warehouses in an industrial zone in Cranbury (center right, 5.7 MW). The red rectangle in the upper right inset shows the location of the background map (10 miles by ten miles) within New Jersey. The photos of the Princeton University and Cranbury sites have the same scale; the photo of Mercer Community College site has been shrunk relative to the other two.

The university's project is the 37th largest solar project in New Jersey by capacity. A sense of the spatial density of large mid-scale projects in central New Jersey is conveyed in Figure A.2, which shows the Princeton project and the two other mid-scale projects of comparable size in the same area, 10 miles by 10 miles. Princeton's project is shown in at the top left. At bottom left, an 8-megawatt installation sits between Mercer County Community College's West Windsor campus and a large park. It is a joint endeavor between the college, the Mercer County Improvement Authority, and SunLight General Capital, LLC.[28] At center right, is a 5.7 megawatt project that has placed panels on the roofs of four adjacent warehouses.

---

[28] http://apps3.eere.energy.gov/greenpower/news/news_template.shtml?id=1856

Panels and Mounts

The Princeton University panel has 16,500 photovoltaic panels. 80 percent of these panels are mounted to provide single-axis tracking (lighter areas in Figure A.1) and the rest have fixed-tilt mounts (darker areas in Figure A.1). Decisions about where to locate the two mountings were driven largely by the shape of the available parcel of land.

The tracking panels are arrayed in rows that move like a seesaw, with the seesaw mount oriented north and south. The compass at the top-right of Figure A.1 confirms this orientation: each of the thin rectangles is one of these seesaws, rotated maximally to the east at sunrise, flat at noon, and rotating toward the west throughout the afternoon. As for the fixed-tilt panels, they face south, making at a 25 degree angle with the horizontal.

The tracking panels are further apart than the fixed panels, because the spacing required to avoid the shadowing of one panel by another is larger for tracking than for fixed panels. The tracking panels occupy approximately three times as much land area as the total active surface area of the panels, and for the fixed panels the multiplier is less, two instead of three.

The panel chosen for the Princeton University project has 96 monocrystalline silicon cells in an 8x12 array, roughly 1.0 meter by 1.5 meters in size. The panel's rated peak power output is 327 watts, the product of 6.0 amps of current and 54.7 operating volts.

Distribution of Initial Construction Costs

The total construction cost of the Princeton University project was approximately 30 million dollars, or about six dollars per watt. This cost is disaggregated in Figure A.3. The PV modules themselves account for 30 percent of the total, and balance of systems costs account for 70 percent. Non-panel hardware accounts for 16 percent; included are tracking equipment, inverters, transformers, mounting structures, motors, concrete, and fencing. There are also substantial site-specific costs, 12 percent of the total, as is usual for large, ground-mounted systems. The location of the solar field across a lake from the main campus electrical substation to which the solar field is attached mandated the placement of a 13 kilovolt cable (Figure A.4) in a conduit under the lake that lies between the solar field and campus. Other costs, including labor, project management, and permitting, account for the remaining 41 percent.

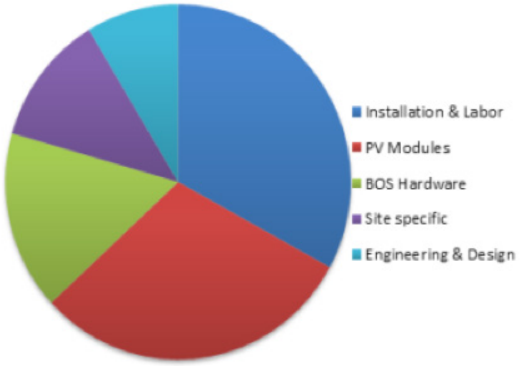

Figure A.3: Cost breakdown for the installation of Princeton University's 5.4-megawatt solar field completed in 2012. BOS is "balance of systems." The total cost was about 30 million dollars.

Figure A.4 Photo shows a cross section of the 13 kilovolt cable that connects the solar field to main campus and which passes underneath Lake Carnegie.

Variability at Various Scales

Variability in the project's electricity production at various time scales can be documented with the help of the detailed performance data that Princeton University is recording.[29] We distinguish four time scales here: seasonal variability, day-to-day variability, variability during the day, and variability in minutes.

Seasonal Variability

Figure A.5 plots the output of both the tracking and the fixed-tilt panels, averaged over a week for the 52 weeks of 2014. For both the tracking and fixed panels, the rate of production of solar power from all of the panels (in kilowatts) is divided by the total area of panels (in square meters). A tracking panel produces more power than a fixed panel in the spring and summer, but the two mountings perform approximately equally in the fall and winter. A tracking panel's output is about six times as large near the summer solstice (and also in one week in April) as near the winter solstice. Although tracking panels are more productive, adding tracking hardware

---

[29] Gokhale, Manali P, 2015. "An Analysis of Princeton University's Photovoltaic Field," Junior Paper, Spring 2015 (unpublished).

(including motors) adds costs. The trade-off is not clear cut. Indeed, utility-scale solar installations are being built today using both tracking and fixed-panel arrays.

The weekly data in Figure A.5 exhibit some bumpiness, rather than fitting a smooth curve. The bumps are rainy and cloudy weeks. Week-to-week variability, of course, is less predictable than the seasonal variation and nonetheless needs to be accommodated on the grid. In this instance, the costs of grid integration can be substantially reduced by good weather forecasting.

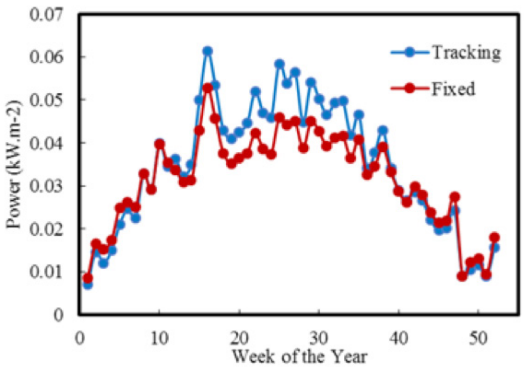

Figure A.5: Average weekly power production of tracking and fixed-tilt panel, divided by total panel surface area, for each week of 2014. Monthly averages were substituted for missing days of data.

Day-to-Day Variability

Figure A.6 shows, for the 365 days of 2014, hourly solar output from the tracking panels. The analogous record for the fixed-tilt panels looks similar, except for a difference in daily load shape that is discussed below. Looking at the entire year reveals intermittency on several different scales. Intermittency within the hours of the day is obvious; it is as predictable as sunrise and sunset. Day-to-day intermittency is much less predictable: a sunny day can be followed by a cloudy one during which hardly any power is produced. Note the four days in a row in the first week of January when almost no power was produced; weather data reveal that these were snowy days, and the panels may have been covered in snow. Deliberately, Figure A.6 does not display seasonal variability, because the peak value for each month has been made the same, 180 watts per square meter of panel area. This scaling enables the patterns in the winter and summer to be equally prominent. In fact, as seen above in Figure A.5, far more power is produced in the summer.

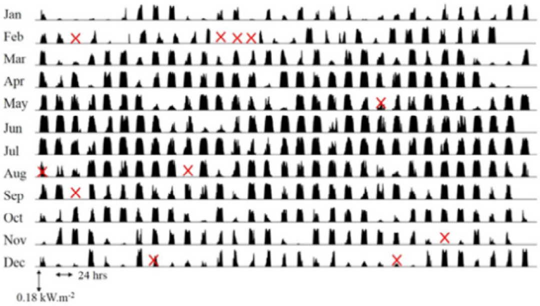

Figure A.6: Power production profiles for tracking panels in Princeton University's 5.4 megawatt solar installation for each day in 2014. Each month's peak hourly output is rescaled to be 180 watts per square meter of panel area. Red crosses indicate days when data are missing for at least part of the day.

Variability During the Day

The hourly profile for solar-panel output is at the heart of the challenge of integrating solar power into the grid, as the Duck Curve (Figure 5.1) makes clear. Detailed insight into these profiles can be obtained from the average hourly output of the identical panels equipped with two different kinds of mountings installed at the Princeton University project. Figure A.7 displays the average profiles for the fixed-tilt and tracking mountings on two specific sunny days (July 6 and December 26, 2014).

The most important feature is shared by both kinds of panels. As in the Duck Curve, the decline in output at the end of the afternoon is very steep – on both the winter and the summer days and for both tracking and fixed-tilt panels. On both days, however, the shape of the curve through the day is quite different for the tracking and fixed panels. On July 6 the output of the tracking panel power has a flat top, constant over much of the day. The Sun is high in the sky and high above the seesaw. By contrast, the fixed panel has a prominent peak at solar noon, when the Sun is highest in the sky, but the output drops off quickly on both sides of noon. Over the day, consistent with Figure A.5, the tracking panel collects more sunlight.

December 26th tells a different story. The fixed-tilt panels do better over the day than the tracking panels. The difference is most pronounced at noon, when the Sun is low in the southern sky, and the single-axis tracking panels are fully horizontal. The fixed panels, because they are tilted 25 degrees toward the south, see the noon Sun at a less oblique angle than the

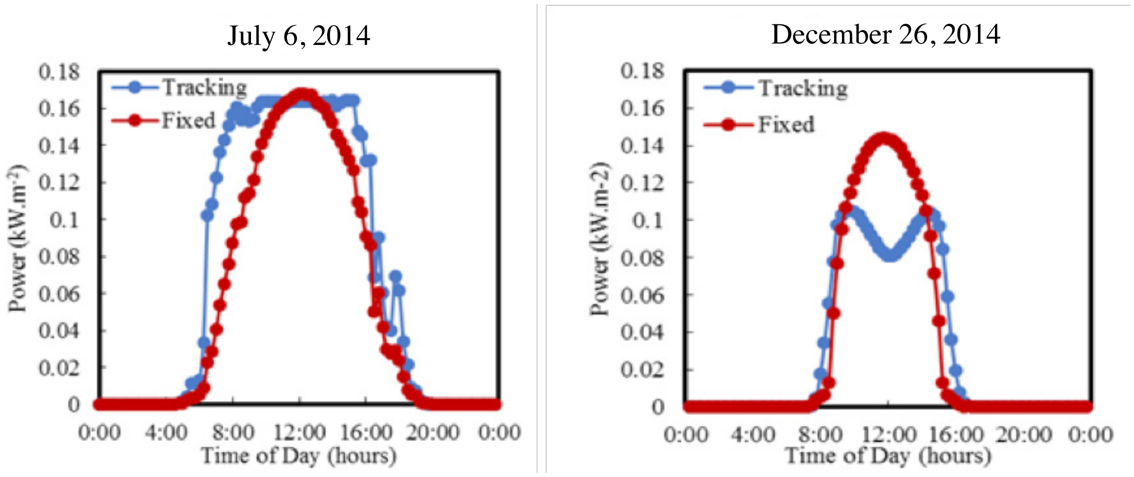

Figure A.7: Power production curves for tracking and fixed-tilt panels on July 6, 2014 (a sunny summer day, left panel) and December 26, 2014 (a sunny winter day, right panel). Power output (in kilowatts) is divided by the total surface area (in square meters) for each type of panel.

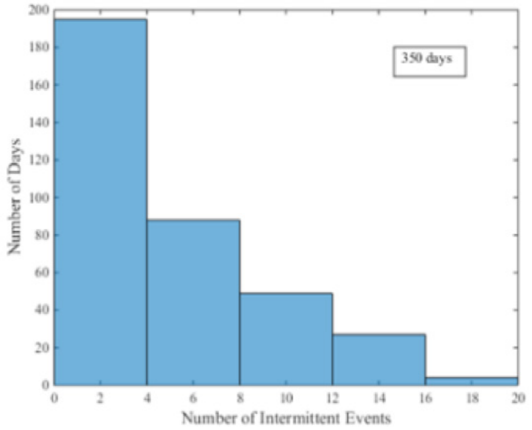

Figure A.8. Histogram of the number of days in 2014 with a certain number of rapid-change events, defined by output changing up or down by more the 30 watts per square meter of panel in 15 minutes. The first bin has zero to three events, the second has four to seven events, etc.

tracking panels and therefore produce more power. Note the "W" shape for the tracking-panel output: the tracking panels see the Sun better on both sides of noon than at noon, in mid-morning they look somewhat upward to the east, and at mid-afternoon somewhat upward to the west.

Variability in Minutes

Less amenable to help from weather forecasting than week-to-week or day-to-day variability are the rapid variations in solar input during partly cloudy days. Figure A.8 quantifies this form of variability by documenting the events during an entire year when power output from the solar field changed rapidly, up or down. Specifically, a rapid-change event was quantified as a change in power output of more than 30 watts per square meter of panel area in a span of 15 minutes – the time interval of our data. Because the maximum rate of production of electricity by the solar field is about 190 watts per square meter of panel area, what we are defining as a rapid change is a change of about 15 percent of the maximum possible change. Our analysis was restricted to the tracking panels. The number of such events per day, for 350 out of the 365 days in 2014 (the other 15 days had missing data), is plotted as a histogram with five bins. Nearly half of the days (roughly 160 days) have four or more rapid-change events. Many of these are the particularly troublesome partly cloudy days of the year.

Disposition of Princeton University's Solar Renewable Energy Credits

Princeton University earns New Jersey Solar Renewable Energy Certificates by operating its solar field. It currently sells these certificates to New Jersey's energy providers in a certificates market, enabling these providers to meet their mandated minimum production of solar power. However, the University is assessing the case for taking its certificates off the market in a few years and foregoing this revenue stream, in order to increase the amount of solar electricity produced in New Jersey. Those putting forth this argument assert that when the University participates in the certificate market, another potential producer of solar power will not do so, because the certificate market is designed to produce a specific amount of solar energy (its “solar carve-out”), not more. Selling its certificates, therefore, does not increase the amount of solar energy produced in New Jersey. Only when the University retires its credits rather than selling them, will other projects come into existence, with equivalent solar energy production, to meet New Jersey’s mandated minimum requirement for solar purchases.

This argument assumes that the University’s project by itself has no effect on the size of New Jersey’s solar carve-out, and even its existence. The argument can be countered, however, if each New Jersey project, including the one at Princeton University, affects the overall level of interest in solar energy in New Jersey, and if that level of interest affects the targets set by the New Jersey’s government. Targets are set many years ahead and will be revised only infrequently, but they can be revised upward or downward and can be challenged in state courts.