Relating Product Prices to Long-Run Marginal Cost: Evidence from Solar Photovoltaic Modules

Stefan Reichelstein*

Graduate School of Business, Stanford University

and

Department of Business Administration, University of Mannheim

Anshuman Sahoo

Steyer-Taylor Center for Energy Policy and Finance, Stanford University

December 2016

*Contact information: reichelstein@stanford.edu; asahoo@stanford.edu

We gratefully acknowledge financial support from the Steyer-Taylor Center for Energy Policy and Finance and a U.S. Department of Energy (DOE) grant administered through the Bay Area Photovoltaics Consortium (BAPVC). We thank seminar participants at Stanford, New York University, Southern Methodist University, the University of Vienna, the University of Waterloo and the Antitrust Division of the U.S. Department of Justice for their helpful comments and suggestions. We also acknowledge helpful comments from two anonymous reviewers, Tim Baldenius, Jeremy Bertomeu, Shane Dikolli (Ad-hoc Editor), Gunther Glenk, Alexander Nezlobin, John Parsons and Thomas Pfeiffer. Finally, we acknowledge the assistance of Aneesh Raghunandan with our data collection.

Relating Product Prices to Long-Run Marginal Cost: Evidence from Solar Photovoltaic Modules

December 2016

Abstract: A basic tenet of microeconomics is that for a competitive industry in equilibrium the market price of a product will be equal to its marginal cost. This paper develops a model framework and a corresponding empirical inference procedure for estimating long-run marginal cost in industries where production costs decline over time. In the context of the solar photovoltaic (PV) module industry, we rely primarily on firm-level financial accounting data to estimate the long-run marginal cost of PV modules for the years 2008–2013. During those years, the industry experienced both unprecedented price declines and significant expansions of manufacturing capacity. We compare the trajectory of average sales prices with the estimated long-run marginal costs in order to quantify the extent to which actual price declines were attributable to reductions in production costs. The trajectory of estimated product costs is then extrapolated to forecast an equilibrium trend line for future PV module prices.

Keywords: Long-run Marginal Cost, Cost Estimation, Learning-by-doing, Price Dynamics

JEL codes: D41, L11, L63, M21, Q42

1 Introduction

Economic theory submits that competition drives the equilibrium price of a product to its long-run marginal cost. This paper proposes and implements a method for estimating long-run marginal cost on the basis of firm-level data obtained from financial statements. The resulting trajectory of cost estimates can be compared to actually observed sales prices, with the resulting difference serving as a measure of the extent of disequilibrium at different points in time. In industries where costs decline over time, the time series of long-run marginal cost estimates can then be extrapolated to forecast future costs and equilibrium product prices.

The identification and measurement of marginal cost remains a matter of debate in industrial organization.[1] The principal issue in identifying a product's long-run marginal cost is the inclusion of capital assets such as facilities, machinery and equipment. One common approach in economic textbooks is to assume that capital is a consumable input, similar to raw materials or labor. In contrast, Jorgenson (1963) and Arrow (1964) pioneered capital accumulation models in which a representative firm makes periodic investments in new capacity and thereby creates a capital stock. In these infinite horizon models models, it becomes possible to identify the marginal cost of one unit of capacity that is available to the firm for one period of time.

The front part of our analysis examines a dynamic model of a competitive industry in which firms make sequential and overlapping capacity investments and subsequently choose their periodic output levels in a competitive fashion, taking market prices as given.[2] While the long-run marginal cost contains components that are sunk in the short-run (e.g, capacity costs), the expected market prices will nonetheless be equal to the long-run marginal cost in equilibrium, because firms are capacity constrained in the short run. Furthermore, firms will earn zero economic profits on their capacity investments if the market prices in future periods are equal to the long-run marginal cost in those future periods. These characterizations apply in particular to industries in which costs decline over time.

We describe a procedure for estimating the long-run marginal cost of a product based

[1] See, for instance, Carlton and Perloff (2005), Pittman (2009), McWatters and Zimmerman (2015) and Rogerson (2011).

[2] Recent work in economics and accounting has built upon the framework of Jorgenson (1963) and Arrow (1964) in connection with managerial performance evaluation and profitability analysis; see, for example, Rogerson (2011), Rajan and Reichelstein (2009), Nezlobin (2012) and McNichols, Rajan, and Reichelstein (2014).

primarily on firm-level financial data. We then apply the proposed procedure in the context of the solar photovoltaic (PV) module industry which has experienced sharp price declines and rapid output growth in recent years. Figure 1 plots the history of (the logarithm of) average sales prices against (the logarithm of) cumulative output for the years 1979 to 2010. The corresponding price trajectory conforms closely to an 80% constant elasticity learning curve, an observation that is frequently attributed to Swanson (2011).[3] Accordingly, prices drop by 20% with every doubling of cumulative output, measured in megawatts (MW).

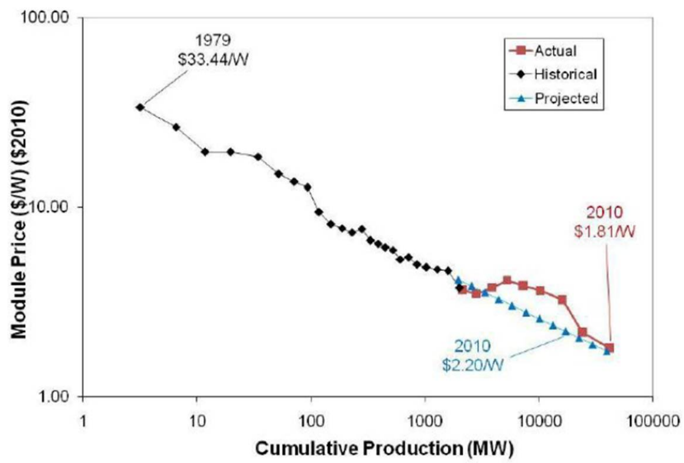

Figure 1: *Plot from Swanson (2011)*

Figure 2 extends the original Swanson plot beyond 2010 and shows that between 2011 and 2013 the decline in average sales prices (ASP) for PV modules was substantially steeper than that predicted by the historical 80% learning curve. Particularly noteworthy is the 40% price drop in 2011 alone and the rebound in prices for late 2013. Industry analysts have pointed out that the steep price declines in recent years may reflect at least in part that the additions to industry-wide manufacturing capacity were excessive.[4]

In most manufacturing industries, including solar PV modules, long-run marginal cost

[3]Figure 1 shows that for the years 2008-2009, ASPs were distinctly above the trend line suggested by the 80% learning curve. Most industry observers attribute this discrepancy to an acute polysilicon shortage which temporarily increased the raw material cost of silicon wafers.

[4]Recent studies, like Candelise, Winskel, and Gross (2013), have pointed out the difficulty in attributing the dynamics of observed sales prices to intrinsic cost reductions as opposed to broader industry level effects. In particular, these authors state: "Overall, it is not straightforward to fully disentangle module price reductions due to reduced production costs related to device and production process improvements and economies of scale along the PV module chain from market demand/supply dynamics, including manufacturers strategies...." (page 100).

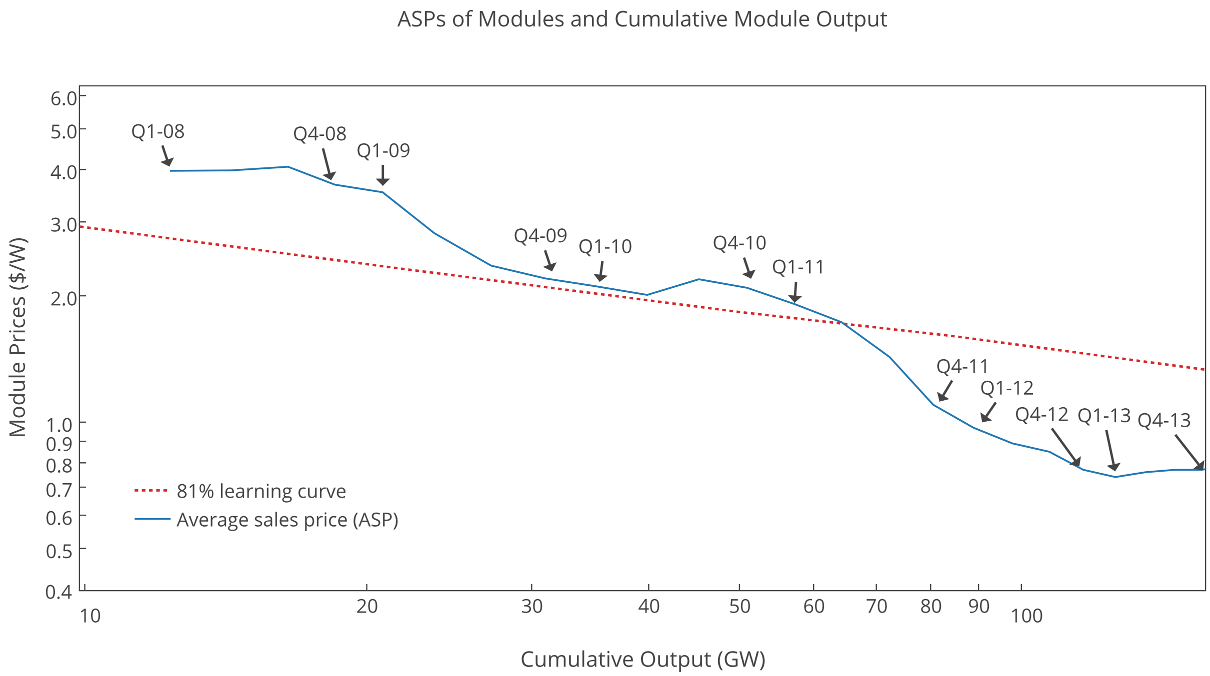

Figure 2: *Predicted and observed ASPs, 2008 – 2013. All prices are in 2013 U.S. dollars.*

comprises capacity related costs for machinery and equipment, current manufacturing costs for materials, labor and overhead as well as periodic costs related to selling and administrative expenses. For the firms in our sample, we infer production costs from quarterly financial statements, primarily based on cost of goods sold, finished goods inventory balances, capital expenditures and SG&A expenses. In addition, our cost inference procedure relies on quarterly data for manufacturing capacity and product shipments.

In applying our cost estimation procedure to solar PV modules, we obtain a close match between average sales prices and the estimated long-run marginal costs for the years 2008 - 2010.[5] Beginning in late 2011, however, the dramatic decline in the observed ASPs for most of the quarters in 2012-2013 result in average sales prices significantly below our estimates of the contemporary long-run marginal costs.[6] In other words, the sharp drop in ASPs for those time periods points to excessive additions in manufacturing capacity rather than solely to

[5]The solar PV industry satisfies the criteria of a competitive industry insofar as a large number of firms in the industry supply a relatively homogeneous product. To note, the median market share of firms in this industry was less than 1% in 2012.

[6]This conclusion is corroborated by the sharply negative earnings and declining share prices that firms in the industry experienced during those two years.

cost reductions. The difference between the ASPs and our long-run marginal cost estimates quantifies the price effect that is attributable to excess capacity in the industry at different points in time.

Despite our conclusion that observed sales prices were not consistent with equilibrium predictions for some of the quarters in our observation window, our econometric results also point to a rate of cost reductions that is faster than suggested by the 80% learning curve. In particular, our estimates point to a 62% constant elasticity learning curve for *core manufacturing costs*, which comprise materials, labor, and manufacturing overhead, excluding depreciation charges. At the same time, we find that capacity related costs for machinery and equipment have fallen at a rate which, given recent industry output, also outperforms the 80% learning rate benchmark. Taken together, these results yield a forecast for the trajectory of future long-run marginal cost, and therefore also for future equilibrium module prices, that is steeper than the traditional learning curve associated with this industry.[7]

Our research design is applicable beyond the solar PV industry. Since the empirical inference procedure outlined in this paper is based on firm-level accounting data, any long-run marginal cost estimate will, ceteris paribus, become more reliable if the product or service in question is (i) fairly homogeneous across suppliers and (ii) constitutes the dominant line of business for firms in the sample. From that perspective, semiconductors, chemicals, aircraft manufacturing and steel would be other natural candidates. In industries where firms typically deliver a heterogeneous mix of products and/or services, one would need to conduct the empirical tests either for product aggregates or obtain access to line-of-business segment reports.[8]

Our analysis is related to earlier work in managerial accounting that has sought to justify the use of full cost for product pricing purposes; see, for instance, Banker and Hughes (1994), Balachandran, Balakrishnan, and Sivaramakrishnan (1997), Göx (2002), Balakrishnan and

[7]The cost inference method we employ to estimate the ESP complements so-called "bottom-up" cost models, e.g., Powell et al. (2012), Powell et al. (2013), Goodrich et al. (2013a), and Goodrich et al. (2013b). These studies estimate costs by aggregating input requirements and input prices as reported by various industry sources. In contrast to our approach, though, these studies provide a snapshot of different costs at particular point in time, rather than a dynamic cost model in which equilibrium prices reflect anticipated future cost reductions. Other studies examining the costs and prices of solar PV modules include Pillai and McLaughlin (2013) who examine the mark-up that firms charge over and above cost of goods sold (COGS).

[8]Previous studies in accounting, like Sridharan (2015), have used firm-level accounting information to predict market trends and volatility. Our approach relies on firm-level income statements and balance sheet information to identify economically relevant costs in a particular industry.

Sivaramakrishnan (2002), and Narayanan (2003). A common theme in this literature is that, while capacity related costs are sunk in the short-run, they are nonetheless "relevant" for operational decisions like product pricing. These studies typically have focused on a firm with monopoly power in a static setting, that is, the firm is assumed to make a one-time decision regarding its capacity choice.[9] In contrast our dynamic model of a competitive industry departs from the hypothesis that the long-run marginal cost of a product should be equal to the expected equilibrium price. While the long-run marginal cost does include capacity related costs, it does not coincide with the measure of full cost, as usually defined in managerial accounting textbooks.

Finally, our findings contribute to the well established literature on learning-by-doing; see, for instance, Spence (1981), Lieberman (1984), Dick (1991) and Benkard (2000). Common to these studies is the notion that production costs decline over time and that the rate of decline is driven by cumulative industry output. For the most part, these studies seek to infer production costs from observed prices, assuming that in an oligopolistic equilibrium firms will be able to charge certain mark-ups on their costs. Our approach, in contrast, is to infer production costs from firm-level financial reports and to compare these inferences with actually observed product prices.[10]

The remainder of the paper is organized as follows. Section 2 formulates the model and derives the long-run marginal cost for an industry with declining production costs. Section 3 describes our inferential procedure for deriving the long-run marginal cost from firm-level accounting data. We then apply this procedure in the context of solar module manufacturing in Section 4, and compare our cost estimates to the observed ASPs. Section 5 presents our econometric estimates of recent learning effects in PV module manufacturing and applies these estimates to extrapolate a trajectory of future costs. We conclude in Section 6. The Appendix presents proofs, data sources, descriptive statistics and robustness checks.

[9]The rationale for full cost pricing has also been examined in connection with transfer pricing, as many companies appear to rely on full cost as the basis for valuing intracompany transfers; see, for instance, Pfeiffer, Schiller, and Wagner (2011), Dutta and Reichelstein (2010), Bouwens and Steens (2016), and the references therein.

[10]Some textbooks in managerial accounting, e.g., Hilton (2005), emphasize the need for cost analysis to incorporate learning effects in order to obtain better forecasts of efficiency improvements.

2 Long-Run Marginal Cost and Competitive Pricing

2.1 Base Model

The model framework developed in this section identifies the trajectory of long-run marginal costs for an industry characterized by declining production costs. We consider a model of an industry composed of a large number of suppliers who behave competitively. Firms are assumed to be capacity constrained in the short-run, such that output supplied to the market in a particular period is limited to the overall capacity that the firm has installed in previous periods. Production capacity available at any given point in time thus reflects the cumulative effect of past investments.

Our model feature of overlapping capital investments is in the tradition of the capital accumulation literature, as pioneered by Jorgenson (1963) and Arrow (1964). In their model formulations the cost of acquiring new capacity is assumed to be constant over time.[11] In addition to capacity costs, our model also includes variable production costs as well as fixed operating costs and corporate income taxes. The model presented here is arguably closest to Reichelstein and Rohlfing (2015), except that their framework effectively allows for only one stage of upfront capacity investments. In our infinite horizon framework with periodic capacity investments, in contrast, the long-run marginal cost of a product must reflect the anticipated decline in the expenditures required for future capacity additions.

In the base version of the model, firms can accurately predict future demand. Let $P_t^o(Q_t)$ denote the aggregate willingness-to-pay (inverse demand) curve at time t, where Q_t denotes the aggregate quantity supplied at date t. Market demand is assumed to be decreasing in price and, in addition, we postulate that demand is expanding over time in the sense that:

$$P_{t+1}^o(Q) \geq P_t^o(Q), \tag{1}$$

for all $t \geq 1$ and all Q. The significance of this condition is that if firms make investments sufficient to meet demand in the short-run, they will not find themselves with excess capacity in future periods.[12] This condition appears plausible in the context of solar PV modules,

[11]Rogerson (2008, 2011) generalizes the earlier studies by allowing for productive capacity to diminish in an arbitrary (rather than geometric) fashion. Rogerson's model also allows for the acquisition cost of new capacity to decline over time. All of these models examine optimal investment decisions for a representative firm, without solving for an industry equilibrium.

[12]Rogerson (2008) refers to (1) as the No-Excess Capacity (NEC) condition. See also Rajan and Reichel-

particularly for the time period covered in our empirical analysis.

In order to break even on their capacity investments, firms will, in equilibrium, realize a stream of revenues that cover their long-run marginal costs at each point in time. This cost comprises capacity related costs, periodic operating costs, and costs related to income tax payments. At the initial date 0, the industry is assumed to have a certain capacity stock in place. To acquire one unit of production capacity, firms must incur an investment expenditure of v at the initial date 0. We allow for technological progress to lower the cost of capacity acquisitions over time. For reasons of tractability, though, we confine attention to a single "technological progress parameter", η, leading to a pattern of geometric declines such that $\eta^t \cdot v$ denotes the acquisition cost for one unit of capacity at time t, with $\eta \leq 1$.[13] Accordingly, investment decisions and the subsequent level of aggregate capacity in the market are conditional on firms' expectation of future decreases in capacity costs.

Investments in capacity represent a joint cost insofar as one unit of capacity acquired at time t will allow the firm to produce one unit of output in each of the next T years.[14] To identify equilibrium prices in terms of costs, it will be useful to introduce the marginal cost of one unit of capacity *made available for one period of time.* As shown by Jorgenson (1963), Arrow (1964) and Rogerson (2008), this effectively amounts to "levelizing" the initial investment expenditure. To that end, let $\gamma = \frac{1}{1+r}$ denote the applicable discount factor. The marginal cost of one unit of capacity in period t then becomes:

$$c_t = \frac{\eta^t \cdot v}{\sum_{\tau=1}^{T} (\gamma \cdot \eta)^\tau}. \tag{2}$$

An intuitive way to verify this claim is to assume that firms in the industry can rent capacity services on a periodic basis. Assuming this rental market is competitive and capacity providers have the same cost of capital, it is readily verified that the capacity provider who invests in one unit of capacity at time t and then rents out that capacity in each of the next T years for a price of $c_{t+\tau}$ would exactly break even on his initial investment of $\eta^t \cdot v$.

stein (2009) and Dutta and Reichelstein (2010).

[13] Decreases in capacity cost as a function of time can be attributed to improvements in manufacturing equipment.

[14] For simplicity, we adopt the assumption that productive capacity remains constant over the useful life of a facility. In the regulation literature, this productivity pattern is frequently referred to as the "one-hoss shay" model; see, for instance, Rogerson (2008) and Nezlobin, Rajan, and Reichelstein (2012).

Accordingly, the earlier capital accumulation literature refers to c_t as the *user cost of capacity.*

In any given period, firms are assumed to incur fixed operating costs, e.g., maintenance, rent and insurance, in proportion to their incumbent production capacity. Like past investment expenditures, these costs are assumed to be "sunk" after date t because they are incurred regardless of capacity utilization. Formally, let f_t represent the fixed operating cost per unit of capacity available at time t, with $f_{t+1} \leq f_t$ for all $t \geq 1$. Finally, production of one unit of output entails a constant unit variable cost, w_t, which again is assumed to be weakly decreasing over time, that is, $w_{t+1} \leq w_t$ for all $t \geq 1$. In contrast to the fixed operating costs, variable costs are avoidable in the short-run if the firm decides not to utilize its available capacity.

Corporate income taxes affect the long-run marginal cost of production through depreciation tax shields and debt tax shields, as both interest payments on debt and depreciation charges reduce the firm's taxable income. Following the standard corporate finance approach, we ignore the debt related tax shield provided the applicable discount rate, r, is interpreted as a weighted average cost of capital. The depreciation tax shield is determined by both the effective corporate income tax rate and the allowable depreciation schedule for the facility. The effective corporate income tax rate is represented as α (in %), and d_t denotes the percentage of the initial asset value that is the allowable tax depreciation charge in year t, $1 \leq t \leq T$.

The assumed useful life of an asset for tax purposes is usually shorter than the asset's actual economic useful life, which we denote by T in our model. Accordingly, we set $d_t = 0$ for those periods that exceed the useful life of the asset for tax purposes. As shown below, the impact of income taxes on the long-run marginal cost can be summarized by a *tax factor* which amounts to a "mark-up" on the unit cost of capacity, c_t.

$$\Delta = \frac{1 - \alpha \cdot \sum_{t=1}^{T} d_t \cdot \gamma^t}{1 - \alpha}. \tag{3}$$

To provide intuition for the expression in (3), we note that corporate income taxes would not affect the economically sustainable price if taxable income were calculated on cash flow basis. In that (hypothetical) case, both the present value of pre-tax cash flows and taxable incomes would be equal to zero if the product price is equal to the economically sustainable price. Thus firms would break even both on a pre-tax basis and after taxes. Accordingly, Δ

would be equal to 1 if the tax code were to allow for immediate full expensing of investments and therefore taxable income would be equal to pre-tax cash flow. However, since the tax code only allows for a delayed write-off of the capital expenditure, corporate income taxes will generally introduce an additional cost factor.[15]

We are now in a position to introduce the following overall measure of unit cost:

$$LMC_t = w_t + f_t + c_t \cdot \Delta. \tag{4}$$

To justify the label long-run marginal cost (LMC_t) in (4), we demonstrate below that for a competitive industry in equilibrium the prevailing product price will indeed be equal to LMC_t at each point in time. We note that the first two terms on the right-hand side of (4) are cash costs, while the capacity related term represents an allocated cost.[16] We note that the expression in (4) is rather similar to the levelized product cost measure introduced in Reichelstein and Rohlfing (2015). The key difference, though, is that they consider a finite horizon model with a single upfront investment decision. As a consequence, capacity costs do not change over time so that the unit cost, c, remains time-invariant.

Given our assumption of a competitive fringe of suppliers, the investment and capacity levels of individual firms remain indeterminate. Denoting the aggregate industry-wide investment levels by I_t, the "one-hoss shay" assumption that productive assets have undiminished productivity for T periods implies that the aggregate capacity at date t is given by:

$$K_t = I_{t-T} + I_{t-T+1} + \ldots + I_{t-1}. \tag{5}$$

Equation (5) holds only for $t > T$. If $t \leq T$, then $K_t = I_0 + I_1 + ... + I_{t-1}$.

Firms choose their actual output in a manner that is consistent with competitive supply behavior. Since capacity related costs and fixed operating costs are sunk in any given period,

[15]Specifically, $\Delta = 1$ if $d_0 = 1$ and $d_t = 0$ for $t \geq 1$. Holding α constant, a more accelerated tax schedule tends to lower Δ closer to one. To calibrate the magnitude of this factor, for a corporate income tax rate of 35%, and a tax depreciation schedule corresponding to a 150% declining balance rule over 20 years, the tax factor will approximately amount to $\Delta = 1.3$.

[16]The main contribution in Rogerson (2008) is that the cap0acity related charges, as introduced in (2), can be expressed in conventional accounting terms as "residual income" charges, that is, as the sum of depreciation- and imputed interest charges, provided capacity investment are written off in accordance with a properly chosen depreciation schedule.

firms will exhaust their entire capacity only if the market price covers at least the short-run marginal cost w_t. Conversely, firms would rather idle part of their capacity with the consequence that the market price will not drop below w_t. Given an aggregate capacity level, K_t, in period t, the resulting market price is therefore given by:

$$p_t(K_t, w_t) = \max\{w_t, P_t^o(K_t)\},$$

while the aggregate output level, $Q_t(K_t, w_t)$ satisfies $P_t^o(Q_t(K_t, w_t)) = p_t(K_t, w_t)$. We refer to the resulting output and price levels as *competitive supply behavior*.

Definition 1 $\{K_t^*\}_{t=1}^{\infty}$ *is an equilibrium capacity trajectory if, given competitive supply behavior, capacity investments have a net present value of zero at each point in time.*

Finding 1 *A capacity trajectory* $\{K_t^*\}_{t=1}^{\infty}$ *that satisfies the pricing condition:*

$$P_t^o(K_t^*) = w_t + f_t + c_t \cdot \Delta, \tag{6}$$

at each point in time t, is an equilibrium capacity trajectory.[17]

The equilibrium price characterization in Finding 1 validates our interpretation of $LMC_t \equiv w_t + f_t + c_t \cdot \Delta$ in (4) as the long-run marginal cost of one unit of output. With additional assumptions, the capacity trajectory identified in Finding 1 is also the unique equilibrium capacity trajectory. This is readily seen if one assumes that capacity investments are reversible or, alternatively, that capacity can be obtained on a rental basis for one period at a time, with all rental capacity providers obtaining zero economic profits. Competition would then force the market price for the product in question to be equal to LMC_t in each period.

In our model framework, the unit variable costs, w_t, and the unit fixed cost, f_t, may decline over time, provided the rate of decline taken is viewed as exogenous. We note that under conditions of atomistic competition, that is, no firm can impact the prevailing market price through its own supply decision, Finding 1 also extends to situations where the unit costs decline as a function of the *cumulative volume* of past output levels. One possible formulation is for $w_t = \beta(\sum Q_t) \cdot w$, where $\beta(\cdot) \leq 1$ is decreasing in its argument and $\sum Q_t \equiv \sum_{\tau \leq t} Q_\tau$.

[17] A formal proof of Finding 1 is presented in Appendix A. An implicit assumption here is that the aggregate capacity in place at the initial date does not amount to excess capacity. Formally, we require $P_t^o(K_0) > LMC_1$.

In concluding this subsection, it is instructive to ask how the LMC_t, as presented in (4), relates to the accountant's measure of full cost. As mentioned in the Introduction, earlier literature has argued that the full cost measure can be relevant for making capacity acquisition decisions.[18] For instance, Cooper and Kaplan (1988) state that "...full cost is meant to be a surrogate for long-run manufacturing cost." As conceptualized in most of the managerial accounting literature, full cost differs from LMC_t as the former includes the aggregate depreciation charge instead of the capacity cost term $c_t \cdot \Delta$ in (4). It is readily verified that if (i) $\eta = 1$ and (ii) depreciation is calculated on a straight line basis, the long-run marginal cost LMC_t will exceed full cost (Reichelstein and Rohlfing, 2015). This inequality reflects that full cost does not properly account either for the time value of money or for the effect of income taxes. Correspondingly, the two measures coincide if (i) and (ii) hold and there is no discounting, that is, $\gamma = 1$. On the other hand, with technological progress, that is, $\eta < 1$, the relation between the two cost measures remains indeterminate and ultimately depends on the growth in past investments as well as the magnitude of the parameters T, r, and η.

2.2 Price Volatility

The characterization of equilibrium in Finding 1 can be extended to environments with price volatility. Suppose that, given the aggregate supply quantity Q_t at date t, the price in period t is given by:

$$P_t(\epsilon_t, Q_t) = \epsilon_t \cdot P_t^o(Q_t),$$

where $\tilde{\epsilon}_t$ reflects volatility in the level of demand and is a random variable with mean 1. The support of $\tilde{\epsilon}_t$ is $[\underline{\epsilon}_t, \bar{\epsilon}_t]$, with $0 < \underline{\epsilon} < 1$. The noise terms $\{\tilde{\epsilon}_t\}_{t=1}^{\infty}$ are assumed to be serially uncorrelated, such that each $\tilde{\epsilon}_t$ is observed by all market participants at the beginning of period t. Competitive supply behavior then requires that:

$$p_t(\epsilon_t, w_t, K_t) = \begin{cases} \epsilon_t \cdot P_t^o(K_t) & \text{if } \epsilon_t \geq \epsilon_t(K_t, w_t) \\ w_t & \text{if } \epsilon_t < \epsilon_t(K_t, w_t), \end{cases}$$

[18]See, for example, Banker and Hughes (1994), Balakrishnan and Sivaramakrishnan (1996), Göx (2002), and Narayanan (2003). In managaerial accounting, the full unit cost of a product is usually defined as Cost of Goods Manufactured (direct materials, direct labor, manufacturing overhead) plus SG&A period costs divided by the number of units produced.

where the threshold level of demand volatility is given by:

$$\epsilon_t(K_t, w_t) = \begin{cases} \bar{\epsilon}_t & \text{if } \bar{\epsilon}_t \cdot P_t^o(K_t) \leq w_t \\ \frac{w_t}{P_t^o(K_t)} & \text{if } \bar{\epsilon}_t \cdot P_t^o(K_t) > w_t > \underline{\epsilon}_t \cdot P_t^o(K_t) \\ \underline{\epsilon}_t & \text{if } \underline{\epsilon}_t \cdot P_t^o(K_t) \geq w_t. \end{cases}$$

Given K_t and w_t, the expected market price in period t then becomes:

$$E\left[p_t(w_t, \tilde{\epsilon}_t, K_t)\right] \equiv \int_{\underline{\epsilon}}^{\epsilon(K_t, w_t)} w_t \cdot h_t(\epsilon_t)\, d\epsilon_t + \int_{\epsilon(K_t, w_t)}^{\bar{\epsilon}} \epsilon \cdot P_t^o(K_t) \cdot h_t(\epsilon_t)\, d\epsilon_t. \tag{7}$$

With risk neutral firms, price volatility will not affect the capacity levels obtained in equilibrium provided firms anticipate that they will exhaust the available capacity even for unfavorable price shocks. To that end, we introduce a condition of *limited price volatility*:

$$\underline{\epsilon}_t \cdot LMC_t \geq w_t.$$

Holding the distributions $h_t(\cdot)$ of $\tilde{\epsilon}_t$ fixed, this condition will be satisfied if the short-run avoidable cost w_t constitutes a relatively small percentage of the long-run marginal cost, LMC_t.[19] The implication of this condition is that even for unfavorable price fluctuations firms will still want to deploy their entire capacity.

Corollary to Finding 1 *With limited price volatility, the trajectory identified in Finding 1 remains an equilibrium capacity trajectory. The expected market prices in equilibrium satisfy:*

$$LMC_t = E\left[p_t(w_t, \tilde{\epsilon}_t, K_t^*)\right]. \tag{8}$$

If the above condition of limited price volatility does not hold, the expected equilibrium price will still be equal to the LMC in period t under additional conditions.[20] A natural

[19]As will become clear in the empirical part in Section 4 below, the limited price volatility condition appears plausible in the context of the solar PV module industry for the years 2008-2013. Our estimates suggest that the unit variable cost, w_t, accounts for less than 65% of the total ESP_t. The limited price volatility condition will therefore hold provided $\underline{\epsilon}_t \geq 0.65$. To be sure, some firms may have idled part of their available capacity during those years, but this observation could be attributed to the industry having been out of equilibrium, at least in parts of 2012 and 2013, rather than to more significant unfavorable price shocks.

[20]Reichelstein and Rohlfing-Bastian (2015) establish this result in their one-shot investment model. In related work, Baldenius, Nezlobin, and Vaysman (2016) examine a model of managerial performance evaluation in which the optimal investment policy is such that in response to negative shocks firms will leave parts of their capacity idle in some periods.

candidate for an equilibrium, regardless of the degree of price volatility, is the sequence $\{K_t^o\}_{t=1}^{\infty}$ implicitly defined by the equations:

$$LMC_t = E\left[p_t(w_t, \tilde{\epsilon}_t, K_t^o)\right]. \tag{9}$$

Clearly, $K_t^o = K_t^*$ if price volatility is limited. Furthermore, the sequence $\{K_t^o\}_{t=1}^{\infty}$ will indeed be an equilibrium capacity trajectory, regardless of whether volatility is limited or not, provided the corresponding capacity levels increase (weakly) over time, i.e., $K_{t+1}^o \geq K_t^o$, so that in equilibrium the industry will always seek to expand the aggregate capacity level.[21] A sufficient condition, in turn, for the K_t^o to increase monotonically is that the expected product price satisfies the condition: $\phi_{t+1}(K) > \phi_t(K)$ for any K, where $\phi_t(K) \equiv E\left[p_t(w_t, \tilde{\epsilon}_t, K)\right]$.

3 Inferring Long-Run Marginal Cost

This section outlines a method for estimating a firm's long-run marginal product cost from financial accounting data in conjunction with select data frequently supplied by industry analysts. Our model framework has conceptualized the long-run marginal cost in each period as the sum of current operating costs and capacity related costs. Since we estimate these cost components primarily from income statements, our approach needs to be cognizant of the income statement separation between *manufacturing* (inventoriable) costs and *period* costs. The former pertain to factory-related costs, including materials, labor and manufacturing overhead which, in turn, includes depreciation charges. Inventoriable costs are reported either as part of Cost of Goods Sold (COGS) or as additions to inventory on the balance sheet. In contrast, period costs are expensed and comprise selling as well as general and administrative (SG&A) expenses, including advertising and R&D. Conceptually, we think of the cost components w_t and f_t from Section 2 as having two components each: $w_t = w_t^+ + w_t^-$ and $f_t = f_t^+ + f_t^-$, with the "+" part referring to manufacturing (inventoriable) costs and the "-" part referring to period costs.[22]

[21] Of course, this constraint is of no importance if one postulates a competitive fringe of contract manufacturers that effectively provide a rental market for manufacturing capacity, as assumed in parts of the investment literature; see, for instance, Abel and Eberly (2011).

[22] The cost inference procedure outlined in this section is not limited to manufacturing industries, but applies equally to a range of service industries in which the provision of services is capacity constrained, e.g., transportation, health care or hospitality services.

3.1 Operating Costs

Core Manufacturing Costs

We use the label *core manufacturing costs* to refer to all manufacturing (inventoriable) costs other than depreciation, that is, $w_t^+ + f_t^+$ in our notation. It will generally be difficult to obtain separate estimates for the fixed and variable components of the core manufacturing costs, though for the purposes of estimating the LMC only the sum of fixed and variable inventoriable costs matters. The same need for aggregation applies to the period costs, that is, $w_t^- + f_t^-$.

Our inferences regarding core manufacturing costs will be based on the firm-specific variables shown in the following table.[23]

Variable	**Description**	**Units**
$Sales_{it}$	Sales Revenue	\$
$COGS_{it}$	Cost of goods sold	\$
Inv_{it}	Finished Goods Inventory	\$
D_{it}	Depreciation charge	\$
s_{it}	Sales volume	*Output units*
RD_{it}	Research and development expense	\$
$SG\&A_{it}$	General and administrative expenses	\$

Table 1: *Variables used to infer core manufacturing costs.*

The variables in Table 1 are obtained from firms' financial statements, except for the number of units sold (s_{it}) by firm i in period t. Fortunately, that variable will frequently be reported as supplementary information by the firms themselves or by industry observers. Our key variable for gauging the core manufacturing cost is Cost of Goods Manufactured (COGM), calculated as the unit cost times the quantity of modules produced (q_{it}) in the current quarter plus current depreciation charges for the use of equipment and facilities:

$$COGM_{it} = Core\ Manufacturing\ Costs + Depreciation \equiv (w_{it}^+ + f_{it}^+) \cdot q_{it} + D_{it}. \quad (10)$$

The only variable that will generally be directly observable in (10) is the depreciation charge, D_{it}. To infer $w_{it}^+ + f_{it}^+$ in (10), we rely on several accounting identities. The chart in Figure

[23]The subscript i indicates that the variables are firm-specific.

3 shows the sequence of steps for this inference procedure. To begin with, the quantity of output produced, q_{it}, equals the number of units sold plus the difference in inventory between the current and the prior period, i.e., $n_{it} - n_{it-1}$:

$$q_{it} = n_{it} - n_{it-1} + s_{it}. \tag{11}$$

Units sold in period t come from current production or inventory left from the prior quarter. Given average costing for inventory valuation purposes, the unit manufacturing cost of firm i in period t is given by:

$$ac_{it} = \frac{Inv_{it-1} + COGM_{it}}{n_{it-1} + q_{it}}. \tag{12}$$

Here, ac_{it} is effectively the average manufacturing cost per module available for sale by firm i in quarter t, taking the arithmetic mean between the beginning balance and the current period addition in both the numerator and the denominator. The left-hand-side of (12) can be inferred immediately from Cost of Goods Sold (COGS) and units sold since:

$$COGS_{it} = s_{it} \cdot ac_{it}. \tag{13}$$

We also make use of the following expression for the balance of ending inventory:

$$Inv_{it} = ac_{it} \cdot n_{it}. \tag{14}$$

Upon initializing the sequence via n_{i0} ($n_{i0} = \frac{Inv_{i0}}{ac_{i0}}$), the identity in (14) gives us the entire sequence of production and inventory levels, q_{it} and n_{it}. This, in turn, identifies the values of $COGM_{it}$ in equation (12), since the remaining four variables are either observed directly or have been identified in previous steps. The flow chart in Figure 3 illustrates the linkage among the variables required to infer both core manufacturing costs and period costs.

Period Costs

As with our inference procedure for core manufacturing costs, we cannot identify the remaining components w^{-} and f^{-} separately. However, since period costs are primarily comprised of research and development (R&D) expenses and sales, general, and administrative (SG&A) expenses, these costs are likely to be fixed for the most part. We treat R&D costs as an unavoidable fixed cost that provide firms with an "entrance ticket" to participate in industry-wide cost reductions. Firm level R&D and SG&A expenses are taken directly

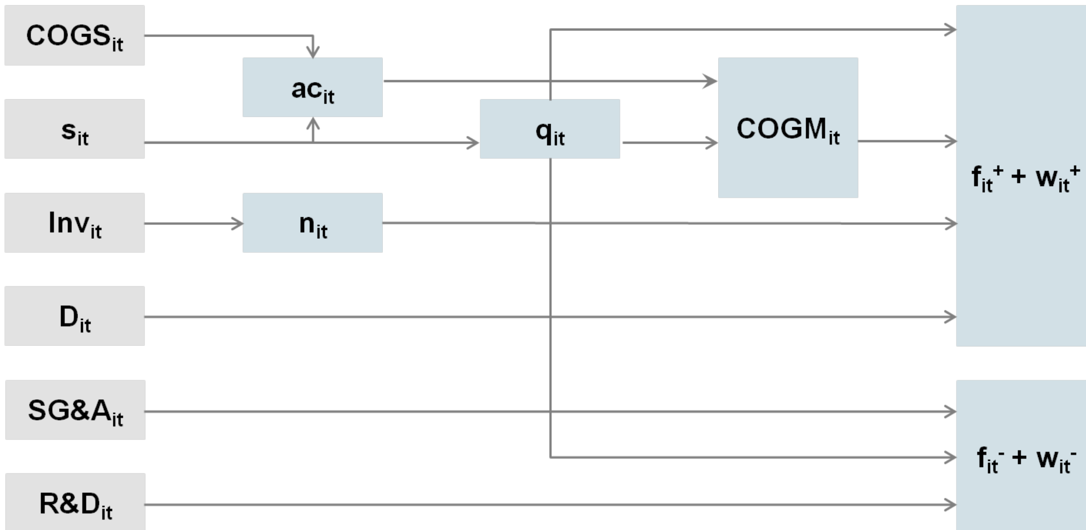

Figure 3: *Flow Chart: Inferring Operating Costs*

from the income statements. We divide these figures by the number of units of output produced by the firm in the given quarter. Thus

$$w_{it}^{-} + f_{it}^{-} = \frac{R\&D_{it} + SG\&A_{it}}{q_{it}}. \tag{15}$$

3.2 Capacity Costs

The model in Section 2 has conceptualized capacity costs as the current "user cost of capacity", c_t, multiplied by the tax factor, Δ, which effectively summarizes the income tax consequences of investments in productive capacity. For firm i in period t, we then obtain:

$$c_{it} \cdot \Delta = \eta^t \cdot \frac{v_i}{\sum_{\tau=1}^{T} (\gamma \cdot \eta_i)^\tau} \cdot \Delta. \tag{16}$$

Equation (16) reflects an implicit assumption that the firms in the industry in question employ the same discount factor, γ, face the same statutory tax rate and are subject to the same industry-wide learning parameter, η.[24]

Given an η estimate, the unit cost of new capacity acquisitions v_i can be gauged at the individual firm level. We estimate capacity acquisition costs by the relation:

[24]In our analysis of the solar PV module industry, we will rely on estimates regarding the dynamics of equipment acquisition costs as reported by the research analyst Greentech Media (GTM). The dynamics they project allow us to estimate an industry-wide learning parameter, η.

$$CAPX_{it} = (v_i \cdot \eta^t) \cdot I_{it},$$

where, in the notation of Section 2, I_{it} denotes gross capacity additions to capacity by firm i in period t. If the window of observations opens at $t = 0$ and covers the years $t = 1$ through $t = \hat{t}$, the parameter v_i can be inferred from the aggregate relation:

$$v_i = \frac{\sum_{t=1}^{\hat{t}} CAPX_{it} \cdot \eta^{-t}}{\sum_{t=1}^{\hat{t}} I_{it}}. \tag{17}$$

The η^{-t} term in the numerator of (17) "scales-up" the impact of future capital expenditures since such investments yield larger capacity additions per dollar spent. In order for equation (17) to be operational, the gross capacity additions I_{it} need to be known. If only the time-series of capacity levels K_{it} has been reported, either by the firms in supplementary information or in industry reports, the relation in (17) can be replaced by:

$$v_i = \frac{\sum_{t=1}^{\hat{t}} CAPX_{it} \cdot \eta^{-t}}{K_{i,\hat{t}+1} - K_{i1}}, \tag{18}$$

provided the amount of "old" capacity that went off-line during the observation window is negligible. This was definitely the case for the solar PV module manufacturers in our sample during the years 2008-2013, largely because almost all the existing capacity has been added after 2005.

4 Application to Solar PV Module Manufacturing

We now apply the cost inference procedure described in the previous section to solar photovoltaic module manufacturers. The major manufacturing steps include the sequential production of polysilicon, ingots, wafers, cells, and modules. There appears to be consensus in the industry that opportunities for continued cost reductions remain at each step.

Polysilicon is primarily produced via the so-called Siemens process, and the main cost reduction opportunities include improvements in energy efficiency and an increase in the scale of the Siemens chemical vapor depositor reactor. Polysilicon is used to grow ingots.

The production of larger ingots has the potential to reduce overall manufacturing costs. Additional savings can be realized through the use of quasicrystalline ingots, as these eliminate a downstream processing step in which active silicon material is discarded. Ingots are subsequently sliced into wafers. The bulk of silicon losses occur at this step, and the magnitude of losses depends on the thickness of the wire saws used.

The most capital intensive step of the module production process is the conversion from wafers to cells. During this step, the wafers are etched and doped with impurities to achieve a desired level of electrical conductivity, metallized to facilitate the transfer of charges, and treated with an anti-reflective coating (Lux Research, 2012b). Finally, cells are strung together, enclosed, and appended with a junction box to build a solar module. Since module assembly requires only one essential piece of equipment, this last step has traditionally been labor intensive. Automation continues to reduce labor requirements (Lux Research, 2012b).

4.1 Data and Industry Specific Considerations

Our sample includes ten major module manufacturers with a combined market share of approximately 35%. Since these firms are listed on U.S. stock exchanges, their financial statements have been prepared in accordance with U.S. GAAP. The firms in our sample are Yingli Green Energy, Trina Solar, Suntech Power, Canadian Solar, LDK Solar, Hanwha SolarOne, JA Solar, ReneSola, Jinko Solar, and China Sunergy. We access financial data through the Bloomberg terminal system, which compiles data readily available from the firms' annual statements. In addition, we obtain the variables s_{it}, K_{it}, and I_{it} introduced above from either quarterly or annual reports by the firms in our sample, or from press releases.

Since financial databases such as Compustat do not cover detailed production data such as production capacity for ingots, wafers, cells, and modules, we also rely on two widely used data sources provided by industry analysts: Lux Research and GreenTech Media (GTM).[25] Almost 300 other firms supply the solar PV module market (Lux Research, 2012a), though we excluded manufacturers based on four criteria: (i) lower than 0.5% share of global capacity in 2012, (ii) privately held or embedded within large conglomerates, (iii) listed on exchanges outside of the U.S, and (iv) relying on thin-film rather than crystalline silicon technology.

[25] Table 5 in Appendix A.2 provides summary details about the firms. The U.S. based firm SunPower is not in our sample because its sizable downstream solar development business makes it difficult to infer manufacturing costs from reported financial information.

Our data span 24 quarters from Q1 2008 to Q4 2013, yielding a sample of 214 cost observations.[26] Earlier studies have pointed out that Chinese solar photovoltaic manufacturers have enjoyed select price subsidies on some of their input purchases; see, for instance, Goodrich et al. (2013a). Our cost estimates would reflect any such subsidies without distorting the conclusions to the extent that all the firms in our sample were eligible for these subsidies and that the magnitude of these subsidies did not change over time.

Following industry convention, module output is measured by the overall peak power (in Watts (W) delivered) that can be generated by the module. For numerical convenience, output is usually stated either in terms of Megawatts (MW) and or Gigawatts (GW). Tables 6 and 7 in the Appendix provide descriptive statistics for the variables in our data set for the beginning and final quarters, that is, Q1-08 and Q4-13, covered in our analysis.

4.2 Estimation of Operating Costs

In order to apply the inference procedure of Section 3, we account for the sale of intermediate products, e.g., cells and wafers. Generally, these upstream products account for a small share of the firms' overall revenue.[27] Our adjustment translates firm-wide shipment levels to module-equivalent figures. To do so, we multiply shipment levels for intermediate products (i.e., wafers and cells) by the ratio of their average selling prices in a quarter to that for modules:[28]

$$s_{it}^{ME} = \zeta_{Wafer_t} \cdot s_{it}^{Wafer} + \zeta_{Cell_t} \cdot s_{it}^{Cell} + \zeta_{Module_t} \cdot s_{it}^{Module}, \tag{19}$$

where $\zeta_{Wafer_t} = \frac{ASP_{Wafer_t}}{ASP_{Module_t}}$, $\zeta_{Cell_t} = \frac{ASP_{Cell_t}}{ASP_{Module_t}}$, and $\zeta_{Module_t} = 1$. The Online Appendix lists the multipliers we use.[29] To illustrate, consider a firm that ships 200MW of modules and 100MW of wafers in a given quarter. Assuming $\zeta_{Wafer_t} = 0.38$, we would record the firm's module equivalent shipment quantity as $200 + 100 \cdot 0.38$ or 238MW. Upon doing so, we can use firm-wide COGS and inventory figures to estimate $w_{it}^{+} + f_{it}^{+}$.

[26]We do not have 240 observations given the timing of listing and delisting events. For LDK Solar, Suntech Power, and Jinko Solar, our observation windows include Q2 2009 to Q4 2012, Q1 2008 to Q1 2012, and Q3 2010 to Q4 2013, inclusive and respectively.

[27]Between 2008 and 2012, 96.5% to 98.7% of Yingli's revenues were from sales of modules.

[28]Units of time refer to quarters in this subsection; in contrast, a time period refers to a year in the next subsection dealing with capacity costs.

[29]See http://stanford.io/1ov1kdQ.

Where breakdowns of inventory into finished goods, work-in-progress, and raw materials are unavailable on a quarterly basis, we assume that the split of inventory into finished and work-in-progress goods during the first, second, and third quarters is similar to the fourth quarter splits reported by the firm for the current and previous year. We impute values for finished goods and work-in-progress for Q1, Q2, and Q3 by calculating a weighted average of the fourth quarter data points. The Q1, Q2, and Q3 estimates weigh the previous year's annual data by 75%, 50%, and 25%, respectively. Finally, to recover each firm's sequence of production and inventory levels, we index one quarter in each firm's data series to $t = 0$. The initial period is Q4-07 for most firms.[30]

4.3 Estimation of Capacity Costs

The estimate of the unit capacity costs, c_t, is anchored to the capacity acquisition cost, v, required to put in place the manufacturing capacity for one unit of output over the next T years. This expenditure is then "levelized" in accordance with equation (2) in Section 2 to arrive at the cost of one unit of capacity made available for unit of output. This cost takes into account the technological progress parameter η, which causes the unit cost of capacity acquisitions to decrease geometrically over time. Since we expect the rate of technological progress to differ between manufacturing equipment and facilities, we split capacity-related costs into two pools: manufacturing equipment (eq) and facilities (fc). In accordance with equation (2), we have:

$$c_t = c_{t,eq} + c_{t,fc} = \eta_{eq}^t \cdot \frac{v_{eq}}{\sum_{\tau=1}^{T_{eq}} (\gamma \cdot \eta_{eq})^\tau} + \eta_{fc}^t \cdot \frac{v_{fc}}{\sum_{\tau=1}^{T_{fc}} (\gamma \cdot \eta_{fc})^\tau}. \tag{20}$$

In evaluating (20), we assume $T_{eq} = 10$ for equipment and $T_{fc} = 30$ for facilities.

Ideally, we would use firm-level data on fixed assets, depreciation, capital expenditures, and total capacity available to construct a quarterly panel of capacity acquisition costs. However, the data available entail several complications. First, it is unclear whether investment expenditures were directed at capacity upgrades or capacity additions. Second, the proportion of expenditure applied to investments in facilities as opposed to equipment is ambiguous. For these reasons, we rely on data from GTM (2012) to estimate both the

[30]The exceptions are LDK Solar and Jinko Solar, for which the initial periods are Q1-09 and Q2-10, respectively.

technical progress parameter η_{eq} and a baseline, industry-wide capacity acquisition cost estimate, v_{eq}. In the Appendix (Section A.5), we perform the calculation of firm-level capacity acquisition costs consistent with equations (17) and (18) in Section 3 and find that, despite some variation in these firm-level cost estimates, the average values generate a reasonably good match with the industry-wide values reported by GTM. [31]

Component	**2009**	**2010**	**2011**	**2012**	**2013**	**2014**	**2015**	**2016**
Facility	\$0.07	\$0.07	\$0.06	\$0.06	\$0.06	\$0.06	\$0.06	\$0.06
Equipment								
Ingot	\$0.21	\$0.17	\$0.11	\$0.09	\$0.06	\$0.05	\$0.04	\$0.04
Wafer	\$0.24	\$0.20	\$0.16	\$0.13	\$0.08	\$0.07	\$0.06	\$0.06
Cell	\$0.45	\$0.30	\$0.20	\$0.16	\$0.10	\$0.08	\$0.08	\$0.07
Module	\$0.12	\$0.09	\$0.07	\$0.05	\$0.03	\$0.03	\$0.02	\$0.02
Total	\$1.02	\$0.76	\$0.54	\$0.43	\$0.27	\$0.23	\$0.20	\$0.19

Table 2: *Facility related capacity acquisition cost estimates based on Powell et al. (2013). Equipment related capacity acquisition cost estimates based on GTM (2012).*

Table 2 shows equipment related acquisition cost estimates, v_{fc}, as reported by GTM (2012). Furthermore, GTM (2012) provides equipment related costs broken down along the major manufacturing steps. For the much smaller manufacturing facility acquisition costs, v_{eq}, we rely on Powell et al. (2013).

To estimate the technological progress parameter, η_{eq}, we run the simple regression:

$$v_t = \eta_{eq}^t \cdot v_0 + \xi_t,$$

where ξ_t is assumed to be a log-normally distributed error term with $E\left[ln(\xi_t) \mid t\right] = 0$. Then, $ln(\frac{v_t}{v_0}) = t \cdot ln(\eta_{eq}) + ln(\xi_t)$. Setting 2009 as year 0, the regression estimate yields $\eta_{eq} = 0.76$ with a standard error of 0.01.

Following GTM (2012), we use the reported capacity acquisition cost for 2012, that is v_{eq} =\$0.43, as our estimate for the fourth quarter and backcast (or forecast) capacity acquisition costs for the other quarters using our estimate of η_{eq}. Using the quarter-specific capacity acquisition costs, we finally obtain quarter-specific equipment-related capacity costs,

[31]GTM generates its cost estimates by consulting industry sources on both the supply- and demand side. In particular, GTM solicits input from major equipment manufacturers such as Centrotherm and Schmid and module manufacturers such as GCL-Poly, Renesola, Suntech Power, China Sunergy, and Canadian Solar.

i.e., the term $c_{t,eq}$ in (20). To illustrate, for the Q4 2013, we obtain an equipment capacity cost of c_{eq} = \$0.16/W.

In contrast to equipment related costs, we set the technological progress parameter for facilities, η_{fc}, equal to one. The facility costs in Table 2 are based on an estimate by Powell et al. (2013) that a plant with an annual capacity of 395MW entails costs of approximately \$53M. After accounting for a 50% discount on capital investment in China (Goodrich et al., 2013a), we derive a v_{fc} =\$0.066/W. The numbers in Table 2 reflect an adjustment in the facility acquisition cost to reflect efficiency increases over time, by adjusting for the power output (in Watts) per unit of space occupied by facilities. As efficiency increases, the same physical area can generate a larger amount of output and correspondingly the facility capacity acquisition cost per Watt decreases over time.[32] These adjustments explain the slight decrease in v_{fc} over time despite setting η_{fc} =1. Upon levelizing the v_{fc} costs in Table 2, we derive a Q4 2013 facility capacity cost of c_{ef} =\$0.01/W.

Regarding the tax factor Δ in the estimation of LMCs, we employ separate tax factors, Δ_{eq} and Δ_{fc}, since equipment and facility assets have different useful lives. To calculate these figures in the context of the Chinese module manufacturers in our sample, we follow Goodrich et al. (2013b) and apply a tax rate of $\alpha = 15\%$ and a (weighted average) cost of capital of $r = .13$. Under Chinese tax rules, the useful life of equipment is 10 years and that of facilities is 20 years (PWC, 2012). Finally, these assets are depreciated on a straight-line basis for tax purposes. Taken together, the resulting values are $\Delta_{eq} = 1.08$ and $\Delta_{fc} = 1.11$.

We summarize our estimates regarding capacity costs in Finding 2:

Finding 2 *We estimate the technological progress parameter for equipment capacity costs to be η_{eq} = 0.76, implying a 24% annual reduction in equipment capacity costs. Our estimated 2013 facility and equipment capacity costs are \$0.01/W and \$0.16/W, respectively.*

4.4 Long-Run Marginal Cost Estimates for PV Modules

The preceding estimates can now be aggregated to calculate the long-run marginal cost, LMC, for each firm and quarter in our sample. These values, in turn, yield an industry-wide figure, LMC_t, by calculating the weighted average of the firm-specific LMCs in that period:

[32] Appendix A.2 details this adjustment.

$$LMC_t = \sum_i \sigma_{it} \cdot LMC_{it}. \tag{21}$$

The weights σ_{it} in (21) are derived on the basis of firm i's share of module-equivalent shipments in quarter t.

Figure 4 compares our LMC estimates with observed and in-sample ASPs.[33] We note that up to the first quarter of 2011 the ASPs and LMCs generally stayed within a narrow band of each other. The only exception appears to be the aforementioned polysilicon shortage in late 2008 and early 2009. However, starting in early 2011, the ASP curve starts to diverge from the LMC curve.

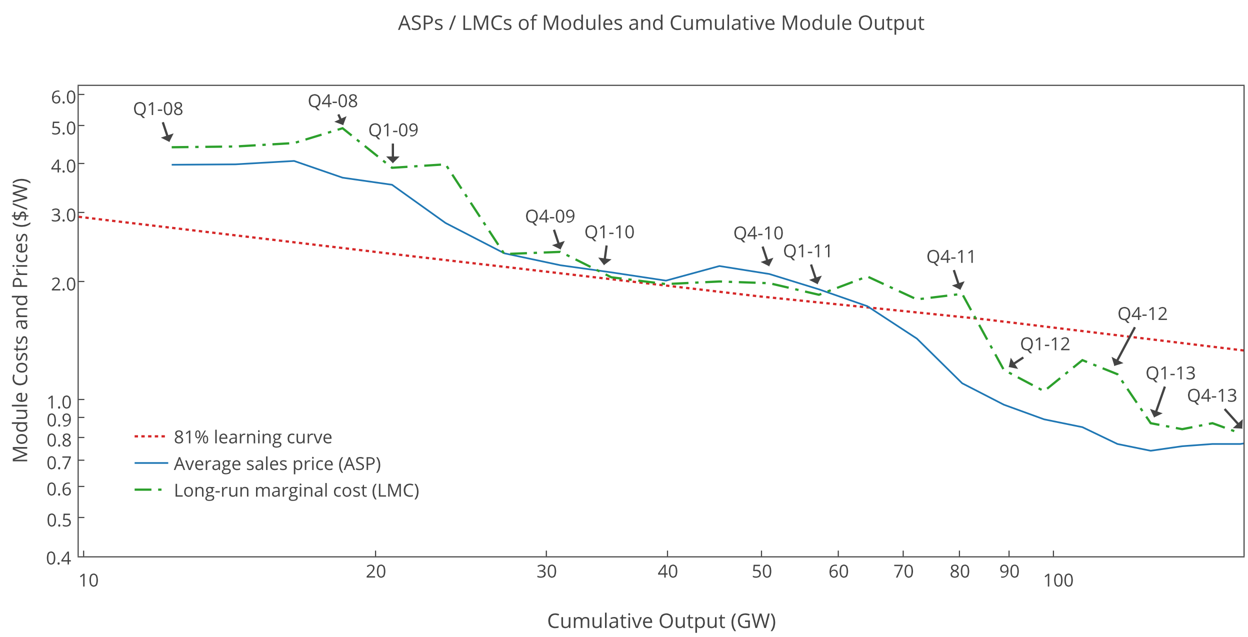

Figure 4: *LMCs and ASPs between Q1-08 and Q4-13. All prices are in 2013 U.S. dollars.*

Statistical inference allows us to make a formal claim regarding the pattern shown in Figure 4. To the extent that firm-specific LMCs and ASPs can be interpreted as draws from a distribution around the "true" market-wide LMC and ASP, the weighted standard deviation serves as a measure of the standard error around our LMC and ASP estimates and permits a statistical test of their equality.[34] For a given quarter, our procedure tests the null

[33]See Appendix A.2 for details on the ASP series.

[34]We only use firm-specific ASPs, which are defined as the ratio of firm-specific revenues to module-equivalent shipments (see Appendix A.2).

hypothesis that $ASP_t = LMC_t$. We follow Afshartous and Preston (2010) in calculating confidence intervals around the LMC and ASP measures, and perform the equivalent of a t-test of the null hypothesis, with the alternative hypothesis being that $ASP_t \neq LMC_t$.[35] We summarize our inference regarding LMCs with the following result.

Finding 3 *Our estimated LMCs are statistically significantly different from the observed ASPs in Q3-11, Q1-12, Q2-12, Q3-12, Q1-13, Q2-13, and Q3-13.*[36]

The Online Appendix details the weighted means, standard errors, and degrees of freedom used in our statistical tests.[37] The statistical results also suggest that, despite a tight polysilicon market in 2008, cost and price data from that year are consistent with a module market in equilibrium.

In interpreting our findings, we note that in any given quarter an average ASP significantly above the LMC could, of course, simply be the consequence of an unfavorable shock to demand. Yet, the recurrence of this finding in six of the eight quarters in the years 2012-2013 point to the alternative explanation that the industry had accumulated too much capacity during that period of time. The difference between the green (LMC) and the blue (ASP) curves in Figure 4 provides a measure of the price effect that can be attributed to overcapacity during our sample period. While our findings support the explanation of overcapacity, we also note a significant fall in LMCs from \$1.75/W in Q1-11 to \$0.82/W in Q4-13. This drop strongly suggests that significant intrinsic cost reductions explain a significant part of the price decreases observed during those six years.

To conclude this section, we relate our LMC measure to the so-called "Minimum Sustainable Price" (MSP) estimates in Powell et al. (2013) and Goodrich et al. (2013a,b). The

[35]Per standard econometric texts, the non-inclusion of zero in a confidence interval for a given significance level of a random variable is equivalent to a formal parametric test at the same level (Greene, 2003). Thus, the test suggested here can be implemented by comparing confidence intervals around our mean LMC and ASP measures.

[36]At the 5% significance level, we reject the equality of the ASP and LMC based on the following values: Q3-11 ($p = 0.041$), Q1-12 ($p < 0.010$), Q2-12 ($p < 0.010$), Q3-12 ($p = 0.017$), Q1-13 ($p < 0.010$), Q2-13 ($p < 0.010$), and Q3-13 ($p < 0.010$). At a less stringent 10% significance level, we also reject the null hypothesis of equality between the ASP and LMC for Q4-13.

[37]Inference using the simple arithmetic average of ASPs and LMCs broadly agrees with our results in Finding 3. At the 5% level, tests based on the simple arithmetic average of the LMC_{it} reject the equality of the ASP and LMC in Q4-09, Q4-10, Q2-11, Q3-11, Q1-12, Q2-12, Q3-12, Q4-12, Q1-13, Q2-13, Q3-13, and Q4-13. At the 10% level, the test additionally rejects the null hypothesis in Q4-11.

Minimum Sustainable Price also seeks to identify a cost-based sales price that provides an adequate return to investors. In contrast to our top-down approach based on firm-level financial data, Goodrich et al. (2013a,b) rely on a bottom-up cost model in which individual cost components are assessed in 2012 on the basis of various information sources available from industry observers. The MSP is then calculated as the derived manufacturing cost plus a profit mark-up. These approaches essentially complement our top-down approach to the derivation of LMCs which incorporate anticipated future cost reductions due to technological progress.

5 Long-Run Marginal Cost Forecasts for PV Modules

This section develops a prediction model in the form of a trajectory of future LMCs assuming that learning effects observed in the past will persist in the future. Given our estimate of η_{eq} from Section 3, the projection of capacity costs is straightforward: for any time period, capacity costs are determined by the time elapsed since the period in which the baseline cost of capacity was calculated. We project future core manufacturing costs by estimating a constant elasticity learning curve and then combine these estimates to project LMCs through 2020.

5.1 Estimation of Core Manufacturing Costs

In the following regression analysis, the dependent variable is the aggregate core manufacturing cost, $w_{it}^{+} + f_{it}^{+}$. We assume that this cost measure adheres to a constant elasticity learning curve of the form:

$$w_{it}^{+} + f_{it}^{+} = (w_{i1}^{+} + f_{i1}^{+}) \cdot (\frac{Q_t}{Q_1})^{-b_c} \cdot e^{\mu_{it}} \tag{22}$$

In (22), Q_t is the industry-wide cumulative production level in period t, b_c is the learning elasticity, and μ_{it} is an idiosyncratic error term, with $E\left[\mu_{it} \mid Q_t\right] = 0 \, \forall \, i, t$. The corresponding learning curve parameter, L, is given by $L = 2^{b_c}$. Accordingly, every doubling of cumulative output results in core manufacturing costs of only L% of what these costs were before. Given a projected cumulative industry production level at time t and an original manufacturing cost, $w_{i1}^{+} + f_{i1}^{+}$, we can use L to forecast the core manufacturing cost at time t.

We note that the assumed dynamics in (22) follows the specification of earlier studies on learning-by-doing which have centered on cumulative production as the main driver of cost reductions. This stands in partial contrast to the specification that the decline in capacity costs, in particular the acquisition cost for machinery and equipment, is assumed to be a function of calendar time. The latter choice is definitely convenient from an analytical perspective insofar as it allows for a relatively simple characterization of the marginal cost of one unit of capacity in a dynamic environment. At the same time, we note that the learning processes underlying these two cost categories may be qualitatively different: reductions in core manufacturing reflect efficiency improvements in the four main steps of the photovoltaic module manufacturing process, as summarized at the beginning of this section, while reductions in capacity costs reflect changes in the upstream operations of the equipment suppliers.

Equation (22) presents our base specification. Earlier empirical studies have found that the effect of static scale economies is small in comparison to cumulative learning effects (Lieberman, 1984). Nonetheless, we also control explicitly for changes in the scale of manufacturing facilities. Specification 2 in Table 3 below introduces scale effects by allowing for the possibility that manufacturing costs change exponentially with the scale of plants.[38] In particular, we introduce a term, $\Delta Scale_{it}$, equal to the difference between scale at time t and scale in Q1-08. $Scale_{it}$ is measured in MW/year and is defined as the average capacity per manufacturing site operated by firm i. The regression coefficient on $\Delta Scale_{it}$ is denoted by b_s. Thus:

$$w_{it}^{+} + f_{it}^{+} = (w_{i1}^{+} + f_{i1}^{+}) \cdot (\frac{Q_t}{Q_1})^{-b_c} \cdot e^{b_s \Delta Scale_{it}} \cdot e^{\mu_{it}}. \tag{23}$$

The recent work of Pillai (2015) points to the significance of polysilicon prices for decreases in the unit value of COGS.[39] Our Specifications 3 and 4 below seek to exploit variations in the slope of polysilicon prices over time. Though Specification 3 is structurally the same as Specification 2, it uses data only from the time periods over which polysilicon prices

[38] We have tested an alternative form of Specification 2 with scale effects structurally similar to the learning effects. This alternative form follows Lieberman (1984) and Stobaugh and Townsend (1975); it yields a poorer fit to the observed LMC data than the one reported in Table 4 below.

[39] Pillai (2015) has identified several explanatory variables for the decrease in COGS among module manufacturers from 2005 to 2012: (i) a reduction in input polysilicon prices, (ii) a shift in production to China, (iii) technological innovations, including a lower polysilicon utilization rate and higher module efficiencies, and (iv) lower capacity costs associated with larger capacity orders.

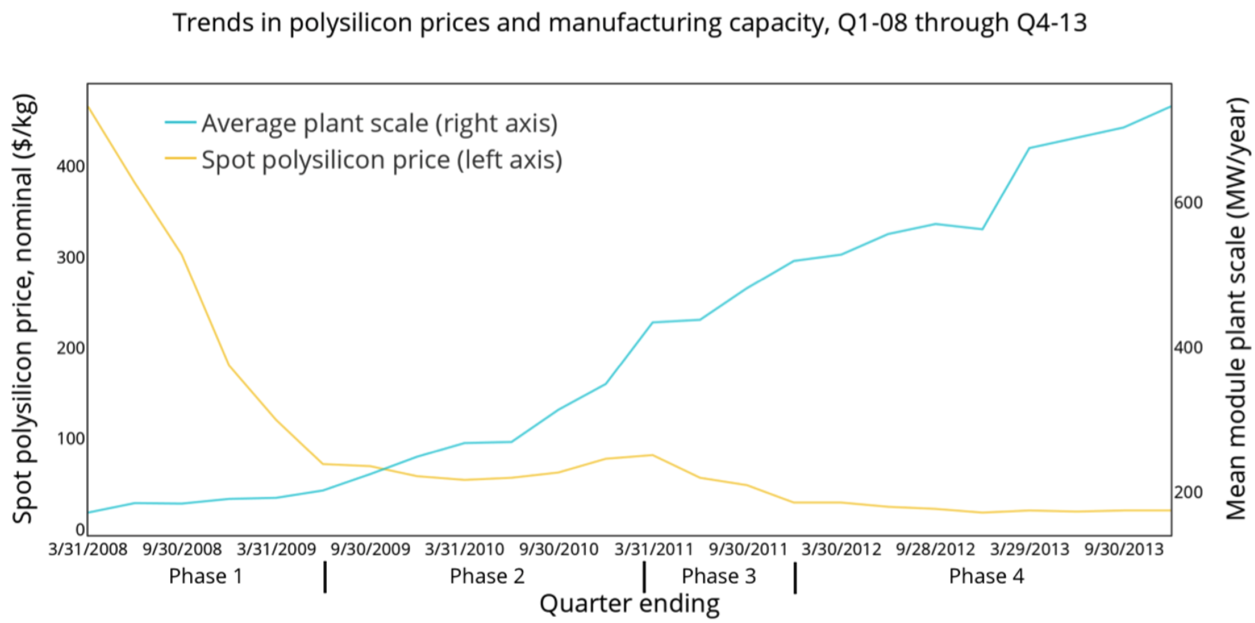

Figure 5: *Between 2008 and 2013, the average size of module manufacturing facilities increased and the price of polysilicon decreased.*

remained relatively constant; these are labeled as (pricing) phases 2 and 4 in Figure 5.[40] Specification 4 adds a phase 4 dummy term to Specification 3 and yields a learning curve parameter estimate that can be interpreted as an upper bound. By including this dummy term, we effectively 'remove' all cost reductions between the end of Phase 2 and the start of Phase 4 from the estimate of the learning elasticity. If observed cost reductions were due only to polysilicon price decreases, we would expect the coefficient on cumulative output to be statistically indistinguishable from 0 in Specification 4.

Stating the equations in logarithmic form, Specification 2, for instance, becomes:

$$ln(w_{it}^{+} + f_{it}^{+}) = ln(w_{i1}^{+} + f_{i1}^{+}) - b_c \cdot ln(\frac{Q_t}{Q_1}) + b_s \cdot \Delta Scale_{it} + \mu_{it} \tag{24}$$

Across all specifications, we set Q1-08 equal to $t = 1$. Our cumulative output measure is based on the quarterly industry-wide production estimates by Lux Research (2014). Regarding changes in scale, we use firm-level plant capacity data from Lux Research (2014). Across all four specifications, we assume the idiosyncratic error term has mean zero and is uncorrelated with the explanatory variables.

[40]Pricing phases 1, 2, 3, and 4 include quarters Q1-08 through Q2-09, Q3-09 through Q4-10, Q1-11 through Q4-11, and Q1-12 through Q4-13, respectively.

Table 3 presents our regression results.[41] The intercept should be interpreted as the average of the natural logarithm of the Q1-08 core manufacturing cost across firms. Across all specifications, the coefficient on cumulative output is significant, while that on scale is not.[42] Taken together, Specifications 1 through 4 point to significant and sustained decreases in core manufacturing costs. Furthermore, these decreases cannot be attributed exclusively to decreases in polysilicon prices.

	Dependent variable: Log Core Manufacturing Cost/Watt			
Specification	**1**	**2**	**3**	**4**
Intercept	1.216***	1.220***	0.974***	0.699***
	(0.102)	(0.100)	(0.112)	(0.093)
Cumulative Production (b_c)	-0.773***	-0.785***	-0.699***	-0.445***
	(0.052)	(0.058)	(0.064)	(0.087)
Firm Scale (b_s)	–	0.000	0.000	0.000
		(0.000)	(0.000)	(0.000)
Dummy, PS phase 4	–	–	–	-0.355*
				(0.129)
Learning Curve Parameter ($L = 2^{b_c}$)	58.5%	58.0%	61.6%	73.5%
Adjusted R^2	0.8012	0.8013	0.8128	0.8303
N	213	213	125	125
Firm fixed effects?	Yes	Yes	Yes	Yes

Table 3: *Estimated coefficients for a constant elasticity learning curve. Entries in parentheses are Driscoll-Kraay standard errors.*
Key to statistical significance: ***: ≤ 0.001; **: ≤ 0.01; *: ≤ 0.05.

Specification 3 is our preferred variant because it accounts for large swings in polysilicon prices during Phase 1 and controls for scale. Though Specification 4 does the same, we believe it is too conservative in excluding all cost reductions that occurred while polysilicon prices declined in Phase 3, especially since over 90% of the demand for polysilicon is from the solar market. By excluding these reductions, we would also risk removing other changes, such as those in module efficiency and polysilicon utilization, which Pillai (2015) documents as significant drivers of reductions in COGS. Nonetheless, if one believes either that the

[41] Our results reflect the exclusion of the observation of $w_{it}^{+} + f_{it}^{+}$ for SOL in Q4 2008, given the outlier value of \$57.92/W. Including this observation, we estimate a learning curve parameter of 57.4% and 57.8% in Specifications 1 and 2, respectively.

[42] The magnitude of estimated coefficients and standard errors on firm scale is indeed smaller than 0.001.

polysilicon and solar module markets are insufficiently linked or that polysilicon price dynamics are likely to fundamentally differ from those observed from Q2-09 through Q4-13, Specification 4 provides a more appropriate rate of cost declines for core manufacturing costs.

Finding 4 *Controlling for plant scale and excluding periods with large polysilicon price declines, we estimate a 62% learning curve for core manufacturing costs over the period 2008–2013.*

In interpreting our estimates, we note that measurement errors introduced by our cost inference procedure are assumed to be normally distributed with mean zero and uncorrelated with our explanatory variables. Since we use data from firms listed on U.S. exchanges, our estimates are potentially subject to sample selection bias. Our estimates of learning effects should be interpreted as conditional on public listing on a U.S. exchange.

With regard to the adjustments of standard errors and our inference procedure, we note that in order to account for lack of homoskedasticity, auto-correlation within firms and cross-sectional dependence across them, we report standard errors suggested by Driscoll and Kraay (1998) and implemented by Hoechle (2007).[43] Given the small size of our dataset, we correct the standard errors by scaling the asymptotic estimates by $\sqrt{\frac{N}{N-1} \cdot \frac{T-1}{T-k}}$, where N, T, and k are the number of firms, time periods, and coefficients, respectively. Moreover, our inference is based on a t-distribution with (N-1) degrees of freedom to account for our small sample size.

Appendix A.6 presents robustness checks on our inference procedure. One of these checks presents an estimation in which we explicitly adjust for increases in solar cell efficiency. Accordingly, output is measured in $\$/m^2$, rather than on a $\$/W$ basis.

5.2 Long-Run Marginal Cost Forecasts

We are now in a position to project the future trajectory of long-run marginal costs. Our cost forecasts include the three components: (i) capacity costs, (ii) period costs and (iii) core manufacturing costs. For projected capacity costs, we extrapolate the estimated Q4 2013 capacity cost, using our η_{eq}, estimate to the years 2014 through 2020. Finally, we add R&D and SG&A costs on a unit basis that are equal to the 2013 shipment-weighted average

[43] Although the calculation of these standard errors relies on large sample asymptotics, the Driscoll-Kraay errors have better small-sample properties than common alternatives, such as cluster robust variance estimators (CRVE), when cross-sectional dependence exists (Hoechle, 2007).

of firms' median R&D and SG&A costs from Q1-08 to Q4-13. For core manufacturing costs, we use the estimated intercept and coefficient on cumulative production, i.e., b_c, from Specification 3 above to project these manufacturing costs for the years 2014 through 2020. A plausible alternative would be to project core manufacturing costs from the most recent observation in Q4-13. However, this alternative approach would put exclusive weight on the LMC observation from Q4-13 by shifting the forecast LMC line to overlap with this one observation rather than the predictions based on the full LMC history from Q3-09 to Q4-13 that was used to estimate Specification 3.

The model framework in Section 2 suggests that ASPs should converge to the long-run marginal cost over time. However, given our findings in Section 4 indicating that the industry over-invested in capacity in 2011 and 2012, market demand needs to "catch up" to the aggregate manufacturing capacity in place in order for the LMCs to become the market clearing prices. For 2014 and 2015, there were 45 GW and 56 GW of new PV module production, respectively (Bloomberg Terminal). To capture the sensitivity of LMC projections to variations in demand, we present LMC forecasts contingent on annual industry output of either 50 GW/year or 60 GW/year in future years.

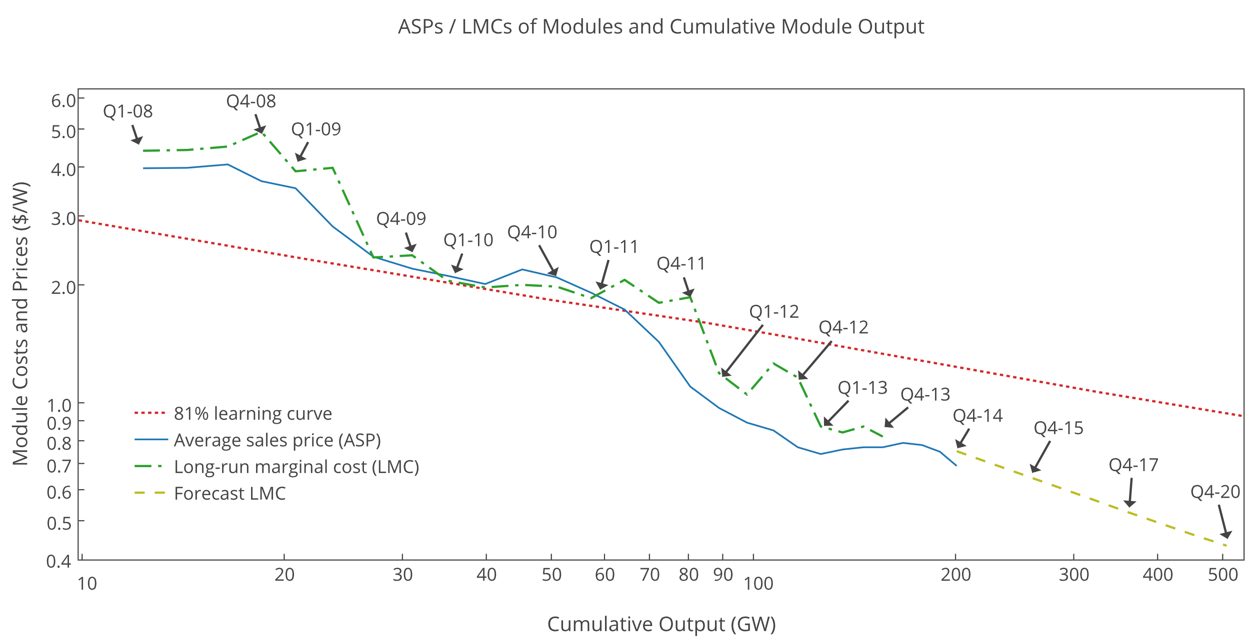

Figure 6: *Projected LMCs through 2020, assuming a constant yearly addition of 50GW between 2016 and 2020. All prices are in 2013 U.S. dollars.*

Figure 6 depicts the forecast trajectory of LMCs through 2020. This curve reflects both the 62% learning curve for core manufacturing costs and the 76% annual geometric decline in capacity costs. This trajectory represents our benchmark of the industry's production cost fundamentals and can be interpreted as a trend-line to which the ASPs are expected to converge over time as market demand catches up with installed capacity. Table 4 presents sensitivity analysis for our estimates by means of 95% prediction intervals for the years 2017 and 2020 under the two alternative industry output scenarios.[44]

	2017 LMC			**2020 LMC**		
Demand	PE^{-}	**PE**	PE^{+}	PE^{-}	**PE**	PE^{+}
50 GW	\$0.49	\$0.53	\$0.58	\$0.39	\$0.43	\$0.48
60 GW	\$0.48	\$0.52	\$0.57	\$0.38	\$0.42	\$0.47

Table 4: *LMC projections under different assumptions about annual demand for solar PV modules. All figures are in 2013 dollars.*

The results in Figure 6 and Table 4 speak directly to the so-called SunShot Initiative, articulated by the U.S. Department of Energy in 2011. According to this policy goal, the U.S. government envisions a market price of \$0.50/W for PV modules by 2020. Our econometric estimates indicate, that even if the industry were to continue to produce at a rate of 50 GW in the intervening years, and not experience any further output growth, the DOE target is likely to be met provided market prices are given by the long-run marginal cost. In fact, our confidence intervals suggest that the DOE price goal for modules is likely to be met ahead of schedule in either 2018-2019, at least in terms of 2013 dollars.

6 Conclusion

This paper has presented a model framework and an empirical inference procedure for the long-run marginal cost in an industry characterized by declining production costs. We have focused our analysis on solar photovoltaic modules, an industry in which a large number of firms supply a fairly homogeneous product. Our model framework predicts that in a dynamic competitive equilibrium suppliers choose their aggregate capacity investments so that the resulting market prices will in expectation be equal to the declining long-run marginal cost.

[44] In the table, "PE" indicates "point estimate", "PE^{+}" denotes the upper bound of the prediction interval and "PE^{-}" denotes the lower bound of the prediction interval.

Since the corresponding trajectory of market prices would allow firms to recover their periodic operating- and capacity related costs in the long-run, these prices are minimally sustainable from an economic perspective.

Our cost inference procedure is based on firm-level financial accounting data, in particular COGS, SG&A expenses and inventory balances. In addition, we rely on select data from industry analysts regarding manufacturing capacity and output shipments by individual firms in our sample. Applying our cost inference method to data from solar PV module manufacturers enables us to estimate long-run marginal costs on a quarterly basis and to contrast them with observed ASPs. While our findings suggest that the ASPs and LMCs are statistically indistinguishable early on in our sample period, they are significantly different in most quarters of the years 2012 and 2013. Our findings support the argument that during those years the observed dramatic price reductions reflect at least in part excess capacity in the industry rather than cost reductions only. Furthermore, the resulting difference between average LMCs and ASPs provide a measure of the price effect associated with excess industry capacity.

Our cost inferences also generate panel data that allow us to extrapolate how the long-run marginal cost of PV modules will change as a function of time and experience. Controlling for plant scale and significant drops in polysilicon prices, our findings lead to a 62% learning curve for core manufacturing costs. Combined with our estimates for the annual capacity cost declines, we arrive at an overall learning curve, that appears much steeper than the traditional 80% learning curve, provided the industry continues to add at least 50 GW of output annually.

The methods and findings of this paper have several immediate policy implications that could be explored in future research. The pricing of solar PV modules, in particular by Chinese suppliers, has been challenged on legal grounds in recent years. The complaint of “dumping” modules is akin to that of predatory pricing in domestic pricing disputes and generally refers to pricing below cost. Our results indicate that the prices in the 2011-2013 window were frequently below the LMC at the time. Yet, as argued in Section 2 above, the long-run marginal cost includes several components that are likely to be considered “sunk” in the short-run. If the relevant cost benchmark in legal disputes is the short-run average variable cost, our measure of core manufacturing costs would provide an upper bound. Our findings give no indication that at any point in time the firms in our sample were charging prices below their core manufacturing costs.

The demand for solar PV modules over the past decade has in significant part been driven by public policy support in the form of feed-in tariffs, investment tax credits and renewable energy portfolio standards. As solar system prices have dropped rapidly in recent years, governments in many countries have scaled back these subsidies. For instance, the federal investment tax credit in the U.S. is scheduled to decrease gradually from the current 30% to 10% by 2024. One of the key questions in the ongoing debate is the magnitude of tax credits or, alternatively feed-in tariffs in some countries, to sustain the recent pace of new solar PV deployments. The debate about these policy support question hinges in large part on whether the recent acceleration in cost reductions associated with the production of solar modules is likely to persist.

Another promising direction for future research is to examine the link between R&D spending and subsequent reductions in core manufacturing costs at the individual firm level. Our framework has essentially assumed that by spending a sufficient amount on R&D, firms gain access to the cost reduction opportunities that are available for the industry. Over longer time horizons, it is, of course, particularly plausible that the speed of cost reductions for an individual firm is at least partially linked to that firm's earlier R&D spending.

The research approach taken in this paper is principally applicable in industries other than solar photovoltaic modules. Our analysis has taken advantage of several features that apply to the market for PV modules, including a large number of price taking firms, a fairly homogeneous product and a significant number of "pure play" manufacturers for whom the production of modules is the dominant line of business. In industries where the latter condition is not met, our cost inference procedure would be diluted if one were to rely on firm-wide financial reports. Instead one would need to either aggregate the cost inferences across multiple products or rely on segment reports, e.g, income statement and operating assets for a particular division within the firm. In industries where firms have significant pricing power, the analysis in this paper could be enriched by inferring not only long-run marginal costs but also the corresponding price mark-ups.

A Appendix

A.1 Proof of Finding 1

We verify that the sequence of K_t^* given by:

$$P_t^o(K_t^*) = LMC_t \equiv c_t \cdot \Delta + w_t + f_t,$$

is indeed implementable by a sequence of non-negative investments I_t^* if

$$K_t = I_{t-T} + I_{t-T+1} + \ldots + I_{t-1},$$

and $K_0 \leq K_1^*$. The non-negativity constraints are met if $K_{t+1}^* \geq K_t^*$ for $t \geq 1$. This follows from the observation:

$$P_{t+1}^o(K_{t+1}^*) = LMC_{t+1} < LMC_t = P_t^o(K_t^*),$$

combined with the NEC condition requiring that $P_{t+1}^o(K) > P_t^o(K)$ for all K.

It remains to verify that, given the aggregate capacity levels $\{K_t^*\}_{t=1}^{\infty}$, firms will break-even on their investments. Without loss of generality, assume that a particular firm invests in one unit of capacity at time t. The prevailing equilibrium market price in the next T periods is given by $P_{t+\tau}^o(K_{t+\tau}^*) = LMC_{t+\tau}$, with $\tau \in [1, T]$. The firm utilizes this capacity over the next T periods and the pre-tax cash flows of the investment are given by:

$$CFL_t = -v \cdot \eta^t,$$

and for $1 \leq \tau \leq T$,

$$CFL_{t+\tau} = LMC_{t+\tau} - w_{t+\tau} - f_{t+\tau} = c_{t+\tau} \cdot \Delta$$

since, by definition, $LMC_t = w_t + f_t + c_t \cdot \Delta$. Taxable income in period $t + \tau$ becomes:

$$I_{t+\tau} = CFL_{t+\tau} - d_\tau \cdot v \cdot \eta^t.$$

Given a corporate income tax rate of α, the overall NPV of the investment is:

$$NPV_t = \sum_{\tau=1}^{T} \left[CFL_{t+\tau} - \alpha \cdot I_{t+\tau}\right] \gamma^\tau - \eta^t \cdot v. \tag{25}$$

To see that the expression in (25) is indeed zero, we substitute the expressions for $CFL_{t+\tau}$ and $I_{t+\tau}$:

$$NPV_t = (1-\alpha) \cdot \sum_{\tau=1}^{T} \Delta \cdot c_{t+\tau} \cdot \gamma^{\tau} + \alpha \cdot \sum_{\tau=1}^{T} d_{\tau} \cdot \eta^{t} \cdot v \cdot \gamma^{\tau} - \eta^{t} \cdot v. \tag{26}$$

The second term on the right-hand side of (26) denotes the depreciation tax shield. Dividing by $1-\alpha$ and collecting terms yields:

$$\frac{1}{(1-\alpha)} \cdot NPV_t = \Delta \cdot \sum_{\tau=1}^{T} c_{t+\tau} \cdot \gamma^{\tau} - \eta^{t} \cdot v \left[\frac{1 - \alpha \cdot \sum_{\tau=1}^{T} d_{\tau} \cdot \gamma^{\tau}}{1-\alpha} \right]. \tag{27}$$

The tax-factor, Δ, was defined in the main text as:

$$\Delta = \frac{1 - \alpha \cdot \sum_{\tau=1}^{T} d_{\tau} \cdot \gamma^{\tau}}{1-\alpha}. \tag{28}$$

Therefore equation (27) reduces to:

$$\frac{1}{(1-\alpha)} \cdot NPV_t = \Delta \left[\sum_{\tau=1}^{T} c_{t+\tau} \cdot \gamma^{\tau} - \eta^{t} \cdot v \right].$$

It remains to show that $NPV_t = 0$, which follows from the construction of the unit cost of capacity because:

$$\sum_{\tau=1}^{T} c_{t+\tau} \cdot \gamma^{\tau} \equiv \sum_{\tau=1}^{T} \frac{\eta^{t+\tau} \cdot v}{\sum_{j=1}^{T} (\gamma \cdot \eta)^{j}} \cdot \gamma^{\tau} = \eta^{t} \cdot v.$$

$\blacksquare$

A.2 Data and Adjustments for Cost Inferences

1. Sample of PV Module Manufacturers

Firm	Ticker	2012 Capacity %	2012 Production %
Yingli Green Energy	NYSE: YGE	3.7	6.4
Trina Solar	NYSE: TSL	3.7	4.6
Suntech Power	NYSE: STP	4.0	5.9
Canadian Solar	NASDAQ: CSIQ	3.4	4.8
LDK Solar	NYSE: LDK	2.7	1.5
Hanwha Solar One	NASDAQ: HSOL	2.4	2.6
JA Solar	NASDAQ: JASO	3.3	3.6
ReneSola	NYSE: SOL	1.6	1.7
JinkoSolar	NYSE: JKS	1.9	2.7
China Sunergy	NASDAQ: CSUN	1.3	1.1

Table 5: *Firms included in sample, including stock tickers, capacity, and market share as of 2012 (Lux Research, 2012a).*

2. Price Data

Prominent sources of module price data (e.g., Bloomberg New Energy Finance (BNEF) and pvXchange) began collecting price data no earlier than Q3-09. The average sales price (ASP) measure we use in our graphs is therefore a composite of several indexes. The measure equally weighs our estimates of in-sample ASPs and a composite of price indexes that we obtain either from Swanson (2011) or the Bloomberg terminal system. For each firm and quarter, we derive firm-specific ASPs as the quotient of revenues and the sales volume:

$$ASP_{it}(firm - specific) = \frac{Revenue_{it}}{s_{it}^{ME}} \tag{29}$$

These figures are aggregated into a quarter-specific average to obtain the in-sample ASP:

$$ASP_t(in - sample) = \sum_i \omega_{it} \cdot ASP_{it}, \tag{30}$$

where the weights, ω_{it} in the above summation are in proportion to the firms' share of module-equivalent shipments across all firms in our sample for that quarter.

Our composite of indexes reflects the data available for a particular period. Prior to Q1-10, we use price data included in Swanson (2011). After Q4-10, we use a Bloomberg

New Energy Finance index for multi-crystalline silicon module prices. To bridge the gap between the data from Swanson and that available from BNEF, we use the pvXchange Crystalline Modules China Price available from the Bloomberg terminal. We chose this index for Q1-10 through Q3-10 among those from PVXchange and PVinsights and accessible from the Bloomberg terminal system because it offered the best match with the BNEF multi-crystalline silicon module price index over the time periods in which we could observe both indexes. The ASPs on our graphs equal the simple average of our composite index and in-sample ASP measures.

We note that our tests about whether the market was equilibrium in a given quarter use only the firm-specific ASP estimates. By using the firm-specific ASP estimates, we can compare a distribution of estimated ASPs and inferred LMCs across firms. Overall, we observe a relatively close match between our in-sample ASPs and the index price data across the 24 quarters in our sample (median difference of about 12% between the two numbers.)

A.3 Descriptive Statistics

Tables 6 and 7 summarize the key variables considered in our sample.[45] For 2008, we include eight firms in the sample since only eight of ten firms released quarterly data for Q1 2008, given the timing of their listing (except $CAPX_{it}$ and K_{it}, where the summary includes 10 firms). For 2013, we include 8 firms in the sample since only eight of ten firms released quarterly data for Q4 2013. All dollar figures are in nominal terms.

[45]Note: The minimum LMC reported in Table 6 is an outlier; the 25th percentile of LMCs in Q1-08 across firms is \$3.95/W.

Variable	**Units**	**Min**	**Median**	**Max**	**Notes**
$COGS_{it}$	\$ (M)	69.94	131.90	338.11	
s_{it}^{ME}	MW	19.15	41.79	106.69	
Inv_{it}	\$ (M)	13.82	53.02	119.42	
D_{it}	\$ (M)	1.48	2.82	10.40	
$SG\&A_{it}$	\$ (M)	3.66	9.03	31.77	
RD_{it}	\$ (M)	0.30	0.59	2.79	
ac_{it}	\$/W	1.49	3.16	3.65	
n_{it}	MW	4.78	16.32	54.84	
q_{it}	MW	18.55	42.45	10.45	
$COGM_{it}$	\$ (M)	68.94	130.61	360.44	
$f_{it}^{+} + w_{it}^{+}$	\$/W	1.16	3.13	3.64	
$f_{it}^{-} + w_{it}^{-}$	\$/W	0.04	0.27	0.33	
ASP_{it}	\$/W	1.97	4.15	4.67	
LMC_{it}	\$/W	1.80	3.98	4.50	
$CAPX_{it}$	\$ (M)	44.08	144.97	1297.90	Annual figure
K_{it}	MW/year	40	210	680	

Table 6: *Descriptive statistics for Q1-08*

Variable	**Units**	**Min**	**Median**	**Max**	**Notes**
$COGS_{it}$	\$ (M)	118.93	345.82	538.41	
s_{it}^{ME}	MW	223.36	582.83	937.18	
Inv_{it}	\$ (M)	33.49	157.41	271.21	
D_{it}	\$ (M)	5.67	22.22	58.09	
$SG\&A_{it}$	\$ (M)	12.82	40.32	154.69	
RD_{it}	\$ (M)	1.07	4.92	18.08	
ac_{it}	\$/W	0.48	0.58	0.67	
n_{it}	MW	62.89	249.77	420.22	
q_{it}	MW	198.36	592.04	773.47	
$COGM_{it}$	\$ (M)	109.15	336.42	423.42	
$f_{it}^{+} + w_{it}^{+}$	\$/W	0.45	0.52	0.60	
$f_{it}^{-} + w_{it}^{-}$	\$/W	0.06	0.08	0.22	
ASP_{it}	\$/W	0.56	0.67	0.84	
LMC_{it}	\$/W	0.72	0.80	0.89	
$CAPX_{it}$	\$ (M)	23.13	69.45	196.74	Annual figure
K_{it}	MW/year	1155	1800	2800	

Table 7: *Descriptive statistics for Q4-13*

A.4 Adjusting Facility Costs for Physical Efficiency Gains

In Section 4.3, we arrived at a v_{fc} estimate of \$0.066/W based on analysis reported by Powell et al. (2013). Our estimate reflects an assumption that the manufacturing plant produces modules with a 13.6% efficiency, where efficiency is defined as the ratio of power capacity per physical area (i.e., $efficiency = \frac{power}{m^2}$). As the efficiency of solar modules increases, the same physical area of output contains a greater Watt capacity and therefore the facility capacity acquisition cost per Watt is likely to decrease. Since the efficiency of modules changed over the course of our data window, we adjust v_{fc} to reflect inter-temporal changes in efficiency. In particular, we follow (31) to derive an efficiency-adjusted $v_{t,fc}$:

$$v_{t,fc} = v_{fc} \cdot \frac{eff_{ref}}{eff_t}. \tag{31}$$

The ratio of efficiencies in (31) reflects an adjustment for improvements in cell efficiency over time. Recalling that $\eta_{fc} = 1$, we modify $c_{t,fc}$ from its form in (20) to:

$$c_{t,fc} = \frac{v_{fc} \cdot \frac{eff_{ref}}{eff_t}}{\sum_{\tau=1}^{30} \gamma^\tau}.$$

Here, eff_t and eff_{ref} refer to average efficiency levels in the current and baseline periods, respectively. Table 8 shows solar PV efficiency data from Fraunhofer (2012). Our estimated facility cost estimates in Table 2 for 2013 through 2016 reflect assumptions about efficiency improvements. Given reports by Fraunhofer that 2014 average efficiency levels had reached 16%, we assume a 15.5% efficiency level in 2013 and a steady 16% efficiency level thereafter.

2003	**2004**	**2005**	**2006**	**2007**	**2008**	**2009**	**2010**	**2011**	**2012**
12.0%	12.5%	12.5%	12.7%	13.0%	13.1%	13.1%	13.4%	14.5%	14.7%

Table 8: *Average crystalline silicon module efficiency(Fraunhofer, 2012).*

A.5 Firm-specific Capacity Cost Estimates

While our analysis in the text relied on figures from industry analysts to derive industry-wide capacity acquisition costs, we now illustrate the derivation of firm-level capacity acquisition costs as outlined in Section 4.3. The direct application of (18) would entail two potential challenges in connection with solar PV modules. First, the expression assumes that all capital expenditures were used to expand module manufacturing capacity across all four steps of the value chain (i.e., ingots, wafers, cells, and modules), or what we will refer to as *integrated* module manufacturing capacity. However, firms could have expanded their capacity to produce only some of these components. Second, firms' financial statements do not specify the portion of their capital expenditures that were applied to facility improvements as opposed to investments in new production equipment.

We address the first issue by using an integrated module-equivalent (ME) level of capacity, K^{ME}, that "marks down" capacity additions that did not include all components of module manufacturing by the ratio of the capacity costs for the components actually installed to that for all components. In practice, firms have tended to invest in capacity either for only one of the four steps or for combinations of the four steps that are contiguous to each other. The second observation implies that firms have invested in, for example, cell, wafer, and ingot capacity but not in only cell and ingot capacity. This practical reality implies that there are ten types of what we term *contiguous capacity investment bundles.*[46] Of course, the ten capacity bundles will differ in terms of associated costs. To account for these differences, we calculate an integrated module-equivalent (ME) level of capacity, K^{ME}:

$$K^{ME} = \sum_{j=1}^{j} K_j \cdot \chi_j. \tag{32}$$

Here, the index j refers to the ten contiguous capacity investment bundles. We use quarterly firm-level capacity data from Lux Research (2014) across all steps of the value chain to derive K_j and K^{ME}.[47] We determine the expansion of a particular bundle by (1) taking the minimum of the capacity expansions for all constituent value chain steps and (2) subtracting the expansions recorded for more inclusive bundles. As an example, when calculating the capacity expansion in the "cells and modules" bundle, we know that this increase cannot exceed the observed expansion of either cell or module capacity (i.e., the constituent value chain steps). We thus calculate the minimum capacity expansion level observed across these two steps. To avoid double counting capacity expansions in cells and modules, we subtract the capacity expansion observed across the two more inclusive bundles, namely "modules, cells, wafers, and ingots" and "modules, cells, and wafers."

[46]The ten bundles are investments in (1) ingots only, (2) wafers only, (3) cells only, (4) modules only, (5) wafers and ingots, (6) cells, wafers, and ingots, (7) cells and wafers, (8) modules, cells, wafers, and ingots, (9) modules, cells, wafers, and (10) modules and cells.

[47]We make several amendments to the data based on our findings from firms' press releases and industry analysts. The Online Appendix details the adjustment.

The adjustment factor χ_j "marks down" the capacity additions for bundles that do not include all four components of the value chain. We define χ_j as the ratio of the capacity cost for bundle j to the capacity cost for the integrated module capacity investment. We estimate χ_j as the average of the ratios from 2009 to 2016 implied by Table 2.

The following example illustrates the derivation of K_j, χ_j and K^{ME}. Our data for China Sunergy (CSUN) indicate for the first quarter of 2008 a module capacity of 70 MW/year and a 2014 module capacity of 1155 MW/year, for a change of 1085 MW/year. We also observe a 2008 cell capacity of 220 MW/year and a 2014 cell capacity of 800 MW/year, for a change of 580 MW/year.[48] In this example, we seek to determine the change in $K_{Module,Cell}$, K_{Module} and K_{Cell}. The change in $K_{Module,Cell}$ is $min(1085, 580) = 580$ MW/year, that in K_{Module} is 1085 MW/year - 580 MW/year or 505 MW/year, and in K_{Cell} it is 580 MW/year - 580 MW/year or 0 MW/year. The subtractions to calculate K_{Module} and K_{Cell} adjust for the increase in capacity of the more inclusive bundle, $K_{Module,Cell}$. The capacity acquisition costs in Table 2 imply $\chi_{Module,Cell}$, χ_{Module} and χ_{Cell} coefficients of 0.50, 0.12, and 0.38, respectively. We finally calculate the change in K^{ME} for CSUN as 580·0.5 + 505·0.12 + 0·0.38 MW/year for a total of 348.13 MW/year.

The second issue in implementing (18) above concerns the split between investments in equipment as opposed to facilities. We address this issue by defining a factor β_t that measures the share of equipment costs among the total capacity costs:

$$\beta_t = \frac{v_{eq} \cdot \eta_{eq}^t}{v_{eq} \cdot \eta_{eq}^t + v_{fc}}.$$

The data in Table 2 allow us to calculate β_t on an annual basis. Since GTM does not provide capacity cost data for years preceding 2009, we backcast v_{eq} for 2008 by using our estimate $\eta_{eq} = 0.79$. These two adjustments lead to the following modification of (18):

$$v_{i,eq} = \frac{\sum_{t=1}^{6} \beta_t \cdot CAPX_{it} \cdot \eta_{eq}^{-t}}{K_{i2014}^{ME} - K_{i2008}^{ME}}. \tag{33}$$

Expression (33) allows us to derive a set of $v_{i2007,eq}$ values. Consistent with the model in Section 3, we can levelize the capacity acquisition cost to derive the equipment-related cost

[48] Without the adjustments reflected in (32), we would infer that CSUN expanded its capacity by 1085 MW/year, that $v_{CSUN2007,eq}$ was \$0.52/W instead of \$1.61/W (see Table 9) and that $c_{CSUN2013,eq}$ was \$0.05/W instead of \$0.15/W.

of capacity. We calculate firm-specific values $v_{i,fc}$ in the same manner as (33), except that (1-β_t) substitutes for β_t and $\eta_{fc} = 1$ for η_{eq}.

Table 9 presents our estimates of levelized capacity costs, c_{fc}, $c_{2007,eq}$, and $c_{2013,eq}$. The bottom row adjusts these values to incorporate the tax factors Δ_{fc} and Δ_{eq}. The penultimate row in Table 9 corresponds to weighted averages for v_{eq}, v_{fc}, c_{eq}, and c_{fc}. These weights of firm-specific measures are calculated in proportion to each firm's share of module-equivalent capacity added between 2008 and 2014, relative to the total additions in the sample.[49]

In concluding this subsection, we note that our firm-specific capacity cost estimates yield values that align reasonably well with those in Table 2 in the main text where we followed GTM (2012). In particular, our weighted average of v_{eq} =\$1.88 for 2007 translates into a value of v_{eq} =\$1.08 for 2012, upon applying the technological progress parameter η =0.76 for two years. We recall that the corresponding GTM (2012) estimate in Table 2 is v_{eq} =\$1.01 for 2012.

Firm	v_{fc}, **\$/W**	c_{fc}, **\$/W**	$v_{2007,eq}$, **\$/W**	$c_{2007,eq}$, **\$/W**	$c_{2013,eq}$, **\$/W**	**Notes**
CSUN	0.07	0.01	1.61	0.80	0.15	
YGE	0.12	0.02	2.68	1.33	0.26	
TSL	0.05	0.01	1.17	0.58	0.11	
JKS	0.05	0.01	1.15	0.57	0.14	2009 base year
CSIQ	0.06	0.01	1.52	0.75	0.14	
HSOL	0.11	0.01	2.44	1.21	0.23	
LDK	0.16	0.02	3.71	1.84	0.35	2013 end year for CAPX, module K
JASO	0.07	0.01	1.70	0.84	0.16	
SOL	0.11	0.01	2.58	1.28	0.25	
STP	0.06	0.01	1.37	0.68	0.13	2012 end year for CAPX, module K
Weighted Avg.	0.08	0.01	1.88	0.93	0.17	
With tax factor		0.01		1.01	0.19	

Table 9: *Estimated cost of capacity, facility and equipment for firms in our sample.*

[49]The weighted average statistics exclude estimates for LDK because the reported equipment capacity acquisition cost is much higher than that of its peers.

A.6 Learning Curve Estimation: Robustness Checks

Accounting for Physical Efficiency Gains

Since a time trend would have accounted for core manufacturing cost reductions due to improved quality, as measured by physical efficiency, we repeat Specifications 1 through 4 with data expressed on a dollar per square meter basis. We represent the effect of efficiency improvements in solar cells via the metric $efficiency = \frac{power}{m^2}$. Accordingly, we convert manufacturing costs from dollars per Watt to dollars per square meter by multiplying the former by efficiency. We use average module efficiency levels from Table 8 and change the scale and cumulative output measures to a square meter basis. Table 10 summarizes our estimates; each specification corresponds to the numbering in Table 3.

	Dependent variable: Log Core Manufacturing Cost/m^2			
Specification	**1**	**2**	**3**	**4**
Intercept	-0.844***	-0.841***	-1.128***	0.703***
	(0.106)	(0.104)	(0.104)	(0.098)
Cumulative Production (b_c)	-0.723***	-0.731***	-0.626***	-0.454**
	(0.054)	(0.059)	(0.063)	(0.094)
Firm Scale (b_s)	–	0.000	0.000	0.000
		(0.000)	(0.000)	(0.000)
Dummy, PS phase 4	–	–	–	-0.371*
				(0.128)
Learning Curve Parameter ($L = 2^{b_c}$)	60.6%	60.3%	64.8%	73.0%
Adjusted R^2	0.7617	0.7618	0.7648	0.8293
N	213	213	125	125
Firm fixed effects?	Yes	Yes	Yes	Yes

Table 10: *Estimated coefficients on a constant elasticity learning curve. Entries in parentheses are Driscoll-Kraay standard errors.*
Key to statistical significance: ***: ≤ 0.001; **: ≤ 0.01; *: ≤ 0.05.

Comparing Tables 3 and 10, we observe that the implied learning curve parameters are essentially the same across the two tables, with the parameters on a \$/Watt basis roughly up to $1 - 3\%$ steeper than those on a $\$/m^2$ basis. The estimates suggest that our results are robust to the specification of output on an efficiency-adjusted basis.

Exclusion of Select Firm-Quarter Observations

Since some firms had a small share of modules in their output mix over some of the periods in our panel, we conducted four robustness checks in which we exclude observations for these firms in such periods. In the first, we drop data from CSUN between Q1-08 and Q3-10. The second drops data from JASO between Q1-08 and Q4-11, while the third drops data from SOL between Q1-08 and Q3-10. Finally, we drop all three sets of observations. We do not list the estimates derived upon dropping these observations, but we do not find any material differences between the learning curve parameters estimated from the full sample and those obtained when using the restricted samples. Though the standard errors change, they do not change in a systematic direction with these exclusions.

References

Abel, A., and J. Eberly (2011), "How Q and Cash Flow Affect Investment Without Friction: An Analytic Explanation," *Review of Economic Studies*, 78, 1179–1200.

Afshartous, D., and R. A. Preston (2010), "Confidence intervals for dependent data: equating non-overlap with statistical significance," *Computational Statistics & Data Analysis*, 54(10), 2296–2305.

Arrow, K. (1964), "Optimal Capital Policy, The Cost of Capital and Myopic Decision Rules," *Annals of the Institute of Statistical Mathematics*, 1-2, 21 – 30.

Balachandran, B., R. Balakrishnan, and K. Sivaramakrishnan (1997), "On the Efficiency of Cost-Based Decision Rules for Capacity Planning," *The Accounting Review*, 72(4), 599–619.

Balakrishnan, R., and K. Sivaramakrishnan (1996), "Is Assigning Capacity Costs to Individual Products Necessary for Capacity Planning?" *Accounting Horizons*, 10(3), 1–11.

Balakrishnan, R., and K. Sivaramakrishnan (2002), "A Critical Overview of Full-Cost Data for Planning and Pricing." *Journal of Management Accounting Research*, 14, 3–31.

Baldenius, T., A. Nezlobin, and I. Vaysman (2016), "Managerial Performance Evaluation and Real Options," *The Accounting Review*, 91(3), 741–766.

Banker, R. D., and J. S. Hughes (1994), "Product Costing and Pricing." *The Accounting Review*, 69, 479–494.

Benkard, L. (2000), "Learning and Forgetting: The Dynamics of Aircraft Production," *The American Economic Review*, 90(4), 1034 – 1054.

Bouwens, J., and B. Steens (2016), "Full-Cost Transfer Pricing," *Journal of Management Accounting Research*, forthcoming.

Candelise, C., M. Winskel, and R. J. Gross (2013), "The Dynamics of Solar PV Costs and Prices as a Challenge for Technology Forecasting," *Renewable and Sustainable Energy Reviews*, 26, 96–107.

Carlton, D. W., and J. M. Perloff (2005), *Modern Industrial Organization.*, 4th ed., Pearson/Addison Wesley, New York.

Cooper, R., and R. S. Kaplan (1988), "How Cost Accounting Distorts Product Costs," *Management Accounting*, 69, 20 – 28.

Dick, A. R. (1991), "Learning by Doing and Dumping in the Semiconductor Industry," *Journal of Law and Economics*, 34(1), pp. 133–159.

Driscoll, J. C., and A. C. Kraay (1998), "Consistent Covariance Matrix Estimation with Spatially Dependent Panel Data," *Review of Economics and Statistics*, 80(4), 549–560.

Dutta, S., and S. Reichelstein (2010), "Decentralized Capacity Management and Internal Pricing," *Review of Accounting Studies*, 15(3), 442 – 478.

Fraunhofer (2012), "Photovoltaics Report," Tech. rep., Fraunhofer Institute for Solar Energy Systems.

Goodrich, A., P. Hacke, Q. Wang, B. Sopori, R. Margolis, T. L. James, and M. Woodhouse (2013a), "A wafer-based monocrystalline silicon photovoltaics road map: Utilizing known technology improvement opportunities for further reductions in manufacturing costs," *Solar Energy Materials and Solar Cells*, 114, 110–135.

Goodrich, A. C., D. M. Powell, T. L. James, M. Woodhouse, and T. Buonassisi (2013b), "Assessing the Drivers of Regional Trends in Solar Photovoltaic Manufacturing," *Energy & Environmental Science*, 6, 2811 – 2821.

Göx, R. (2002), "Capacity Planning and Pricing under Uncertainty." *Journal of Management Accounting Research*, 14, 59–78.

Greene, W. H. (2003), *Econometric Analysis*, Pearson.

GTM (2012), "GTM Research PV Technology, Production and Cost Outlook: 2012 – 2016," Tech. rep.

Hilton, R. (2005), *Managerial Accounting: Creating Value in a Dynamic Business Environment*, Irwin/McGraw-Hill, New York., 6th ed.

Hoechle, D. (2007), "Robust Standard Errors for Panel Regressions with Cross-Sectional Dependence," *Stata Journal*, 7(3), 281.

Jorgenson, D. (1963), "Capital Theory and Investment Behavior," *American Economic Review*, 53(2), 247–259.

Lieberman, M. B. (1984), "The Learning Curve and Pricing in the Chemical Processing Industries," *The RAND Journal of Economics*, 15(2), 213 – 228.

Lux Research (2012a), "Q4 2012 Supply Tracker," Available by subscription to Lux Research.

Lux Research (2012b), "Turning Lemons into Lemonade: Opportunities in the Turbulent PV Equipment Market," Available by subscription to Lux Research.

Lux Research (2014), "Q1 2014 Supply Tracker," Available by subscription to Lux Research.

McNichols, M., M. Rajan, and S. Reichelstein (2014), "Conservatism Correction for the Market-to-Book Ratio and Tobin's q," *Review of Accounting Studies*, 58(5), 980–995.

McWatters, C., and J. Zimmerman (2015), *Management Accounting in a Dynamic Environment*, Routledge.

Narayanan, V. (2003), "Activity-based pricing in a Monopoly." *Journal of Accounting Research*, 41, 473–502.

Nezlobin, A. (2012), "Accrual Accounting, Informational Sufficiency, and Equity Valuation," *Journal of Accounting Research*.

Nezlobin, A., M. V. Rajan, and S. Reichelstein (2012), "Dynamics of Rate-of-Return Regulation," *Management Science*, 58(5), 980–995.

Pfeiffer, T., U. Schiller, and M. Wagner (2011), "Cost-Based Transfer Pricing," *Review of Accounting Studies*, 16(2), 219–246.

Pillai, U. (2015), "Drivers of cost reduction in solar photovoltaics," *Energy Economics*, 50, 286–293.

Pillai, U., and J. McLaughlin (2013), "A model of competition in the solar panel industry," *Energy Economics*, 40, 32–39.

Pittman, B. (2009), "Whom Are You Calling Irrational? Marginal Costs, Variable Costs, and the Pricing Practices of Firms," *Working Paper, Department Of Justice, Washington DC*.

Powell, D., M. Winkler, H. Choi, C. Simmons, D. Needleman, and T. Buonassisi (2012), "Crystalline silicon photovoltaics: a cost analysis framework for determining technology pathways to reach baseload electricity costs," *Energy & Environmental Science*, 5, 5874 – 5883.

Powell, D., M. Winkler, A. Goodrich, and T. Buonassisi (2013), "Modeling the Cost and Minimum Sustainable Price of Crystalline Silicon Photovoltaic Manufacturing in the United States," *IEEE Journal of Photovoltaics*.

PWC (2012), "Financial Engineering as a Means to Support Jawaharlal Nehru Solar Mission," Tech. rep., PWC.

Rajan, M., and S. Reichelstein (2009), "Depreciation Rules and The Relation Between Marginal and Historical Cost," *Journal of Accounting Research*, 47, 1 – 43.

Reichelstein, S., and A. Rohlfing-Bastian (2015), "Levelized Product Cost: Concept and Decision Relevance," *The Accounting Review*, 55, 117 – 127.

Rogerson, W. (2008), "Inter-Temporal Cost Allocation and Investment Decisions," *Journal of Political Economy*, 116, 931 – 950.

Rogerson, W. (2011), "On the Relationship between Historic Cost, Forward Looking Cost and Long Run Marginal Cost," *Review of Network Economics*, 10(2).

Spence, A. M. (1981), "The learning curve and competition," *The Bell Journal of Economics*, 49–70.

Sridharan, S. (2015), "Volatility Forecasting Using Financial Statement Information," *The Accounting Review*, 90(5), 2079–2106.

Stobaugh, R. B., and P. L. Townsend (1975), "Price forecasting and strategic planning: the case of petrochemicals," *Journal of Marketing Research*, 19–29.

Swanson, R. (2011), "The Silicon Photovoltaic Roadmap," The Stanford Energy Seminar.