Online Appendix for:
Demanding Innovation: The Impact of Consumer Subsidies on Solar Panel Production Costs

Todd D. Gerarden*

June 2022

Contents

*Charles H. Dyson School of Applied Economics and Management, Cornell University, Warren Hall, Ithaca, NY 14853. Email: gerarden@cornell.edu. Website: toddgerarden.com.

Appendix A Additional Figures and Tables

Figure A.1: Comparison of Two Solar Panels

Output (Watts)	275	330
Size (cells)	60	72
Energy Conversion Efficiency (%)	16.8	17.0
Relative Price ($/W)	**1.00**	**1.01**

Notes: Individual solar panels are rated at different output levels, come in a few different standardized physical sizes (measured here by the number of solar cells), and can be different colors. Despite the differences between these two example solar panels in physical size, power output, and appearance, their prices are very similar when measured using the industry convention of $/Watt.

Data Source: Details for two Canadian Solar models from ressupply.com, accessed on September 24, 2017.

Figure A.2: Raw Data used in Demand Estimation

Notes: These figures plot the raw data used in demand estimation: quarterly observations of each market's quantity of solar panel sales (GW), average price ($/W), and subsidy level. There is no clearly-defined subsidy for Other, and so the effect of subsidies is not modeled for that market.

Data Source: IHS Markit and the International Energy Agency.

Figure A.3: Variation in the Input Prices used as Instrumental Variables

(a) Weighted Average Solar Panel Price

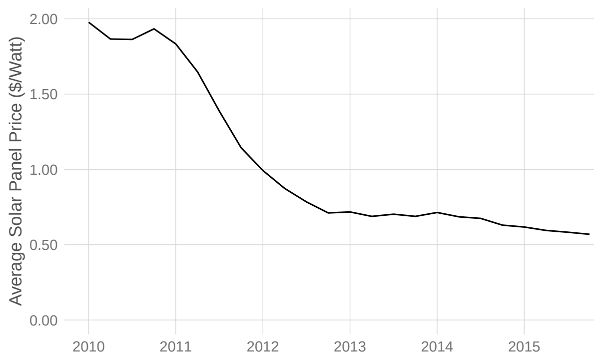

(b) IV: Polysilicon Price

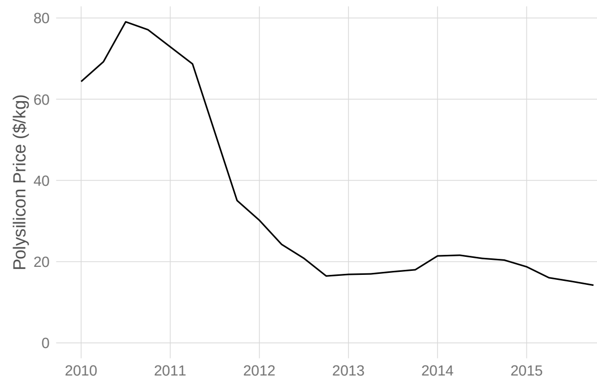

(c) IV: Silver Price

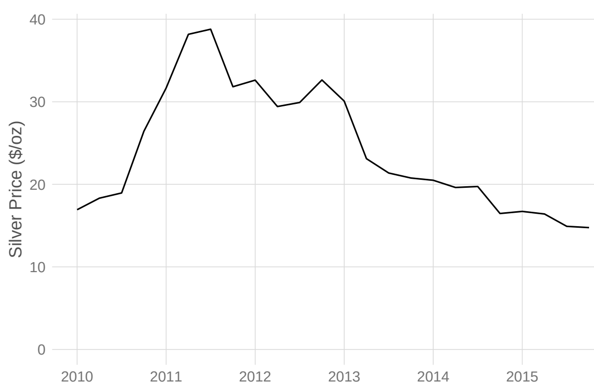

(d) IV: Aluminum Price

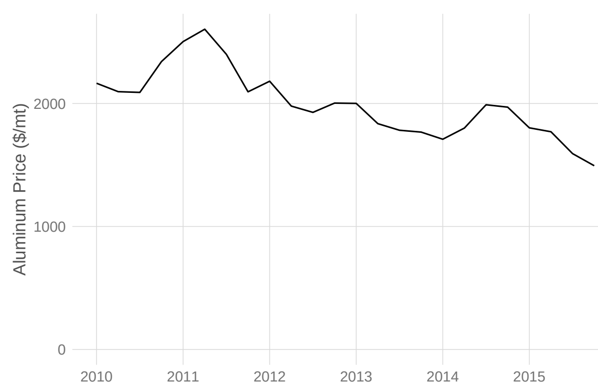

Notes: These figures summarize time series variation in the price of solar panels as well as three production inputs used as instrumental variables for solar panel prices: polysilicon, silver, and aluminum. Solar panel prices vary over time and across markets. The prices for inputs vary over time but not across markets.

Data Source: IHS Markit, The Silver Institute, and FRED.

Table A.1: Regional Market Shares for Each Firm

	Firms' Regional Shares			
Firm	Germany	Japan	ROW	USA
1	13.9	17.6	48.6	19.9
2	11.0	8.8	76.5	3.7
3	12.7	3.8	56.1	27.4
4	12.8	22.0	48.0	17.2
5	12.2	4.1	83.4	0.3
6	6.5	23.3	65.5	4.7
7	7.3	4.1	74.6	14.1
8	7.2	76.1	12.1	4.7
9	18.2	14.1	46.7	20.9
10	6.8	14.8	67.3	11.2
11	8.3	69.3	14.3	8.1
12	28.8	2.9	32.3	35.9
13	16.2	10.2	58.8	14.7
14	11.3	9.7	53.3	25.7
15	22.7	12.0	49.8	15.5

Notes: This table presents the share of each firms' sales (in Watts) in each market in percentage points. The rows all sum to 100. Comparisons of market shares across firms within a single column summarize the extent of variation in firms' exposure to each market (in relative terms, after normalizing to account for differences in the size of each firm). Firm names are masked to avoid disclosure of confidential data.

Data Source: Author's calculations using data described in Section 4.

Figure A.4: Estimated Demand States with and without Subsidies

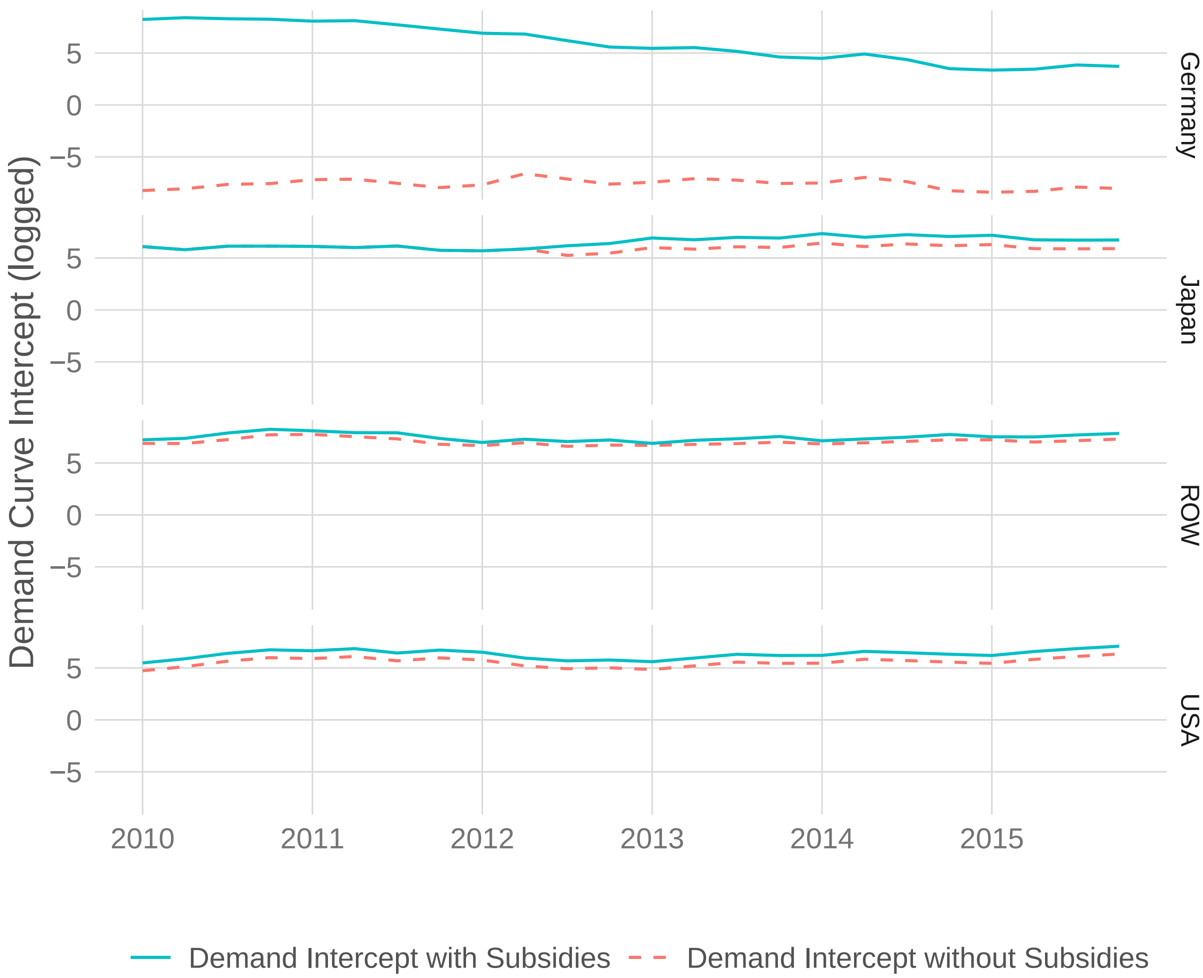

Notes: This figure plots estimated demand states with and without subsidies. The solid lines are subsidy-inclusive demand states (i.e., $\hat{d}_{mt}$) recovered from estimating equation 2 and inserting the estimates into equation 3. The dashed lines are the counterfactual demand states had the subsidies not been in place. Both are in terms of the natural logarithm of quantity, as they represent the demand curve intercept from equation 2. ROW denotes the residual market ("Rest of the World").

Data Source: Author's calculations using data described in Section 4 and model described in Section 5.

Table A.2: Estimates of Exogenous State Transitions

	w	Germany	Japan	ROW	USA
Constant	-0.04	-0.08	0.89	2.21	1.82
	(0.02)	(0.25)	(0.68)	(1.18)	(0.93)
lag(w)	0.96				
	(0.03)				
lag(Germany)		0.98			
		(0.04)			
lag(Japan)			0.87		
			(0.10)		
lag(ROW)				0.71	
				(0.16)	
lag(USA)					0.72
					(0.15)
$\hat{\Sigma}_{\xi}$	0.003	0.106	0.066	0.068	0.076
Observations	23	23	23	23	23

Notes: This table presents coefficients from estimation of the vector autoregression model for the exogenous states described by equation 5. The estimates are consistent with a stationary process for the exogenous states, as the point estimates of the lag coefficients are all less than one in absolute value.

Data Source: Author's calculations using data described in Section 4 and model described in Section 5.

Table A.3: Investment Policy Function Estimates

	Investment (x_{it})
η_{it}	31.2
	(69.7)
$\eta_{it}\bar{\eta}_t$	38.4
	(67.3)
Intercept	-36.6
	(20.3)
Observations	345
Log Likelihood	-179.0

Notes: This table presents policy function estimates from a logit model of whether a firm advanced its energy conversion efficiency as a function of its state. The regressors were selected in a first stage using penalized maximum likelihood estimation (i.e., lasso) of a logit model of whether a firm advanced its energy conversion efficiency as a flexible function of all state variables and their interactions.

Data Source: Author's calculations using data described in Section 4 and model described in Section 5.

Appendix B Additional Demand Estimation Results

Table B.1: Estimated Subsidy Coefficients

	OLS	IV: Input Prices	IV: Metal Prices	IV: Other Prices
	(1)	(2)	(3)	(4)
Australia (AUD/kWh)	2.29***	2.19***	2.16***	2.32***
	(0.72)	(0.73)	(0.69)	(0.73)
China (CNY/kWh)	1.61***	1.66***	1.67***	1.60***
	(0.41)	(0.41)	(0.41)	(0.41)
Germany (EUR-ct/kWh)	4.14***	4.25***	5.94***	4.08***
	(0.55)	(0.64)	(1.16)	(0.56)
India (INR/kWh)	0.58***	0.59***	0.59***	0.58***
	(0.10)	(0.10)	(0.11)	(0.10)
Italy (EUR-ct/kWh)	0.83***	0.82***	0.82***	0.83***
	(0.09)	(0.09)	(0.08)	(0.09)
Japan (JPY/kWh)	0.16*	0.25***	0.25***	0.21***
	(0.09)	(0.09)	(0.09)	(0.08)
Other (N/A)				
USA (N/A)				
Region FE	X	X	X	X
ln(Price)	X	X	X	X
Min. F-stat		18.96	5.98	53.6
Within R^2	0.7	–	–	–
Observations	192	192	192	192
Note:				*p<0.1; **p<0.05; ***p<0.01

Notes: This table presents coefficients on the inverse hyperbolic sine of each country's feed-in tariff (i.e., $\hat{\alpha}_{sm}$ from equation 2, with $f(\cdot) = \text{arcsinh}(\cdot)$). Each row corresponds to a different market (m) and each column presents estimates from a different estimation strategy. For Other, which aggregates demand in all other markets, there is no clearly defined subsidy; thus, the effects of subsidies in Other are not modeled and as a result there are no coefficients. There are no coefficients for USA because the ITC and accelerated depreciation are a time-invariant fraction of system prices rather than a time-varying price for electricity output; thus, their effects are identified using variation in prices rather than subsidies. Min. F-stat is the minimum F-statistic across market-specific tests of the hypothesis that the excluded instruments are jointly irrelevant. Heteroskedasticity-consistent standard errors are in parentheses.

Data Source: Author's calculations using data described in Section 4 and model described in Section 5.

Figure B.1: Estimated Demand States with and without Subsidies: Before Aggregation

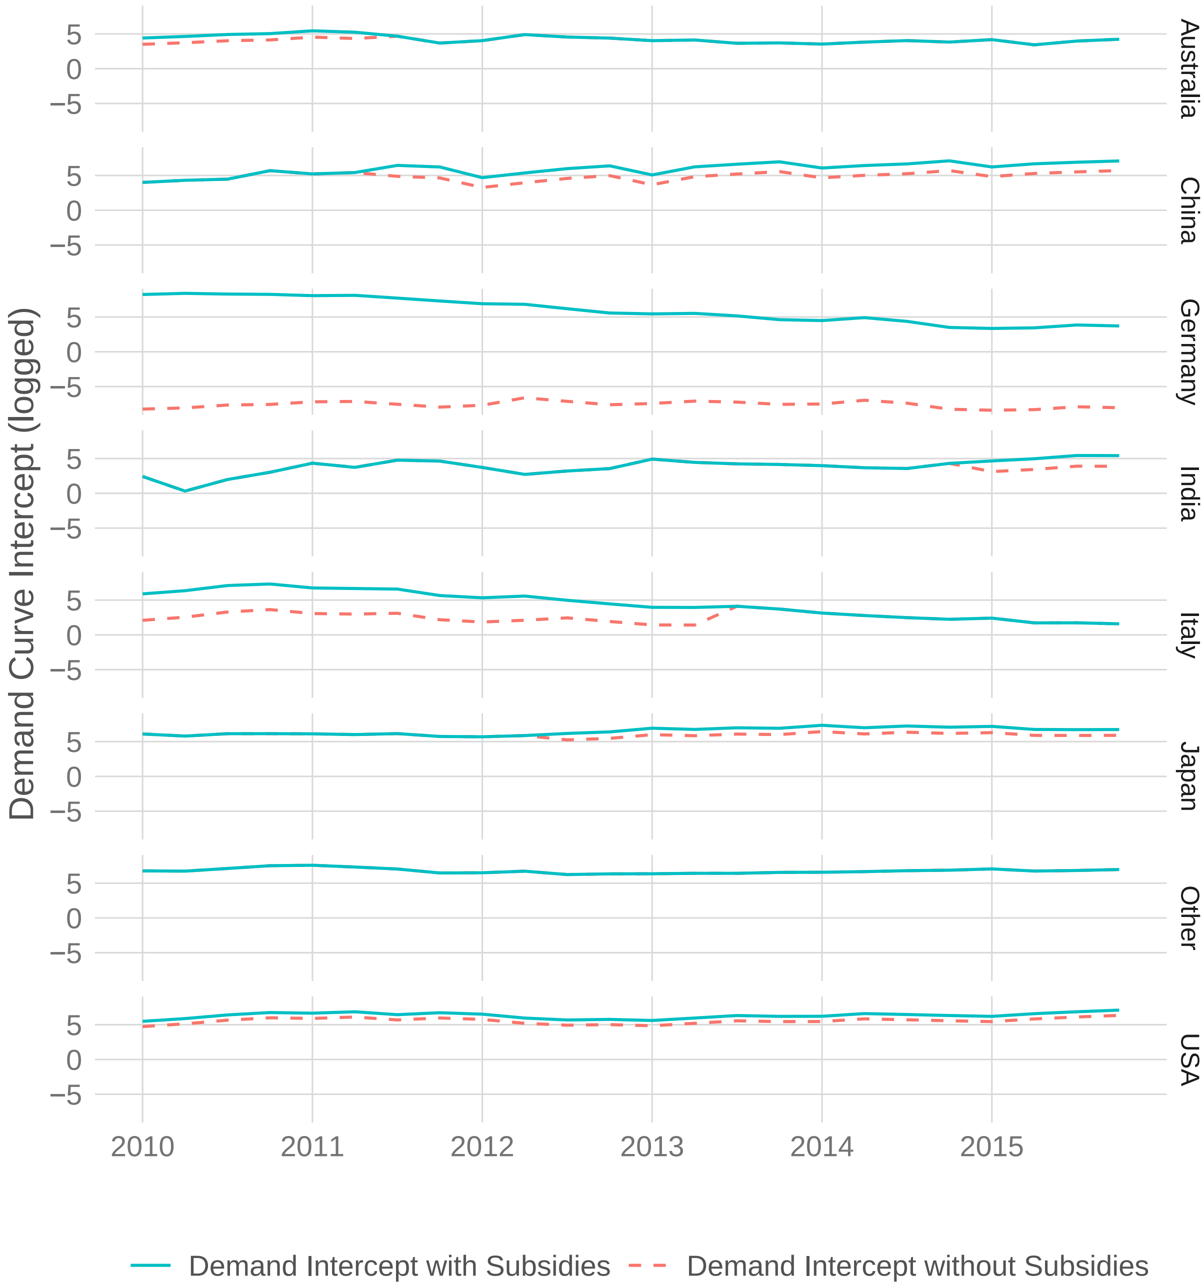

Notes: This figure plots estimated demand states with and without subsidies before aggregation into four regional markets for modeling competition among manufacturers. The solid lines are subsidy-inclusive demand states (i.e., $\hat{d}_{mt}$) recovered from estimating equation 2 and inserting the estimates into equation 3. The dashed lines are the counterfactual demand states had the subsidies not been in place. Both are in terms of the natural logarithm of quantity, as they represent the demand curve intercept from equation 2. "Other" denotes the combined demand from all other markets, for which there is no clearly defined subsidy; thus, the effects of subsidies in Other are not modeled and as a result the two lines coincide.

Data Source: Author's calculations using data described in Section 4 and model described in Section 5.

Table B.2: Estimated Demand Elasticities:

Alternative Subsidy Parameterization: $\ln(S_{mt})$ rather than $\text{arcsinh}(S_{mt})$

	OLS	IV: Input Prices	IV: Metal Prices	IV: Other Prices
	(1)	(2)	(3)	(4)
Germany	-2.53^{***}	-2.65^{***}	-4.38^{***}	-2.46^{***}
	(0.51)	(0.56)	(1.14)	(0.51)
Japan	-2.27^{***}	-2.24^{***}	-2.34^{***}	-2.23^{***}
	(0.12)	(0.11)	(0.13)	(0.14)
ROW	-0.85^{***}	-0.84^{***}	-0.82^{***}	-0.86^{***}
	(0.25)	(0.25)	(0.28)	(0.25)
USA	-1.49^{***}	-1.39^{***}	-1.48^{***}	-1.48^{***}
	(0.26)	(0.26)	(0.32)	(0.26)
Region FE	X	X	X	X
ln(Subsidy)	X	X	X	X
Min. F-stat		48.87	6.39	154.35
Within R^2	0.42	–	–	–
Observations	128	128	128	128
Note:				*p<0.1; **p<0.05; ***p<0.01

Notes: This table presents estimated price elasticities of demand (i.e., $\hat{\alpha}_{pm}$ from equation 2). Each row corresponds to a different market (m) and each column presents estimates from a different estimation strategy. The first column presents estimates of equation 2 using ordinary least squares. The second column presents IV estimates that use the prices of polysilicon, silver, and aluminum as instruments for the price of solar panels. The third column uses only the prices of silver and aluminum as instruments for the price of solar panels. The final column uses the prices of solar panels in markets other than the market of interest as instruments for the price in the market of interest. The coefficients are not statistically distinguishable from the baseline estimates presented in Table 2. The coefficients for the U.S. market are unchanged because the subsidies for that market are a fraction of the solar panel price as described in Section 5. Min. F-stat is the minimum F-statistic across market-specific tests of the hypothesis that the excluded instruments are jointly irrelevant. Heteroskedasticity-consistent standard errors are in parentheses.

Data Source: Author's calculations using data described in Section 4 and model described in Section 5.

Table B.3: Estimated Demand Elasticities:

Alternative Subsidy Parameterization: $\ln(1 + S_{mt})$ rather than $\text{arcsinh}(S_{mt})$

	OLS	IV: Input Prices	IV: Metal Prices	IV: Other Prices
	(1)	(2)	(3)	(4)
Germany	-2.54^{***}	-2.66^{***}	-4.50^{***}	-2.47^{***}
	(0.51)	(0.57)	(1.19)	(0.52)
Japan	-1.47^{***}	-1.11^{***}	-1.09^{***}	-1.27^{***}
	(0.31)	(0.29)	(0.32)	(0.26)
ROW	-1.37^{***}	-1.33^{***}	-1.31^{***}	-1.38^{***}
	(0.24)	(0.24)	(0.22)	(0.24)
USA	-1.49^{***}	-1.39^{***}	-1.48^{***}	-1.48^{***}
	(0.26)	(0.25)	(0.32)	(0.26)
Region FE	X	X	X	X
ln(1 + Subsidy)	X	X	X	X
Min. F-stat		19.42	6.17	54.63
Within R^2	0.7	–	–	–
Observations	192	192	192	192
Note:			$^{*}p<0.1$; $^{**}p<0.05$; $^{***}p<0.01$	

Notes: This table presents estimated price elasticities of demand (i.e., $\hat{\alpha}_{pm}$ from equation 2). Each row corresponds to a different market (m) and each column presents estimates from a different estimation strategy. The first column presents estimates of equation 2 using ordinary least squares. The second column presents IV estimates that use the prices of polysilicon, silver, and aluminum as instruments for the price of solar panels. The third column uses only the prices of silver and aluminum as instruments for the price of solar panels. The final column uses the prices of solar panels in markets other than the market of interest as instruments for the price in the market of interest. The coefficients are not statistically distinguishable from the baseline estimates presented in Table 2. The coefficients for the U.S. market are unchanged because the subsidies for that market are a fraction of the solar panel price as described in Section 5. Min. F-stat is the minimum F-statistic across market-specific tests of the hypothesis that the excluded instruments are jointly irrelevant. Heteroskedasticity-consistent standard errors are in parentheses.

Data Source: Author's calculations using data described in Section 4 and model described in Section 5.

Appendix C Marginal Cost Estimate Robustness

Table C.1: Results under Alternative Production Cost Specifications

(a) Includes Time Period x Chinese Manufacturer FEs

	(1)	(2)	(3)	(4)	(5)
$\ln(\tilde{\eta}_{it})$	-5.43^{***}	-0.91^{***}	-0.16^{*}	-1.81^{***}	-0.34^{**}
	(0.16)	(0.10)	(0.10)	(0.13)	(0.16)
$\ln(w_t)$		0.66^{***}		0.60^{***}	
		(0.01)		(0.01)	
Period x China FE			X		X
Firm FE				X	X
Observations	1,352	1,352	1,352	1,352	1,352
Adjusted R^2	0.47	0.89	0.91	0.90	0.92

Note: *p<0.1; **p<0.05; ***p<0.01
Data include 24 periods (T) for 4 markets (M).

(b) Includes Scale Measure

	(1)	(2)	(3)	(4)	(5)
$\ln(\tilde{\eta}_{it})$	-5.10^{***}	-0.87^{***}	-0.15	-1.76^{***}	-0.27^{*}
	(0.16)	(0.10)	(0.09)	(0.14)	(0.16)
$\ln(w_t)$		0.65^{***}		0.60^{***}	
		(0.01)		(0.01)	
Time Period FE			X		X
Firm FE				X	X
Observations	1,352	1,352	1,352	1,352	1,352
Adjusted R^2	0.49	0.89	0.91	0.90	0.92

Note: *p<0.1; **p<0.05; ***p<0.01
Data include 24 periods (T) for 4 markets (M).

(c) Includes Experience Measure

	(1)	(2)	(3)	(4)	(5)
$\ln(\tilde{\eta}_{it})$	-3.18^{***}	-0.60^{***}	-0.08	-1.16^{***}	-0.30^{*}
	(0.17)	(0.10)	(0.09)	(0.15)	(0.16)
$\ln(w_t)$		0.61^{***}		0.54^{***}	
		(0.01)		(0.01)	
Time Period FE			X		X
Firm FE				X	X
Observations	1,352	1,352	1,352	1,352	1,352
Adjusted R^2	0.60	0.90	0.91	0.91	0.92

Note: *p<0.1; **p<0.05; ***p<0.01
Data include 24 periods (T) for 4 markets (M).

(d) Mean instead of Max Energy Conversion Efficiency

	(1)	(2)	(3)	(4)	(5)
$\ln(\tilde{\eta}_{it})$	-6.32^{***}	-1.17^{***}	-0.23^{*}	-1.65^{***}	-0.41^{***}
	(0.16)	(0.12)	(0.13)	(0.14)	(0.16)
$\ln(w_t)$		0.63^{***}		0.60^{***}	
		(0.01)		(0.01)	
Time Period FE			X		X
Firm FE				X	X
Observations	1,239	1,239	1,239	1,239	1,239
Adjusted R^2	0.57	0.88	0.91	0.90	0.92

Note: *p<0.1; **p<0.05; ***p<0.01
Data include 24 periods (T) for 4 markets (M).

(e) Excludes Time Periods with Capacity Constraints

	(1)	(2)	(3)	(4)	(5)
$\ln(\tilde{\eta}_{it})$	-3.72^{***}	-0.94^{***}	-0.14	-2.12^{***}	-0.37^{*}
	(0.19)	(0.12)	(0.12)	(0.17)	(0.22)
$\ln(w_t)$		0.65^{***}		0.58^{***}	
		(0.01)		(0.02)	
Time Period FE			X		X
Firm FE				X	X
Observations	983	983	983	983	983
Adjusted R^2	0.29	0.78	0.89	0.87	0.89

Note: *p<0.1; **p<0.05; ***p<0.01
Data include 17 periods (T) for 4 markets (M).

Notes: These tables present coefficients from alternative specifications of the model in equation 4. The dependent variable is the natural logarithm of estimated marginal cost. The first row contains coefficients on energy conversion efficiency ($\tilde{\eta}_{it}$). The second row contains coefficients on the price of polysilicon (w_t). The finding that energy conversion efficiency lowers production cost is robust: in model 3, the specification used to estimate the dynamic model, the coefficient on energy conversion efficiency in each of these alternative models is statistically indistinguishable from the baseline specification in Table 3.
Data Source: Author's calculations using data described in Section 4 and model described in Section 5.

Figure C.1: Aggregate Capacity Utilization over Time

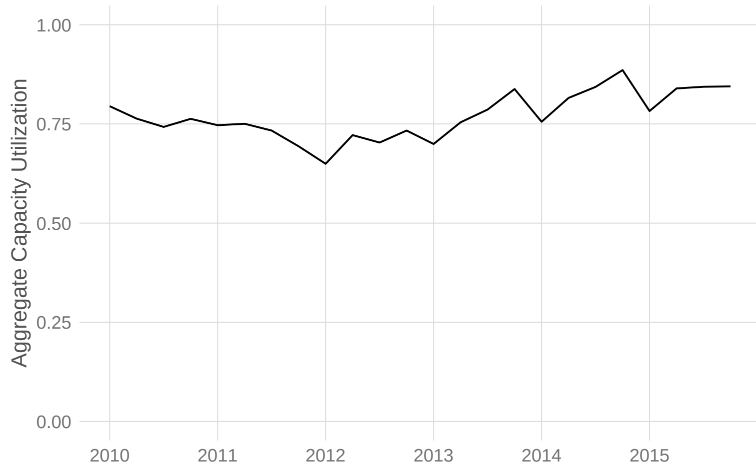

Notes: This figure plots the ratio of total module production to total module manufacturing capacity for firms in the sample over time.

Data Source: Author's calculations using data described in Section 4.

Appendix D Summary of Inferred Marginal Costs

As discussed in Sections 3 and 5, this paper maintains the assumption that firms engage in Cournot competition. If this assumption holds, estimates of firms' marginal costs can be computing using a combination of data (firms' market shares and prices) and demand estimates (price elasticities of demand). To help readers assess this assumption and its implications, this appendix provides additional details on these data and estimates.

First, Figure D.1 presents graphical summaries of market shares and estimated marginal costs over time. Figure D.1a plots each firms' quantity sold as a share of the global market over time. There is significant variation in market shares in the cross section – ranging from just above 0% to nearly 20% – as well as variation in the relative market shares of firms over time.[1]

Figure D.1: Summary of Firms' Market Shares and Marginal Costs over Time

(a) Market Shares (b) Estimated Marginal Costs

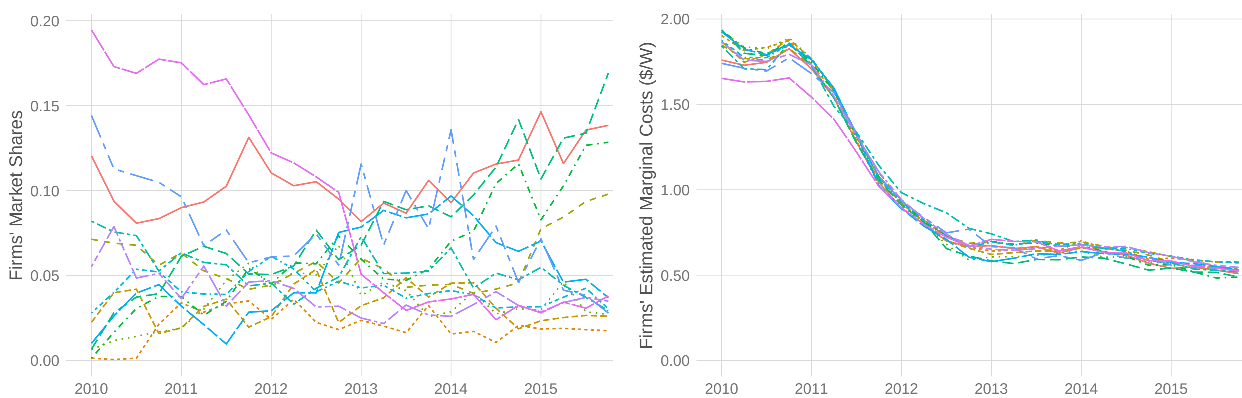

Notes: These figures summarize variation over time in firms' global market shares (a) and estimated marginal costs (b). Figure a is constructed using data without any economic assumptions. Marginal costs in Figure b are inferred from firms' first order conditions for optimal production using data and estimated demand parameters. This produces estimates at the firm-market-time level, and these are averaged across markets to create a time series for each firm.

Data Source: Author's calculations using data described in Section 4 and model described in Section 5.

As described in Section 5, firm i's optimal quantity choice in market m and time period t is implicitly defined by its first order condition for profit maximization. This first order condition can be rearranged to express the firm's unobserved marginal production cost as

[1] Table 1 in the main text presents summary statistics for market shares within each regional market.

a function of its market share, the price, and the price elasticity of demand. This allows estimation of each firm's marginal cost in each market and time period. To summarize this information, I compute the average of these estimated marginal costs across markets for firm i in time period t (weighting by the quantity sold in each market). Figure D.1b plots these firm-specific estimated marginal costs over time. These marginal costs estimates exhibit intuitive patterns. First, they fall significantly over time, as observed equilibrium prices do. Second, firms with higher market shares have lower estimated marginal costs. This is most evident early in the sample period when there was more dispersion in market shares and estimated marginal costs.

While Figure D.1 is informative about variation across firms and over time, it does not present information on variation across markets for a given firm and time period. One potential reason for firms' unobserved marginal costs to differ across markets in a given time period is transport costs. While transport costs are small relative to product prices, they may still affect firms' decisions on the margin. For example, firms may have a lower (unobserved) marginal cost to supply solar panels to the market region(s) in which they produce solar panels. Under Cournot competition, this would lead firms to have higher market shares in the locations in which they manufacture, all else equal.

To assess whether the market share data and marginal cost estimates conform to this logic, I estimate the following regression model

$$y_{imt} = \alpha + \beta \text{ Domestic Manufacturer}_{imt} + \varepsilon_{imt}, \tag{D.1}$$

where Domestic Manufacturer$_{imt}$ is a binary indicator for whether firm i actively manufactures solar panels in market m and time period t. I estimate this model via ordinary least squares using two outcome variables: market shares and marginal cost estimates (in natural logarithms). While the model is meant for purely descriptive purposes, I also include market and time fixed effects in some specifications to account for the possibility of correlated unobservables.

Table D.1 presents estimates of equation D.1. Columns 1 and 2 show that, on average, firms who manufacture within a given market have a higher market share than firms who do not manufacture within that market. Columns 3 and 4 present analogous regression estimates using estimated marginal costs (in natural logarithms) as the dependent variable. The results show that, on average, firms who manufacture within a given market have a lower estimated marginal cost than firms who do not manufacture within that market. Both of these results are consistent with the predictions of the Cournot model as discussed above.

Table D.1: Summary of Firms' Market Shares and Marginal Costs across Markets

	Market Share		ln(Marginal Cost)	
	(1)	(2)	(3)	(4)
Domestic Manufacturer	0.04^{***}	0.10^{***}	-0.06^{**}	-0.09^{***}
	(0.004)	(0.01)	(0.03)	(0.01)
Market FE		X		X
Time Period FE		X		X
Observations	1,352	1,352	1,352	1,352
Adjusted R^2	0.07	0.16	0.003	0.96
Note:		*p<0.1; **p<0.05; ***p<0.01		

Notes: This table presents estimates from ordinary least squares regressions of firms' market shares and marginal costs in a given market and time period on a binary indicator of whether the firm manufactures solar panels within that market and time period. Regressions in columns 2 and 4 include market and time fixed effects.

Data Source: Author's calculations using data described in Section 4 and model described in Section 5.

Appendix E Investment Outcomes

The model in Section 3 assumes that firms make a discrete choice of whether to invest, and that the improvement in energy conversion efficiency conditional on investing, ν_{it}, is i.i.d. across firms and time. I use data on energy conversion efficiency to assess whether this assumption is reasonable. To summarize, the tests that follow fail to reject this assumption, although they have low power due to small sample size.

Comparison of Efficiency Improvements by Firm The assumption on ν_{it} implies that all firms draw from the same distribution. To assess this visually, I plot observed changes in energy conversion efficiency by firm in Figure E.1a. I formalize this comparison by applying the Kolmogorov-Smirnov test to each pair of firms under the null hypothesis that each pair of sets of observed changes in energy conversion efficiency are drawn from the same distribution. I compute p-values, pool them across all the pairwise combinations, and evaluate the distribution of p-values at three quantiles:

Quantile	0.10	0.50	0.90
p-value	0.08	0.45	0.95

While this pooled comparison of individual tests is informal, the distribution of p-values is not heavily skewed toward zero. I also use a Kruskal-Wallis rank sum test for a joint (rather than pairwise) comparison under the null hypothesis that all firm-level samples originate from the same distribution. The result is a p-value of 0.14.

Comparison of Efficiency Improvements by Time Period I plot observed changes in energy conversion efficiency by date in Figure E.1b. I apply the Kolmogorov-Smirnov test to each pair of time periods under the null hypothesis that each pair of sets of observed changes in energy conversion efficiency are drawn from the same distribution. I compute p-values, pool them across all the pairwise combinations, and evaluate the distribution of p-values at three quantiles:

Quantile	0.10	0.50	0.90
p-value	0.20	0.67	1.00

I use a Kruskal-Wallis rank sum test for a joint (rather than pairwise) comparison under the null hypothesis that all time-level samples originate from the same distribution. The result is a p-value of 0.25.

Serial Correlation Perhaps the most important economic implication of the assumption on ν_{it} is that it rules out serial correlation in a given firm's outcomes over time. To assess this, I plot the autocorrelation function by firm in Figure E.2. This visual test relies on a small number of non-missing observations in the time series for each firm because firms rarely invest, so it is only suggestive. Still, I informally interpret the plots as failing to reject the null that ν_{it} is not serially correlated within a firm.

Figure E.1: Investment Realizations

(a) Investment Realizations (ν_{it}) by Firm (i)

(b) Investment Realizations (ν_{it}) by Date (t)

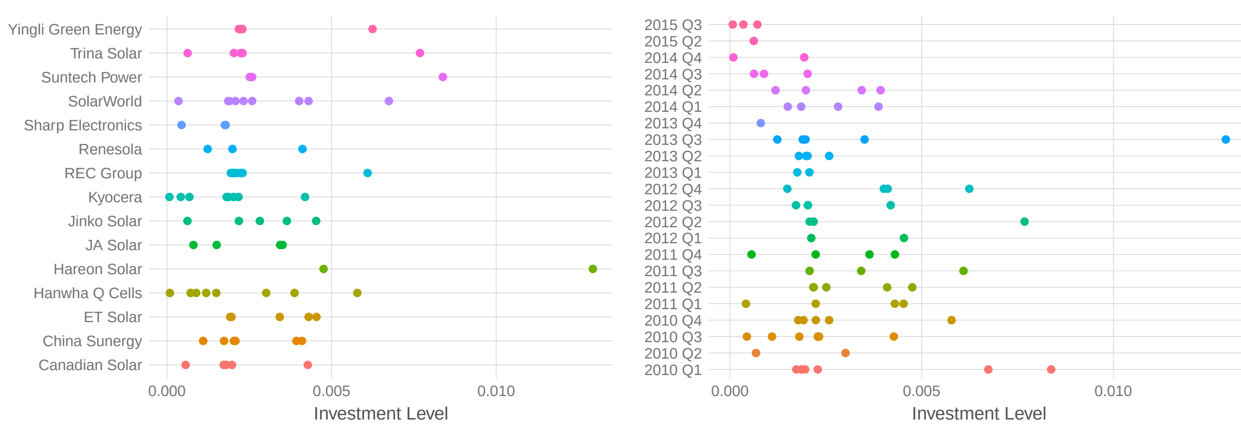

Figure E.2: Autocorrelation of Investment Level (ν_{it}) within a Firm

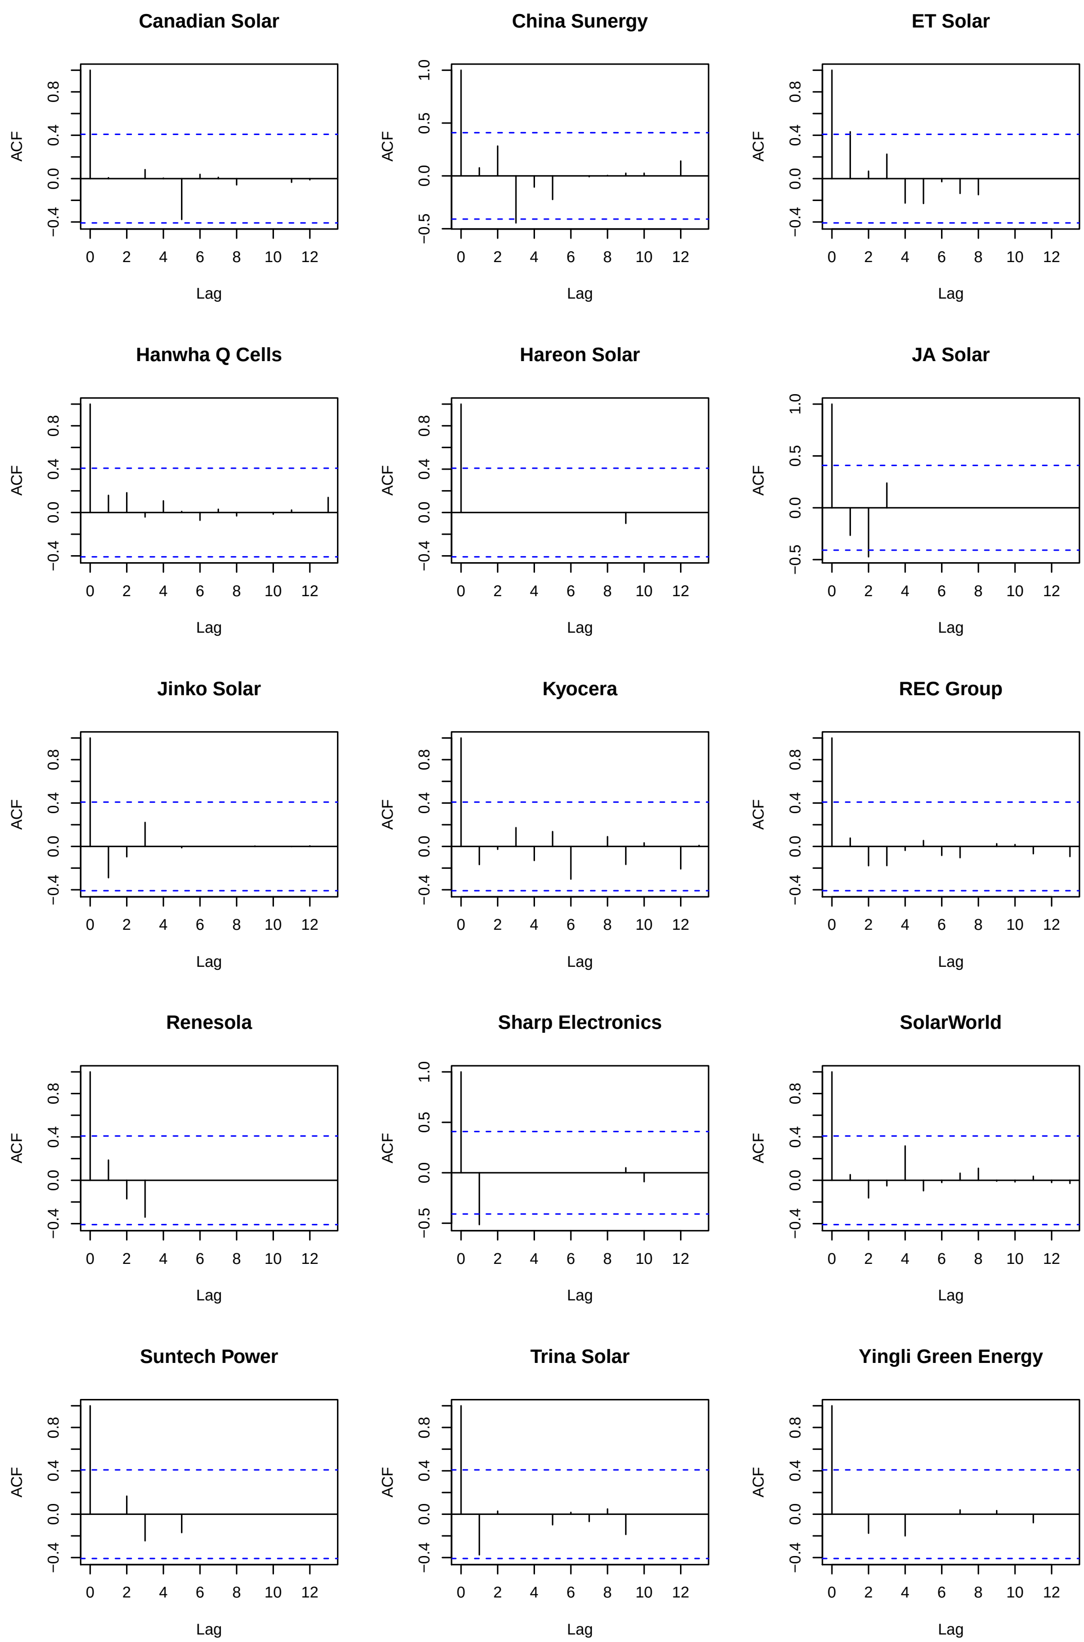

Appendix F Dynamic Parameter Estimation

F.1 Pseudo Maximum Likelihood Estimation

The firm's investment optimization problem is:

$$V_i(s_t) = \max_{x_{it} \in \{0,1\}} \bar{\pi}_i(s_t) - \gamma x_{it} + \sigma \varepsilon_{it}(x_{it}) + \beta E\left[V_i(s_{t+1}) | s_t, x_{it} \right],$$

where the dependence of the value function on strategies (ζ), private shocks (ε_{it}), and parameters (θ) are omitted for clarity. Under the assumption that the choice-specific error terms – $\varepsilon_{it}(x_{it})$ – are i.i.d. Type I extreme value, the *ex-ante* probability of investment is:

$$\Pr(x_{it} = 1) = \frac{\exp(v_1)}{\exp(v_0) + \exp(v_1)} \tag{F.1}$$

where

$$v_0 = \frac{1}{\sigma}\left(\bar{\pi}_i(s_t) + \beta E\left[V_i(s_{t+1}) | s_t, x_{it} = 0 \right] \right) \quad \text{and}$$
$$v_1 = \frac{1}{\sigma}\left(\bar{\pi}_i(s_t) - \gamma + \beta E\left[V_i(s_{t+1}) | s_t, x_{it} = 1 \right] \right).$$

The log-likelihood is:

$$\ln \mathcal{L}(\theta | x) = \sum_i \sum_t (1 - x_{it}) v_0 + x_{it} v_1 - \ln\left(\exp(v_0) + \exp(v_1) \right).$$

Evaluating the log-likelihood requires knowledge of $E\left[V_i(s_{t+1}) | s_t, x_{it} \right]$. I approximate this expectation by combining the forward simulation procedure described in Section 5 with value function approximation.[2] I rewrite $V_i(s_t)$ as $W_i(s_t; \zeta) \cdot \theta$ and forward-simulate expected profits, investments, and private shocks for all the states observed in the data in a preliminary step to estimate $W_i(s_t; \zeta)$. I then use $\widehat{W}_i(s_t; \zeta)$ to evaluate $\widehat{V}_i(s_t)$ given θ.

[2]In principle this could be done entirely via forward simulation, without the use of value function approximation, by drawing from the set of possible next period states for each state observed in the data and forward-simulating the value function (up to parameters) for each of these draws. The computational burden of this approach grows linearly with the number of draws used to compute the multidimensional integral defined by the expectation operator, so I use value function approximation to allow for a large number of draws.

Approximating the value function requires some simplification of the state space due to computational constraints.[3] I use a subset of the state space that contains all of the information needed to characterize the product market equilibrium, including firm i's profits, in time period t: $\phi_i(s_t) = [\ 1 \quad \eta_{it} \quad \bar{\eta}_t \quad d_{1t} \quad d_{2t} \quad ... \quad d_{Mt} \quad w_t\ ]$.[4] I then approximate the value function using a parametric function of these variables, $\widehat{V}_i(s_t;\theta) = \phi_i(s_t)'\lambda$. I estimate λ via ordinary least squares.

I use this parametric approximation to estimate $E\left[V_i\left(s_{t+1}\right)\middle|\, s_t, x_{it}\right]$. To do so, I take 1000 draws from the state transition processes to approximate the distribution of states that can be reached in one period from each state observed in the data (conditional on the firm's investment choice, x_{it}). I then use $\hat{\lambda}$ to predict the value function at each of these possible states and take the mean across states to recover an estimate of $E\left[V_i\left(s_{t+1}\right)\middle|\, s_t, x_{it}\right]$. I substitute these predictions into the log-likelihood function above, yielding a pseudo-log-likelihood. Finally, I search over θ, repeating these steps for each candidate parameter value, to maximize the pseudo-log-likelihood.

Finally, confidence intervals are constructed via bootstrap, resampling residuals from each stage of estimation prior to forward simulation 500 times. In practice it is difficult to separately identify γ and σ across the bootstrap samples using the value function approximation described above. For this reason, I fix σ at 30 based on the central estimate derived from the original sample and do not allow it to vary across bootstrap samples.

[3]This is because the model may fail to converge depending on the approximation used. This is a generic feature of parametric policy iteration, not something specific to my analysis (see, e.g., Benitez-Silva et al., 2000). Because I use forward simulation rather than solve the model for estimation, technically this does not preclude using a more complex approximating function for this estimation procedure. However, for consistency, I use the same function to approximate the value function in estimation as in solving the model for counterfactuals.

[4]This choice was inspired by the moment-based Markov equilibrium concept (Ifrach and Weintraub, 2017). $\phi_i(s_t)$ is a sufficient statistic for computing product market outcomes and firm i's profits under Cournot competition among firms with constant marginal costs if the equilibrium is an interior solution in which all firms produce non-zero quantities and no production constraints bind. See Bergstrom and Varian (1985) for a clear derivation. This is a special case of a broad class of aggregative games. Computing the equilibrium also requires parameters that govern demand and production costs, and the number of firms.

F.2 Sensitivity of Estimates to the Discount Factor, β

Table F.1 presents investment cost estimates using alternative discount factors. As the discount factor increases, the fixed cost of investing in energy conversion efficiency improvements also increases. This is because a higher discount factor puts more weight on future profits that result from technological advancement, and so a higher fixed cost is necessary to rationalize observed innovation patterns.

Table F.1: Sensitivity of Investment Cost Parameter Estimates to the Discount Factor

	Annual Discount Factor		
Parameter	$\beta = 0.85$	$\beta = 0.875$	$\beta = 0.90$
γ	95.0	108.6	129.3
σ	30.0	30.0	30.0

Notes: This table presents estimates of the fixed cost of energy conversion efficiency improvements (γ) under alternative assumptions on the discount factor used by firms. The scale parameter on the private shocks (σ) is fixed at 30 rather than estimated due to the practical challenge of separately identifying γ and σ. All numbers are in millions of dollars. The middle column corresponds to the estimates in Table 4.

Data Source: Author's calculations using data described in Section 4 and model described in Section 5.

Appendix G Counterfactual Computation

Solving the Model To accommodate continuous state variables, I solve the model using value function approximation. I approximate the value function using the same approach as in estimation, described in Appendix F. The state variables used for value function approximation are firm i's energy conversion efficiency η_{it}, the industry-average energy conversion efficiency $\bar{\eta}_t$, the vector of each markets' demands d_t, and the input price w_t. I initialize the model using the states observed in the data.

I use policy iteration to solve the model. Standard policy iteration is infeasible due to the continuous nature of the state space. Thus, I use parametric policy iteration (Benitez-Silva et al., 2000), adapting the methods of Sweeting (2013) to this context. This procedure is similar to standard policy iteration methods, but it relies on the use of value function approximation. Before starting the iterative procedure at state s_t, I construct the vector of inputs to value function approximation for each firm, $\phi_i(s_t)$. I denote the matrix of these values for all firms and states as Φ (subsuming the dependence on s_t here and elsewhere below for clarity). In each iteration j, the following steps are used:

1. Compute *ex-ante* expected flow profits, before firms realize the choice-specific private shocks $\varepsilon_{it}(x_{it})$:

$$E_{\varepsilon_{it}}\left[\pi_i(x_{it}, s_t; \varepsilon_{it})\right] = \bar{\pi}_i(s_t) - \gamma p_i^j(x_{it} = 1|s_t) + \sigma\left[\varkappa - \sum_{x_{it}\in\{0,1\}} p_i^j(x_{it}|s_t)\ln\left(p_i^j(x_{it}|s_t)\right)\right],$$

   where $p_i^j(x_{it} = 1|s_t)$ denotes firm i's probability of investment at state s_t in iteration j. The vector of these choice probabilities for all firms is denoted by P^j, and the vector of expected flow profits for all firms by $\widetilde{\Pi}(P^j)$.[5]

2. Compute $E_{P^j}\Phi$ by taking 1000 draws from the state transition processes to approximate the distribution of states that can be reached in one period from each state observed in the data, conditional on candidate choice probabilities P^j.

3. Compute $\widehat{\lambda}^{P^j} = \left((\Phi - \beta * E_{P^j}\Phi)'\,(\Phi - \beta * E_{P^j}\Phi)\right)^{-1}(\Phi - \beta * E_{P^j}\Phi)'\,\widetilde{\Pi}(P^j)$.

[5] For $j = 1$, this requires a starting value P^1; for this, I use the observed unconditional choice probability.

4. Use $\widehat{\lambda}^{P^j}$ to compute choice-specific value functions for each choice for each firm, and then use these in equation F.1 to compute new choice probabilities, P^{j+1}.

5. If $\max|P^{j+1} - P^j| < 10^{-4}$, the iterations stop. If not, iteration $j+1$ begins at step 1.

This procedure produces conditional choice probabilities in each state, P^*.

Simulating the Model Forward over Time To perform a comprehensive counterfactual analysis, and to simulate the model beyond the sample period, I solve the model many times. I start by solving the model for the observed states as described above. Then, to simulate the model forward over time, I start with period $t = 1$ – the first quarter of 2010 – and simulate forward one period. This consists of using the conditional choice probabilities from solving the model at $t = 1$ ($P^*_{t=1}$) and the parameters of the AR(1) process for d_t and w_t, to compute the expectation of the next state via numerical integration using 1000 draws for each of the state transition processes. This simulated state at $t = 2$ is used to solve the model again, as described above. I repeat this procedure $T = 80$ times to simulate the model forward 20 years.

Implementing the Counterfactual Subsidy Levels Solving the model for each state and simulating the model solution forward over time both require estimates of the transition processes for demand. For the *Baseline* scenario, these are the same estimated demand transition parameters as in estimation of the investment cost parameters. The estimates are reported in Table A.2.

If instead subsidies had not been offered to consumers, the observed states and their transition processes would have been different. To account for this, I re-estimate the demand transition processes before solving the *No Subsidies* counterfactual. First, I recover the counterfactual demand levels in each time period using the parameter estimates presented in Section 6 as follows: $\hat{d}_{mt} - \hat{\alpha}_{sm}\text{arcsinh}\left(S_{mt}\right)$. Then, I estimate equation 5 with these modified states, and use the parameter estimates to solve the model and simulate it forward over time. For the counterfactual scenario without German subsidies, but with subsidies in other markets, I only make this adjustment for Germany.

Summarizing the Counterfactual Results Simulating the model forward over time results in a trajectory of states over time for each counterfactual scenario. Figure G.1 summarizes the evolution of energy conversion efficiencies over time for each of these three scenarios. These states can can be used to compute equilibrium quantities and prices, as well as firms' profits and consumer surplus. I assume that estimated demand curves in the absence of subsidies represent the distribution of consumers' marginal willingness to pay for solar panels, and compute changes in consumer surplus by integrating between different quantity realizations along each curve.

Finally, I use the quantities to perform a back-of-the-envelope calculation using existing estimates of the environmental benefits of solar panel adoption to quantify the change in external benefits attributable to the subsidies. The external social benefits attributable to consumer subsidies depend on the quantity of solar panels adopted due to subsidies, the amount of electricity the solar panels produce, and the external damages associated with alternative electricity generation sources that solar electricity displaces.[6] Electricity generation estimates for each regional market come from PVWatts, a publicly available engineering tool that predicts solar electricity generation for different locations and solar system configurations.[7] I construct estimates of lifetime solar generation based on potential electricity output in one location in each of the four markets.[8] Estimates of the external damages attributable to electricity generation from natural gas come from Muller et al. (2011). I present numbers based on two estimates. The first, denoted EB_{CO_2}, only includes the value of carbon dioxide emissions avoided. I use the central estimates from Muller et al. (2011), which correspond to a social cost of carbon of approximately \$33 per ton in 2010 terms. The second, denoted EB_{ALL}, also includes the central estimates of local air pollution damages from Muller et al. (2011). These local air pollution damage estimates are based on historical data from the United States, so they may not be representative of the benefits from avoided local air pollution in other regions and time periods. All

[6]This calculation of external benefits builds on Gillingham and Tsvetanov (2019).

[7]PVWatts (http://pvwatts.nrel.gov) is developed by the National Renewable Energy Laboratory.

[8]I use the PVWatts default weather data for Germany, Japan, and the United States. For ROW, I use the default weather data from a search for "China" to calculate electricity generated from solar panels in the residual market because it is the median electricity generation prediction within the ROW submarkets, and also because it is the largest of the submarkets.

numbers from Muller et al. (2011) are converted to 2010 dollars. I assume a lifetime of 25 year for the solar panels, and construct the present discounted value of external benefits for each Watt of solar panels adopted in each time period using the same discount factor as in estimating and solving the model (β). Finally, I construct the sum of the present discounted value of the external benefits from solar panels adopted in each time period.

Figure G.1: Long-Run Counterfactual: Energy Conversion Efficiency

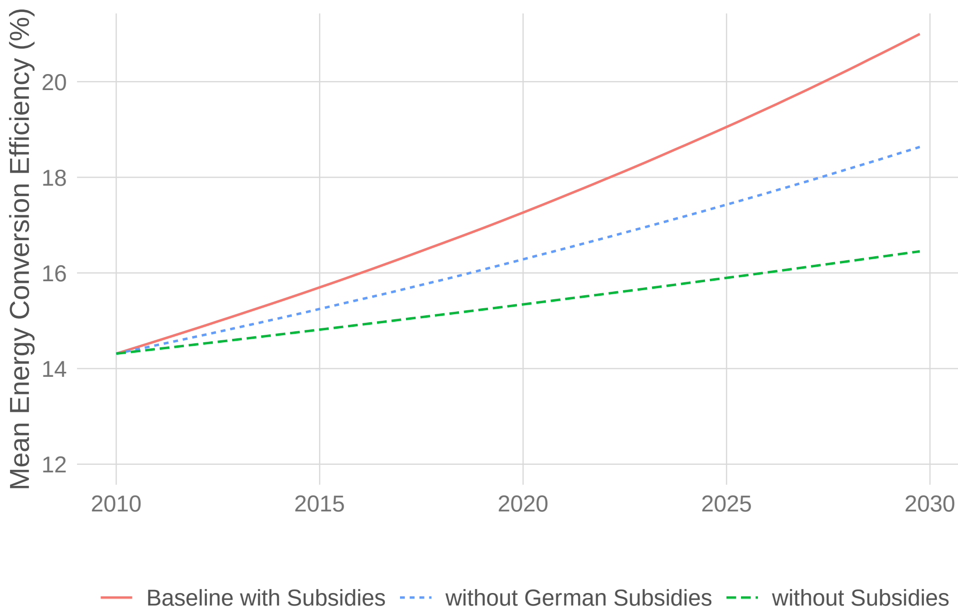

Notes: This figure plots average model-predicted energy conversion efficiencies over time with and without subsidies. The higher, solid line represents predicted efficiency levels from solving the model based on historical demand patterns including subsidies. The lower, dashed lines represent counterfactual efficiency levels based on solving the model after removing subsidies.
Data Source: Author's calculations based on the model estimation described in Section 5 and counterfactual exercise described in Section 7 using data described in Section 4.

Constructing Figure 5 Figure 5 plots quantity and price predictions under different assumptions. The quantity predictions are a sum across markets: $\widehat{Q}_t = \sum_m \widehat{Q}_{mt}$. The price predictions are a weighted average across markets: $\widehat{P}_t = \sum_m \widehat{P}_{mt}\widehat{Q}_{mt} \Big/ \sum_m \widehat{Q}_{mt}$. As described in Section 7, the underlying values $\widehat{Q}_{mt}$ and $\widehat{P}_{mt}$ are computed as follows:

$$\widehat{P}_{mt} = \frac{\sum_i \widehat{mc}_{it}}{I + \frac{1}{\hat{\alpha}_{pm}}} \qquad\qquad \widehat{Q}_{mt} = \exp(\hat{d}_{mt})\widehat{P}_{mt}^{\hat{\alpha}_{pm}}.$$

Baseline plots predictions from solving for the product market equilibrium using the results of demand and production cost estimation: $\hat{d}_{mt}$ and $\widehat{mc}_{it}$.[9]

No Subsidies - Static plots predictions after adjusting the estimated demand states, $\hat{d}_{mt}$, to reflect the removal of subsidies: $\hat{d}^{SCF}_{mt} = \hat{d}_{mt} - \hat{\alpha}_{sm}\text{arcsinh}\left(S_{mt}\right)$, where the superscript SCF stands for "static counterfactual." The estimated production costs remain unchanged. The product market equilibrium is computed using these inputs.

No Subsidies - Dynamic plots predictions after adjusting both the estimated demand states and the estimated production costs. The demand states are the same as in the static case: $\hat{d}^{DCF}_{mt} = \hat{d}^{SCF}_{mt}$, where the superscript DCF stands for "dynamic counterfactual." I adjust the estimated production costs to account for differences in endogenous innovation with and without subsidies in place. Based on the production cost estimates in Section 6, improvements to energy conversion efficiencies constituted 3% to 32% of total cost reductions over the period 2010-2015. Estimation of investment cost parameters relies on the most conservative of these production cost estimates. Furthermore, firms reduce their costs over time through other means which are likely to be endogenous rather than exogenous, but that do not show up in energy conversion efficiency improvements. Thus, using the model to compute equilibrium outcomes due to predicted changes in energy conversion efficiencies yields a lower bound on the total effects of subsidies on equilibrium prices and quantities over time. To account for this, I use the model's predictions about energy conversion efficiency improvements to conduct a back-of-the-envelope calculation that captures other forms of innovation as well. This provide a more comprehensive, albeit somewhat speculative, estimate of the effects of subsidies on innovation by firms.

I implement this back-of-the-envelope calculation by scaling the proportional improvements in cost-indexed energy conversion efficiencies to apply to the entirety of firms' production costs. Production cost predictions for each firm and time period are the product of two terms: $\widehat{mc}_{it} = \widehat{\eta}_{it}\widehat{w}_t$. Cost-indexed energy conversion efficiency predictions, $\widehat{\eta}_{it}$,

[9] An alternative approach would be to plot predictions from forward simulation of the counterfactual scenarios from the full model. The resulting plot would look very similar, except that it would smooth out the variation in Figure 5 that stems from unobserved demand shocks over time. This is because it numerically integrates over many shocks to the demand state transition process, whereas in reality only one shock is realized in each market and time period. Thus, I use the original estimates $\hat{d}_{mt}$ for realism.

are given by $\widehat{\eta}_{it} = \tilde{\eta}_{it}^{\hat{\beta}_1}$.[10] Common input price predictions, $\widehat{w}_t$, are recovered by taking the exponent of the time fixed effects from estimation of equation 4.[11] I use $\widehat{mc}_{it}$ to construct counterfactual cost estimates as follows:

$$\overline{\widehat{mc}}_t^{DCF} = \overline{\widehat{mc}}_1 - \left(\frac{\overline{\widehat{\eta}}_1 - \overline{\widehat{\eta}}_t^{DCF}}{\overline{\widehat{\eta}}_1 - \overline{\widehat{\eta}}_t^{BL}} \right) \times \left(\overline{\widehat{mc}}_1 - \overline{\widehat{mc}}_t \right),$$

where overlines denote means (e.g., $\overline{\widehat{mc}}_t = \sum_i \widehat{mc}_{it}/I$).[12] In words, this formula scales the estimated reduction in industry average marginal cost at time t relative to the initial period $\left(\overline{\widehat{mc}}_1 - \overline{\widehat{mc}}_t \right)$ by the percent reduction in cost-indexed energy conversion efficiencies predicted by forward simulation of the counterfactual solution method for the dynamic counterfactual without subsidies relative to the same predictions from the baseline scenario with subsidies: $\left(\overline{\widehat{\eta}}_1 - \overline{\widehat{\eta}}_t^{DCF} \Big/ \overline{\widehat{\eta}}_1 - \overline{\widehat{\eta}}_t^{BL} \right)$. If solving the dynamic model for the *Baseline* and *No Subsidy* counterfactuals predicted that energy conversion efficiency improvements would have been the same with and without subsidies in place (i.e., $\overline{\widehat{\eta}}_t^{DCF} = \overline{\widehat{\eta}}_t^{BL}$), the scaling factor would be one, and the resulting marginal cost prediction would be $\overline{\widehat{mc}}_t^{DCF} = \overline{\widehat{mc}}_t$. This is the case for $t = 1$, since the counterfactual solution method is initialized at this time period, so that $\overline{\widehat{\eta}}_1^{DCF} = \overline{\widehat{\eta}}_1^{BL} = \overline{\widehat{\eta}}_1$ by construction. If instead solving the dynamic model predicted that energy conversion efficiency improvements would not have occurred at all without subsidies in place (i.e., $\overline{\widehat{\eta}}_t^{DCF} = \overline{\widehat{\eta}}_1^{DCF} = \overline{\widehat{\eta}}_1$), the scaling factor would be zero, and the resulting marginal cost prediction would be $\overline{\widehat{mc}}_t^{DCF} = \overline{\widehat{mc}}_1$. In practice, the results of solving the model for counterfactuals lie in between these two extremes. Finally, $\hat{d}_{mt}^{DCF}$ and $\overline{\widehat{mc}}_t^{DCF}$ are used to compute the product market equilibrium.

[10] $\tilde{\eta}$ is observed energy conversion efficiency, and $\hat{\beta}_1$ is the estimated relationship between conversion efficiency and production cost. See Section 5 for the definition and rationale for using cost-indexed energy conversion efficiencies for estimation of the dynamic model.

[11] Under the specification reported in column 3 of Table 3 that is used for dynamic estimation, equation 4 simplifies to $\ln(mc_{imt}) = \beta_1 \ln(\tilde{\eta}_{it}) + \delta_t + \varepsilon_{imt}^S$. Thus, $\widehat{w}_t = \exp(\hat{\delta}_t)$.

[12] Means are used for convenience because they are a sufficient statistic for the full distribution of costs in the product market model.

References

Benitez-Silva, H., G. Hall, G. J. Hitsch, G. Pauletto, and J. Rust (2000). A Comparison of Discrete and Parametric Approximation Methods for Continuous-State Dynamic Programming Problems. *Manuscript, Yale University*.

Bergstrom, T. C. and H. R. Varian (1985). When Are Nash Equilibria Independent of the Distribution of Agents' Characteristics? *The Review of Economic Studies 52*(4), 715–718.

Gillingham, K. and T. Tsvetanov (2019). Hurdles and steps: Estimating demand for solar photovoltaics. *Quantitative Economics 10*(1), 275–310.

Ifrach, B. and G. Y. Weintraub (2017). A Framework for Dynamic Oligopoly in Concentrated Industries. *The Review of Economic Studies 84*(3), 1106–1150.

Muller, N. Z., R. Mendelsohn, and W. Nordhaus (2011). Environmental Accounting for Pollution in the United States Economy. *The American Economic Review 101*(5), 1649–1675.

Sweeting, A. (2013). Dynamic Product Positioning in Differentiated Product Markets: The Effect of Fees for Musical Performance Rights on the Commercial Radio Industry. *Econometrica 81*(5), 1763–1803.