Optimal Capacity Investments in Oligopolistic Electricity Markets with Learning-by-Doing

Felipe Verastegui
Industrial Engineering & Operations Research
Columbia University
New York, United States
f.verastegui@columbia.edu

Anton Pichler
Institute for Transport and Logistics
Vienna University of Economics and Business
Vienna, Austria
anton.pichler@wu.ac.at

***Abstract*—We study how learning-by-doing affects capacity investment and production decisions in oligopolistic electricity markets. We extend a conjectural variations model to a dynamic setting which includes a learning-by-doing mechanism: current production reduces future investment costs. In this setting, we characterize prices and capacities in closed form and show that, in the competitive equilibrium, market power interacts with learning asymmetrically: post-learning market power is necessary for firms to internalize learning rents, while pre-learning market power governs the strength of the learning response and the reduction of future investment costs. Learning-by-doing amplifies oligopolistic price premia and welfare gaps persist even as conjectural variations vanish, unless learning effects are negligible. In contrast, strategic capacity withholding remains independent of experience and learning. Our results inform policy on balancing market power in contexts where competitive equilibria underinvest in technology experience.**

***Index Terms*—component, formatting, style, styling, insert**

I. INTRODUCTION

Electricity markets are undergoing a profound transformation driven by rapid cost declines in clean energy technologies. Over the past two decades, technologies such as solar photovoltaics, wind turbines, and batteries have experienced steep cost decreases [3], [8], reshaping investment patterns and competitive dynamics across power systems globally. *Learning-by-doing* – the process through which experience reduces costs or improves performance and productivity – has been documented in many sectors, including electricity markets [1], [6], [9]. Endogenous cost declines can influence firms' production and capacity decisions by creating intertemporal incentives: today's production decisions not only affect current profits but also impact future investment costs. Understanding how such changes interplay with market power is essential for evaluating investment efficiency and policy design. Yet how learning-by-doing interacts with strategic behavior and affects social welfare in electricity markets remains an open question.

In this paper, we address this question by building on theoretical models of market outcomes in the context of learning-by-doing [2], [5]. First, we introduce a learning mechanism in a canonical model of oligopolistic competition in electricity markets. We extend the *conjectural variations* model in [11] to a dynamic setting, where we allow for future investment cost reductions to be endogenously determined by current production decisions. This provides a framework to study the interaction of oligopolistic competition with learning-by-doing effects. We characterize equilibrium prices and capacity investment decisions before and after learning and harness these results to identify equilibrium distortions arising from the dynamic feedback between production and investment costs for different levels of market power – ranging from perfect competition to collusion (tit-for-tat).

We find that learning-by-doing distorts oligopolistic behavior in equilibrium through intertemporal incentives. When there is perfect competition after learning, the intertemporal inefficiencies are maximal, and firms fail to realize any welfare improvements from lower investment costs. Market power *before* the learning stage plays an important yet different role: because experience is driven by production, in our model oligopolistic firms effectively compute a *conjectured learning response* – a belief on how other forms competitive response affects learning-by-doing.

While a first large wave of electricity sector reforms focused on deregulating vertically integrated monopolies to enhance market competition [7], a more recent focus have been policies to support the deployment of clean energy [4]. Our results aim at the intersection of these processes, and inform the discussions on optimal market-power mitigation, as well as the design of instruments to support novel energy technologies in competitive electricity markets.

II. MODEL

We study a finitely repeated version of the electricity market presented in [11] with learning-by-doing. Firms $i \in N = \{1, 2, \dots, n\}$ privately choose capacity investments x_{it} and electricity production q_{it} with $0 \leq q_{it} \leq x_{it} \, \forall \, t \in \{1, 2\}$. Capacity does not carry over from one period to the next. This instant obsolescence assumption makes capacity costly at $t = 2$, emphasizing the theoretical impact of learning-by-doing. Vectors of capacity investments and electricity production are $\mathbf{x}_t$ and $\mathbf{q}_t$, both $\in \mathbb{R}^n_+$. Aggregate capacity investment and production are denoted by $Q_t := \mathbf{1^T q_t}$ and $X_t := \mathbf{1^T x_t}$. All firms use the same technology with a linear per unit cost

979-8-3315-8138-1

of $\delta > 0$. Production at each $t \in \{1, 2\}$ is sold at a market clearing price p_t defined as

$$p_t = p(q_{it}, q_{-it}) := \frac{D - Q_t}{\alpha}, \tag{1}$$

where $D > 0$ and $\alpha > 0$ respectively denote the inverse demand function intercept and slope, both which are assumed to be fixed and known in advance. In our model, unlike [11], capacity has a time-varying per unit cost, β_t, due to learning-by-doing. At $t = 1$, the capacity cost is $\beta_1 = \beta > 0$ for some known constant. At $t = 2$, capacity cost is given by

$$\beta_2(z) := \beta - h(z), \tag{2}$$

with *experience* defined as first-stage production $z := Q_1$ and $h : \mathbb{R}^N \mapsto \mathbb{R}^+$ a *learning-by-doing function* that satisfies a bounded and decreasing returns-to-scale condition:

Assumption 1.1. *$h(\cdot)$ is non-decreasing and concave with $h(\mathbf{0}) = 0$ and $0 \leq h(z) < \beta \, \forall z \in \mathbb{R}^N_+$.* [1]

Firm i's utility at each stage $t \in \{1, 2\}$ is the sum of operating revenue minus production and investment costs:

$$u_{i1}(q_{i1}, q_{-i1}, x_{i1}) := \kappa(p_1 - \delta)q_{i1} - \beta_1 x_{i1}, \tag{3}$$
$$u_{i2}(q_{i2}, q_{-i2}, x_{i2}, \mathbf{q}_1) := \kappa(p_2 - \delta)q_{i2} - \beta_2(\mathbf{1}^T\mathbf{q}_1)x_{i2}, \tag{4}$$

where $\kappa > 0$ is an operational scaling constant. Consumer surplus cs_t at each stage $t \in \{1, 2\}$ is computed via the integral of the inverse demand function minus the equilibrium price scaled by κ:

$$cs_t(\mathbf{q}_t) := \kappa \int_0^{\mathbf{1^T q_t}} \Big[\frac{D - s}{\alpha} - \frac{D - \mathbf{1^T q_t}}{\alpha}\Big] \, \mathrm{d}s = \frac{\kappa Q_t^2}{2\alpha}. \tag{5}$$

We assume a common discount factor $\rho \in (0, 1)$ for all agents. Net present utility of firm i and consumer surplus given investment and production decisions, $\mathbf{x} := \{\mathbf{x}_1, \mathbf{x}_2\}$ and $\mathbf{q} := \{\mathbf{q}_1, \mathbf{q}_2\}$, can be written as:

$$u_i(\mathbf{q}, \mathbf{x}) := u_{i1}(q_{i1}, q_{-i1}, x_{i1}) + \rho u_{i2}(q_{i2}, q_{-i2}, x_{i2}, \mathbf{q}_1), \tag{6}$$
$$cs(\mathbf{q}) := cs_1(\mathbf{q}_1) + \rho cs_2(\mathbf{q}_2). \tag{7}$$

We define utilitarian social welfare as the sum of firm utility and consumer surplus for each stage. With this we can define before and after-learning stage welfare as:

$$w_1(\mathbf{q}_1, \mathbf{x}_2) := \sum_{i \in N} u_{i1}(q_{i1}, q_{-i1}x_{i1}) + cs_1(\mathbf{q}_1), \tag{8}$$
$$w_2(\mathbf{q}_2, \mathbf{x}_2, z) := \sum_{i \in N} u_{i2}(q_{i2}, q_{-i2}x_{i2}, \mathbf{q_1}) + cs_2(\mathbf{q}_2). \tag{9}$$

The net present social welfare given investment and production decisions then corresponds to:

$$w(\mathbf{q}, \mathbf{x}) := \sum_{i=1}^N u_i(\mathbf{q}, \mathbf{x}) + cs(\mathbf{q}). \tag{10}$$

[1] The strict upper bound on $h(\cdot)$ guarantees a non-negative investment cost.

A. Conjectural Variations: Price and Learning Responses

In the competitive outcome, we follow [11] and adopt a *conjectural variations* framework. Under this framework, firms hold symmetric beliefs about how other firms adjust their production in response to their own actions. That is:

$$\frac{\partial q_{it}}{\partial q_{jt}} = \phi_t \, \forall i, j \in [n] : i \neq j \tag{11}$$

for some constant ϕ_t which can be time-dependent.[2] By the chain rule, this modeling assumption immediately implies a *conjectured price response*, that is, a belief on how price is affected given a change in firm i's production. It follows that:

$$\begin{aligned}\frac{\partial p(q_{it}, q_{-it})}{\partial q_{it}} &= -\alpha^{-1}(1 - \sum_{j \neq i} \frac{\partial q_{jt}}{\partial q_{it}}), \\ &= -\alpha^{-1}(1 + (n-1)\phi_t). \end{aligned} \tag{12}$$

Unlike [11], in our setting the conjectural variations framework also implies a *conjectured learning response*, i.e., a belief on how future investment costs are affected given a change in firm i's production. Let $h_z(z) := \frac{\mathrm{d}h(z)}{\mathrm{dz}}$. By the chain rule:

$$\frac{\partial h(z(q_t))}{\partial q_{it}} = (1 + (n-1)\phi_t)h_z(z(q_t)). \tag{13}$$

It is useful to write Equations (11)–(13) using firms' conjectured price response. Let $\theta_t := \alpha^{-1}(1 + (n-1)\phi_t)$:

$$\frac{\partial p(q_{it}, q_{-it})}{\partial q_{it}} = -\theta_t, \tag{14}$$
$$\frac{\partial z(q_t)}{\partial q_{it}} = \alpha \, \theta_t, \tag{15}$$
$$\frac{\partial h(z(q_t))}{\partial q_{it}} = h_z(z(q_t)) \, \alpha \, \theta_t. \tag{16}$$

In a perfectly competitive market, one firms' change in production is absorbed by all remaining firms $\phi_t = \frac{-1}{n-1}$, making the right-side of Equations (14)–(16) zero. In other words, under perfect competition, a single firms' actions do not affect equilibrium prices *nor learning-by-doing outcomes*. We harness this conjectural variations framework to then model different levels of oligopolistic competition, adjusting the value of ϕ_t to produce non-zero values for θ_t. For more discussion on this point, we refer the reader to [11].

III. Social Planner's Solution

First, we study the model from the perspective of a utilitarian social planner. The planner's problem consists on choosing investment and production schedules $\mathbf{q}^*, \mathbf{x}^*$ which maximize welfare subject to feasibility constraints:

$$\begin{aligned} \text{(SP):} \; \max_{\mathbf{q}, \mathbf{x}} \quad & w(\mathbf{q}, \mathbf{x}) \\ \text{s.t.} \; & 0 \leq \mathbf{q} \leq \mathbf{x}. \end{aligned}$$

[2] This time dependency will play an important role in the subsequent analysis. Learning-by-doing introduces an asymmetry in the role of market power before and *after* learning on incentives and equilibrium outcomes.

979-8-3315-8138-1

We proceed via backward induction. Due to the finitely repeated nature of the game, in the second stage, there are no social benefits from learning-by-doing. Thus, after the learning stage, the planner takes $\mathbf{q_1}$ as given and solves an optimization problem that pins down the socially efficient capacity and production as a function of first-stage experience z.

Proposition 1. *The socially optimal capacity investments, production, and market clearing price at $t = 2$ are*

$$q^*(z) = \frac{D\kappa - \alpha(\delta\kappa + \beta)}{\kappa} + \alpha\frac{h(z)}{\kappa},$$
$$p^*(z) = \delta + \frac{\beta}{\kappa} - \frac{h(z)}{\kappa},$$

and $x^*(z) = q^*(z)$.

Proof. Follows from fixing $\mathbf{q_1}$ and taking first order conditions over $\max_{\mathbf{q_2},\mathbf{x_2}} w_2(\mathbf{q_2}, \mathbf{x_2}, z)$ with $0 \le \mathbf{q_2} \le \mathbf{x_2}$. □

From Proposition 1 we observe that the socially optimal capacity investments and production are non-decreasing and concave in first-stage experience z. In the socially optimal solution, the learning-by-doing mechanism lowers the market clearing price. This price reduction then leads to higher second-stage consumption and higher consumer surplus.

Because the planner's solution is unique for any level of experience z, we can write an indirect welfare function by combining Proposition 1 and Equation (9):

$$\hat{w}_2(z) := \max_{\mathbf{q_2},\mathbf{x_2}} \{w_2(\mathbf{q_2}, \mathbf{x_2}, z) \mid \mathbf{0} \le \mathbf{q_2} \le \mathbf{x_2}\}.$$

This enables recasting SP using only first-stage variables:

Proposition 2. *The solution to the problem*

$$\text{(SP-II):} \quad \max_{\mathbf{q_1},\mathbf{x_1},z} w_1(\mathbf{q_1}, \mathbf{x_1}) + \rho\hat{w}_2(z)$$
$$\text{s.t.} \quad 0 \le \mathbf{q_1} \le \mathbf{x_1},$$
$$z = \mathbf{1^T q_1}.$$

optimally solves (SP).

Proof. Follows from uniqueness of $q^*_{sp}(z), x^*_{sp}(z)$ and the SPE condition. □

The reformulated problem SP-II explicitly portrays the intertemporal trade-off faced by the social planner under learning-by-doing. Boosting first-stage production $\mathbf{1^T q_1}$ yields higher surplus after learning, and this effect enters the first-stage decisions via the discounted indirect welfare function. We now provide a formal characterization of how future welfare changes with first-stage experience.

Proposition 3. *The rate of change of the socially optimal welfare in $t = 2$ with respect to z can be written as*

$$\frac{\mathrm{d}\hat{w}_2(z)}{\mathrm{d}z} = q^*(z)\, h_z(z)$$

Proof. Follows from replacing $q^*(z)$ and $x^*(z)$ in $w_2(\mathbf{q_2}, \mathbf{x_2}, z)$ and taking a derivative over z. □

Proposition 3 shows that future welfare is increasing in experience but eventually faces decreasing returns to scale from learning effects due to the concavity condition of Assumption 1.1. This implies that although the planner does face a dynamic incentive to boost present aggregate production, this incentive is bounded. This, in turn, enables a characterization of the solution to SP-II using a fixed-point equation as follows.

Proposition 4. *The solution to (SP) can be characterized as the solution to the fixed-point equation $z^* = T(z^*)$ where:*

$$T(z) = \underbrace{q^*(0)}_{\text{Social Planner's Base Capacity}} + \underbrace{\rho\frac{\alpha}{\kappa}q^*(z)\, h_z(z)}_{\text{Social Planner's Learning Incentive}}.$$

This fixed-point equation has a unique solution, which we denote by z^*.

Proof. The fixed-point equation follows from taking first order conditions over (SP-II). Existence and uniqueness follow from the concavity of the learning-by-doing function $h(\cdot)$. □

Proposition 4 provides a characterization for the socially optimal first-stage experience z^* under *any* learning-by-doing function $h(\cdot)$ combining a base capacity – the myopic optimum without learning – with an upward adjustment term from learning-by-doing incentives. An impatient planner ($\rho \leftarrow 0$) chooses optimal capacity investments of $q^*(0)$ in the first stage — the no-learning outcome. As future welfare becomes more important (the planner becomes more patient), the role of the upward distortion increases proportionally to the marginal social benefits from learning-by-doing. We harness this characterization as a benchmark to analyze the effects of oligopolistic competition on equilibrium outcomes.

IV. Competitive Equilibrium

Second, we characterize competitive equilibrium outcomes in an oligopolistic electricity market with learning-by-doing. We focus on subgame perfect equilibria (SPE) of the finitely repeated game. SPE is a refinement of the more general Nash Equilibrium space in which firms play equilibrium strategies at every stage of the repeated game. [3] At $t = 1$, firms choose capacity investments and production *schedules* to maximize their net private utility while best responding to each other's actions. This requires that all firms $i \in N$ simultaneously solve a best response problem:

$$(\text{BR})_i : \quad \max_{q_{i1},x_{i1},q_{i2},x_{i2}} u_{i1}(q_{i1}, q_{-i1}, x_{i1}) + \rho u_{i2}(q_{i2}, q_{-i2}, x_{i2}, \mathbf{q_1}) \tag{17}$$
$$\text{s.t. } (q_{i1}, x_{i1}) \in \arg\max_{q_{i1},x_{i1}} u_{i1}(q_{i1}, q_{-i1}, x_{i1}), \tag{18}$$
$$(q_{i2}, x_{i2}) \in \arg\max_{q_{i2},x_{i2}} u_{i2}(q_{i2}, q_{-i2}, x_{i2}, \mathbf{q_1}), \tag{19}$$
$$0 \le q_{i1} \le x_{i1}, \tag{20}$$
$$0 \le q_{i2} \le x_{i2}. \tag{21}$$

[3] Formally, player's equilibrium strategies are required to form a NE of the continuation game at each node of the game tree. This solution concept coincides with the one used in other learning-by-doing models such as [5].

979-8-3315-8138-1

In an SPE, firms' actions at $t = 2$ must form a Nash Equilibrium for any $t = 1$ outcome. Because of this, the equilibrium at $t = 2$ is unique for any first-stage experience z and it coincides with the one derived by [11] in a model with no learning-by-doing.

Proposition 5. *Let $p^{max} := \frac{D}{\alpha}$ be the maximum market price. Competitive equilibrium outcomes after learning are given by:*

$$q_{i2}^*(z) = \big(1 - \psi(\theta_2)\big)\,\frac{q_{sp}^*(z)}{n},$$
$$p_{eq}^*(z) = \big(1 - \psi(\theta_2)\big)\,p_{sp}^*(z) + \psi(\theta_2)\,p_{max},$$

with $x_{i2}^*(z) = q_{i2}^*(z)$ and $\psi(\theta) := \frac{\alpha\theta}{n+\alpha\theta}$ a market power weighting function.

Proof. Follows from fixing $\mathbf{q}_1$ and simultaneously taking first order conditions over (19) subject to (21) for all $i \in N$. $\square$

Proposition 5 shows that, after the learning stage, firms strategically withhold capacity and production compared to a social planner. This distorts the equilibrium market price upwards due to oligopolistic behavior. The learning mechanism still lowers the market clearing price, but to a lesser extent than in the planner's solution — here, the equilibrium price is a convex combination of the socially optimal price and the maximum price, with the maximum price weight given by the market-power weighting function $\psi(\cdot)$, increasing and concave in θ. Since Proposition 5 uniquely identifies the equilibrium outcome, we can write the indirect private utility function:

$$\begin{aligned}\hat{u}_{i2}(z) &= \max_{q_{i2},x_{i2}} u_2(q,x|z)\\ &= u_{i2}(q_{i2}^*(z), q_{-i2}^*(z), x_{i2}^*(z), z)\\ &= \kappa(p_{eq}^*(z) - \delta)q_{i2}^*(z) - \beta_2(z)x_{i2}^*(z).\end{aligned}$$

This allows recasting BR_i using only first stage variables.

Proposition 6. *A collection of solutions to the problems*

$$\begin{aligned}\forall i:\ \textit{(BR-II)}_i =: &\max_{q_{i1}} u_{i1}(q_{i1}, q_{-i1}, x_{i1}) + \rho\hat{u}_{i2}(z(q_1))\\ &\textit{s.t.}\ \ 0 \le q_{i1} \le x_{i1},\end{aligned}$$

solves (BR)$_i$ to optimality for all $i \in N$.

Proof. Follows from uniqueness of $q_{i2}^*(z), x_{i2}^*(z)\,\forall\, i \in N$ and the SPE condition. $\square$

The reformulated problem BR-II explicitly captures the intertemporal trade-off from the learning mechanism, now faced by the private firms. Increasing q_{i1} may enable larger private rents in $t = 2$, affecting capacity and production decisions at $t = 1$. We now provide a formal characterization of how future private utility changes with first-stage experience.

Proposition 7. *The rate of change of private utility at $t = 2$ with respect to z can be written as*

$$\frac{\mathrm{d}\hat{u}_{i2}(z)}{\mathrm{d}z} = 2\big(\psi(\theta_2) - \psi(\theta_2)^2\big)\,\frac{q^*(z)h_z(z)}{n} \tag{22}$$

The rate of change is maximal at $\phi_2 = 1$ or $\theta_2 = \frac{n}{\alpha}$. [4]

[4]This value of ϕ corresponds to collusive behavior or tit-for-tat.

Proof. Follows from replacing the expressions in $q_{i2}^*(z)$ and $p_{eq}^*(z)$ in $\hat{u}_{i2}(z)$ and taking a derivative wrt. z. $\square$

Proposition 7 shows that while firms do internalize some of the benefits of learning-by-doing in the first stage, they do so to a much lesser extent than the social planner (cf. Proposition 3) introducing a dynamic inefficiency. In fact, when market-power is maximal ($\phi_2 = 1$) it holds that $n\,\frac{\mathrm{d}\hat{u}_{i2}(z)}{\mathrm{d}z} = \frac{1}{2}\frac{\mathrm{d}\hat{w}_2(z)}{\mathrm{d}z}$ — even under very favorable future market conditions, firms will not internalize more than half of the socially optimal intertemporal incentives. Finally, because the dependency of the rate of change on first-stage experience again satisfies eventually decreasing returns to scale, we can characterize a solution to BR-II using a fixed-point equation as follows.

Proposition 8. *The solution to $(BR)_i$ can be characterized via the solution to the fixed-point equation $z^* = T_{eq}(z^*)$ where:*

$$\begin{aligned}T_{eq}(z) = &\underbrace{(1 - \psi(\theta_1))\,q^*(0)}_{\textit{Oligopoly Base Capacity}}\\ &+ \underbrace{\psi(\theta_1)\,2\big(\psi(\theta_2) - \psi(\theta_2)^2\big)\ \rho\frac{\alpha}{\kappa}q^*(z)h_z(z)}_{\textit{Oligopoly Learning Incentive}}\end{aligned}$$

The fixed-point equation has a unique solution, which we denote as z_{eq}^*. Firm-level optimal capacity investments are $q_{i1}^* = \frac{z_{eq}^*}{n}\ \forall i \in N$

Proof. The fixed-point equation follows from taking first order conditions over (BR-II). Existence and uniqueness follow from the concavity of the learning-by-doing function $h(\cdot)$. $\square$

Proposition 8 gives a characterization for the competitive equilibrium first-stage experience z_{eq}^* under oligopolistic competition. In equilibrium, first-stage market power reduces the base capacity but increases the learning incentive via the *conjectured learning response*. This trade-off is further modulated by future market power via the dynamic incentive inefficiency verified in Proposition 7.

V. The Effects of Dynamic Strategic Behavior under Learning-by-Doing

We discuss the effects of present and future oligopolistic competition on capacity investments and prices under learning-by-doing. First, we study oligopolistic rents after learning.

Theorem 9 (Oligopoly Rents under Learning-By-Doing)**.** *Second-stage oligopolistic rents are increasing in first-stage experience.*

Proof. Follows from defining oligopoly rents as $\Pi(z) := \frac{q_{eq}^*(z)(p_{eq}^*(z)-\delta)}{q^*(z)(p^*(z)-\delta)}$ and taking a derivative over z. $\square$

Under oligopolistic competition, strategic behavior creates capacity withholding and introduces a price mark-up. The introduction of a learning-by-doing mechanism amplifies this inefficiency in the second-stage. For any non-trivial level of first-stage experience, there is a decrease in future capacity investment costs. This decrease results in larger future oligopolistic rents, as the lower investment cost gives firms a

979-8-3315-8138-1

better position to exercise any available market power, which is the intuition behind the result in Theorem 9.

Next, we discuss the dynamic incentive misalignment:

Theorem 10 (Incentives from Learning-By-Doing)**.** *There is a dynamic incentive distortion which is decreasing and convex in future market power:*

$$\Delta du_2(z) = 1 - 2(\psi(\theta_2) - \psi(\theta_2)^2),$$

which is independent of the learning-curve specification.

Proof. Follows from defining $\Delta du_2 := \frac{\mathrm{d}\hat{w}_2(z)/\mathrm{d}z - \sum_i \mathrm{d}\hat{u}_{i2}(z)/\mathrm{d}z}{\mathrm{d}\hat{w}_2(z)/\mathrm{d}z}$ and Propositions 3 and 7. □

A key takeaway of Theorem 10 is that, in this model, the anticipation of future market power *alleviates* the dynamic incentive distortion introduced by strategic behavior under learning-by-doing. In fact, Theorem 10 implies that under perfect competition, the dynamic incentive distortion is maximal at $\Delta du_2 = 1$. Recall that the distortion is minimal under second-stage collusion, $\theta_2 = \frac{n}{\alpha}$, which gives $\Delta du_2 = \frac{1}{2}$. Finally, although the distortion is independent of the learning-curve specification and first-stage experience, it is not independent of the *learning-by-doing channel*: without it, the denominator becomes null and the ratio is undefined.

Finally, we characterize the gap in first-stage experience:

Theorem 11 (Capacity Gap under Learning-By-Doing)**.** *First-stage experience gap can be decomposed into ex-ante capacity withholding and a market-attenuated learning incentive:*

$$\Delta T(z) = \psi(\theta_1)q^*(0) + \xi(\theta_1, \theta_2)\rho\frac{\alpha}{\kappa}\, q^*(z)\, h_z(z)$$

where $\xi(\theta_1, \theta_2) := 1 - 2\,\psi(\theta_1)\,(\psi(\theta_2) - \psi(\theta_2)^2)$*. Pre-learning market power increases ex-ante capacity withholding but decreases the attenuated learning incentive gap.*

Proof. Follows from defining $\Delta T(z) := T_{sp}(z) - T_{eq}$ and Propositions 4 and 8. □

Note from Theorem 11 that pre-learning market power imposes an endogenous trade-off between ex-ante withholding and the learning-incentive gap. A welfare-maximizing regulator may therefore tolerate some market power before learning (interior value for θ_1), provided post-learning markets are not perfectly competitive. This is never the case, however, if future market conditions are expected to be perfectly competitive. Second, energy-market transformation programs should target closing the experience gap, and should be carefully designed to account for pre-learning and post-learning market power. Third, the learning-by-doing function plays an important role in the experience gap, with relevant welfare implications. In the presence of rapid-learning technologies – that is, large values of $h_z(z)$ – a careful approach to market-power regulation and transformation measures becomes even more critical.

VI. Discussion

In a two-stage conjectural variations framework, we find that the oligopolistic equilibrium is distorted via the intertemporal incentives created by technological cost improvements. Market power after the learning stage creates local inefficiencies, but is a necessary condition for private firms to ex-ante internalize some of the benefits from cost declines. Because the decrease in future investment cost is driven by past production, market power before the learning stage is critical to support firms' *conjectured learning response* – a belief on how other firms strategically contribute to the learning-by-doing process.

Our model is stylized, and as such, it fails to capture many relevant features in power systems, such as congestion and unit commitment decisions. Yet it provides initial predictions which can be empirically tested to support further work.

First, post-learning oligopolistic rents *increase* with first-stage experience, due to lower investment costs giving firms a better position to exploit consumers' willingness to pay in oligopolistic equilibria. An empirical comparison of electricity markets with similar concentration but different clean energy penetration can provide insights on this mechanism.

Second, the dynamic incentive distortion *decreases* in post-learning market-power and is independent of the learning curve specification. Empirically quantifying the strength of intertemporal incentives from learning-by-doing in electricity investments would be of interest.

Third, a benevolent regulator might optimally choose to allow a *strictly positive* degree of market-power, trading-off the effect of the learning-incentive gap with ex-ante capacity withholding. Empirical work by [10] provides an initial calibration of learning-by-doing curves for energy technologies based on historical data. The rapid cost declines in the energy sector present a valuable opportunity to study regulatory implications of the interaction between innovation and competition.

References

[1] L. Argote and D. Epple, "Learning curves in manufacturing," *Science*, vol. 247, no. 4945, pp. 920–924, 1990.

[2] K. J. Arrow, "The economic implications of learning by doing," *Rev. Econ. Stud.*, vol. 29, no. 3, pp. 155–173, 1962.

[3] M. Bolinger, R. Wiser, and E. O'Shaughnessy, "Levelized cost-based learning analysis of utility-scale wind and solar in the United States," *iScience*, vol. 25, no. 6, 2022.

[4] R. Duke and D. Kammen, "The Economics of Energy Market Transformation Programs", *The Energy Journal*, vol.20, no.4, pp.15-64, 1999.

[5] D. Fudenberg and J. Tirole, "Learning-by-doing and market performance," *Bell J. Econ.*, vol. 14, no. 2, pp. 522–530, 1983.

[6] A. Grübler, N. Nakićenović, and D. G. Victor, "Dynamics of energy technologies and global change," *Energy Policy*, vol. 27, no. 5, pp. 247–280, 1999.

[7] P. Joskow, "Deregulating and regulatory reform in the US electric power sector", *MIT Center for Energy and Environmental Policy Research*, 2000.

[8] M. M. Klemun, G. Kavlak, J. McNerney, and J. E. Trancik, "Mechanisms of hardware and soft technology evolution and the implications for solar energy cost trends," *Nature Energy*, vol. 8, no. 8, pp. 827–838, 2023.

[9] F. Lafond *et al.*, "How well do experience curves predict technological progress? A method for making distributional forecasts," *Technol. Forecast. Soc. Change*, vol. 128, pp. 104–117, 2018.

[10] R. Way, M. C. Ives, P. Mealy, and J. D. Farmer, "Empirically grounded technology forecasts and the energy transition," *Joule*, vol. 6, no. 9, pp. 2057–2082, 2022.

[11] S. Wogrin, B. F. Hobbs, D. Ralph, E. Centeno, and J. Barquin, "Open versus closed loop capacity equilibria in electricity markets under perfect and oligopolistic competition," *Math. Program.*, vol. 140, pp. 295–322, 2013.

979-8-3315-8138-1