sustainability

MDPI

Article

Cost-Based Price Floors, Soft Budgets and Capacity Adjustment in China's Solar Photovoltaic Industry

Chao Cheng [1] and Jinlin Zhang [2,*]

[1] School of Finance, Zhongnan University of Economics and Law, Wuhan 430073, China
[2] International Business School, Hainan University, Haikou 570228, China
* Correspondence: zjlifb@163.com

Abstract

Persistent overcapacity in photovoltaic (PV) manufacturing raises a question for the economic sustainability of renewable energy supply chains: does protecting producers' margins help or hinder capacity adjustment? We evaluate cost-based price floors with an agent-based model of China's PV module industry that features adaptive pricing, market-share-driven investment, credit guarantees and local government rescues. The model is driven by observed polysilicon prices and reproduces eight of ten moments from public company data within a pre-registered tolerance, although it captures only part of the 2024 price decline. Policy effects are paired comparisons across 200 random seeds. A floor at each firm's own average total cost lowers utilisation by 5.8 percentage points and welfare by 1.1% over ten years and, in the main specification, slows exit even under hard budgets; its welfare loss grows when products are more substitutable and turns into a gain only when demand is much less price-sensitive, R&D responds strongly to profits or share pressure is very high. An industry-reference floor stabilises prices but cuts capacity exit by 75%; its 1.8% welfare gain falls to 0.5% when surviving firms are not valued as additional variety. An admission standard for new capacity raises utilisation and, while demand plateaus, welfare. Protecting margins is a poor substitute for policies that change the returns to expansion.

Keywords: photovoltaic manufacturing; overcapacity; price floors; soft budget constraints; industrial policy; agent-based modelling; economic sustainability

1. Introduction

Between 2023 and 2025, China's solar photovoltaic (PV) module industry went through one of the sharpest price wars in its history. For the three leading producers whose module revenue and shipments can be matched, the average realised module price fell by 32.8% in 2024 and by a further 16.4% in 2025. Domestic spot quotations for mainstream TOPCon modules fell by 44% on an annual-average basis in 2024. The two figures differ in scope: the realised price covers all module products and both domestic and export sales and includes deliveries under contracts signed before prices fell, whereas the spot quotation refers to a single product on the domestic market and is not weighted by transaction volume. Four of the six leading listed module producers reported net losses in 2024, and five did so in 2025. Utilisation of disclosed effective module capacity averaged about 66% in 2023–2025 among the four producers that report it. Chinese policy discourse describes this pattern as "involution-style" (neijuan shi) competition: firms keep cutting prices and adding capacity even though the industry as a whole loses money.

Check for updates

Academic Editor: Izabela Jonek-Kowalska

Received: 5 October 2026
Revised: 14 October 2026
Accepted: 21 October 2026
Published: 1 November 2026

Copyright: © 2026 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license.

For sustainability policy, the relevant question is whether a measure supports economically viable industrial adjustment while preserving the benefits of lower prices for buyers. We define the economic sustainability of PV manufacturing as the industry's ability to sustain production, investment and innovation without persistent losses, recurrent fiscal rescue or idle resources, while passing cost reductions on to the buyers who deploy PV. This dimension matters for the energy transition: excess capacity ties up capital, materials and energy in idle plants; price wars that bankrupt firms can cut the investment and R&D on which further cost reductions depend; and module prices feed into the cost of every new installation. Producer profitability alone cannot resolve the question: higher margins can support investment but can also retain inefficient capacity and transfer losses to consumers or public budgets. We therefore assess the economic dimension of PV manufacturing sustainability through capacity use, exit, real resource costs, fiscal exposure and the distribution of welfare. The model does not quantify avoided emissions or life-cycle environmental impacts; Section 7.3 discusses what the results imply for deployment and environmental outcomes.

The policy response has focused on prices. In July 2025, the Central Financial and Economic Affairs Commission called for governing disorderly low-price competition in accordance with the law and for the orderly exit of backward capacity [1]. Later that month, the National Development and Reform Commission published a draft amendment to the Price Law that would extend the prohibition of below-cost dumping [2]. We focus on these price rules because policymakers began legislating them in 2025 and because their effects on capacity are theoretically ambiguous. Capacity-side instruments serve as benchmarks against which the effects of these price floors are evaluated. Can a ban on below-cost sales end involution, and how does it interact with the soft budget constraints that keep loss-making firms alive?

Economic reasoning gives no clear answer. A price floor raises margins and may stop ruinous price cutting, but higher margins also weaken the pressure on high-cost firms to exit and may attract further investment. The answer may depend on the cost the floor refers to—a firm's own variable cost, its own average total cost or an industry reference cost—and on whether budgets are hard: if local governments and banks rescue distressed firms anyway, a floor may simply shift the cost of keeping them alive from taxpayers to buyers. These questions are hard to answer empirically, because the instruments are new, they interact, and their long-run counterfactuals are not observed.

We address them with an agent-based model (ABM) of the module industry. Heterogeneous firms set prices by adjusting markups, invest on the basis of extrapolative demand expectations and market-share targets subject to construction lags, and borrow under leverage limits that implicit local government guarantees relax. Local governments compete for growth with investment subsidies and may rescue distressed firms. Buyers choose among firms and an outside option under logit demand with explicit rationing. Unit costs combine a common silicon cost, driven by the observed 2020–2025 polysilicon price path, with firm-specific processing costs. We calibrate the model by the simulated method of moments [3] to ten moments from public company reports, eight of which it reproduces within a pre-registered tolerance (Section 4.3); the price decline of 2024 is reproduced only in part. Every policy scenario uses the same 200 random seeds, so all effects are paired differences.

We obtain four main results. First, a ban on selling below a firm's own average total cost does not clear excess capacity. Over the ten years following 2025Q3, it lowers capacity utilisation by 5.8 percentage points, slows exit, raises investment and lowers global welfare by 1.1%; at the share pressure of the main specification, it also roughly doubles the time the industry spends in involution. The welfare loss does not depend on how product variety is

valued and grows when products are more substitutable, although exit then rises rather than falls. Second, hard budgets do not change this conclusion. Under a fully hard budget, the floor still slows exit, because without rescues it becomes the main force keeping weak firms alive. A floor at a firm's own full cost thus acts as a soft budget constraint financed by buyers rather than by governments. Third, the cost basis matters. A variable cost floor, the closest analogue to a classic predatory pricing rule, almost never binds. A floor tied to an industry reference cost stabilises prices and cuts capacity exit by 75%; it raises measured welfare by 1.8% under logit demand, by 0.5% when additional surviving firms are not valued as additional variety, and lowers welfare when products are highly substitutable. Fourth, instruments that change the returns to expansion address the problem better than instruments that set prices: an admission standard for new capacity raises utilisation in every specification in which we examine it and raises welfare while demand plateaus, as in the baseline, but lowers it if demand grows as fast as industry projections suggest. Because about 40% of modules are exported, part of what the industry loses in a price war is a transfer to foreign buyers, and welfare rankings also depend on whether their surplus is counted.

Several of these findings confirm existing knowledge: price floors retain inefficient firms [4], below-cost sales laws soften competition [5], and capacity-side instruments outperform price controls. The paper adds four findings that are, to our knowledge, new. A price floor can recreate a soft budget through the market even when fiscal budgets are hard. The variable cost standard, which matches the conventional predatory pricing test, rarely binds in a downturn driven by fixed costs. An industry-reference floor restrains mainly efficient firms rather than weak ones. And the level of subsidies matters more for involution than regional competition for them. Two further results, the advantage of hardening budgets before introducing a floor and the welfare gain from a reference floor, rest on assumptions about R&D behaviour and the valuation of variety, and we present them as conditional.

The paper relates to three strands of literature. The first explains China's recurrent overcapacity by investment herding, local government competition and subsidies [6–9], including simulation models of overcapacity and of PV subsidy withdrawal [10,11]. We take these mechanisms as given and study the newest instrument, cost-based price floors. The second studies price floors and below-cost sales laws, mostly in retail and commodity markets [4,5,12,13]. We study a capital-intensive industry in which floors affect investment, exit and rescue as well as prices. The third concerns soft budget constraints and zombie firms [14–17]. Section 2.2 summarises how the model relates to these approaches. Methodologically, the welfare accounting separates transfers from resource costs, measures realised consumer surplus under rationing and reports domestic as well as global welfare, both with and without a valuation of product variety.

Section 2 reviews the policy background and related literature. Section 3 presents the model and Section 4 its calibration and validation. Section 5 reports the results and Section 6 the robustness checks. Section 7 discusses policy implications, limitations and implications for PV deployment, and Section 8 concludes.

2. Background and Related Literature

2.1. The Downturn and the Policy Response

The downturn combined a collapse in input costs with a capacity glut. In our quarterly index of polysilicon quotations, the price of the main input fell by about two-thirds between the first and the last quarter of 2023, removing a major support for module prices. Domestic spot quotations show module prices falling by 44% in 2024 and by 19% in 2025 on an annual-average basis (InfoLink weekly reports; see Section 4). Among leading listed producers,

losses became the norm, gross margins fell to single digits and utilisation of effective capacity dropped to about two-thirds. Yet expensed R&D intensity rose rather than fell, because revenue declined faster than R&D spending.

The policy response has three strands. The first restricts local support: the Fair Competition Review Regulation, effective August 2024, limits selective tax preferences and subsidies granted without legal basis [18]. The second targets pricing. The draft amendment to the Price Law released on 24 July 2025 keeps the existing intent condition, under which below-cost dumping is prohibited when it aims to squeeze out competitors or monopolise a market, and adds the act of forcing other operators to sell below cost [2]. The revised Anti-Unfair Competition Law of 2025 prohibits platforms from forcing their merchants to sell below cost [19]. The third strand concerns capacity: industry standards for new manufacturing capacity and mergers among producers.

Two features of this framework shape our design. First, no current rule fixes the accounting basis of "cost", and existing law contains intent conditions and exceptions. We therefore study three stylised floors—a firm's own unit variable cost, its own audited average total cost and an industry reference cost—as mechanism counterfactuals, not as implementations of a statute. Second, coordinated minimum prices raise competition law concerns of their own: Article 14(1) of the Price Law prohibits collusive price manipulation, and publishing industry cost information is different from imposing a uniform minimum price. Our industry-reference floor should therefore be read as an upper bound on what cost-information schemes could achieve.

2.2. Related Literature

A large literature explains recurrent overcapacity in China. Lin [6] and Lin et al. [7] attribute it to "wave-like" investment, in which many firms share an optimistic view of a promising industry and overshoot demand together. Geng et al. [8] and Jiang et al. [9] emphasise subsidies, regional competition and institutional distortions that lower firms' private cost of capacity, and promotion tournaments among local officials provide the incentives behind these distortions [20]. Corwin and Johnson [21] document the role of local governments in building China's PV industry, and Yu et al. [22], Chen and Wang [23], Ge et al. [24] and Liu and Lai [25] study the determinants of overcapacity and policies against it. Barwick et al. [26,27] show for shipbuilding how support for entry and capacity can expand output at the cost of fragmentation and low utilisation.

Soft budget constraints are a second root. The expectation of rescue weakens financial discipline [14,15,28], and in market economies the same mechanism appears as zombie lending [16]. For China, Shen and Chen [17] relate zombie firms to manufacturing overcapacity, Zhang et al. [29] and Chang et al. [30] link them to soft budgets and government intervention, and Dai et al. [31], Geng et al. [32] and Mao and Xu [33] document their costs for exit, upgrading and capacity utilisation. Fan et al. [34], Ma et al. [35], Wang and Qiao [36] and Xiang and Xiang [37] study what keeps zombie firms alive.

Price floors have been studied mainly in retail and commodity markets. Resale-below-cost laws can soften competition and raise prices [5], a price floor in gasoline retailing can let less efficient retailers survive [4], and Hernández and Cantillo-Cleves [13] develop a toolkit for setting and evaluating price floors. He and Westerhoff [12] show that price limiters can have counterintuitive effects in a behavioural commodity-market model, and Bortolomiol et al. [38] study price-based regulation of oligopolies under discrete-choice demand. This work does not address two features of the PV case: capacity is built with long lags and financed by credit, and loss-making firms may be rescued by governments.

Our model belongs to the family of agent-based industry models with adaptive firms, credit and policy interventions [39–41]. For Chinese overcapacity, Zhang and Dong [10]

build a neo-Schumpeterian evolutionary model with process innovation, construction lags and local intervention; Zhang et al. [11] use system dynamics to ask whether cancelling subsidies alleviates overcapacity in PV module manufacturing; and Zhang et al. [42], Chen et al. [43] and Wang et al. [44] study subsidies and overcapacity governance with evolutionary games and supply-chain models. Zuo et al. [45] simulate involution in social resource competition with LLM-assisted agents. We treat the formation of overcapacity as established and focus on cost-based floors, budget constraints and welfare incidence, following the validation practice reviewed by Fagiolo et al. [46] and describing the model with the Overview, Design concepts and Details (ODD) protocol [47].

Table 1 contrasts these approaches with ours. Models of overcapacity explain why capacity accumulates but do not contain price-floor instruments, whereas models of price floors abstract from capacity that is built with long lags, financed by credit and protected by rescues. The question studied here, whether a cost-based floor clears capacity when budgets are soft, requires both sets of mechanisms in one framework, together with welfare accounting that separates resource costs from transfers.

Table 1. Modelling approaches to overcapacity and price regulation compared with this study.

Approach and Studies	Main Focus	Firms and Capacity	Finance, Exit and Rescue	Price-Floor Instruments	Welfare and Data
Theory and econometrics of Chinese overcapacity [6–9,17]	Causes of overcapacity: herding, subsidies, regional competition, zombie firms	Industry or aggregate level	Soft budgets and zombie firms as drivers	Not studied	Econometric evidence; no counterfactual welfare
Evolutionary simulation [10]	Process innovation and construction cycles	Heterogeneous firms; construction lags	Local intervention	Not studied	Simulated overcapacity
System dynamics [11]	Subsidy cancellation in PV module manufacturing	Aggregate stocks and flows	Not firm-level	Not studied	Simulated capacity paths
Evolutionary games and supply-chain models [42–44]	Subsidies and governance strategies	Strategy populations or a few chains	Not modelled as credit	Not studied	Equilibrium strategies
Agent-based macro models [39–41]	Growth, cycles and macro policy	Heterogeneous firms; investment	Bank credit; bankruptcy	Not studied	Macro aggregates; stylised facts
Structural industry model [26,27]	Industrial policy in shipbuilding	Entry, exit and investment	Subsidies	Not studied	Estimated welfare effects
Price floors and below-cost laws [4,5,12,13,38]	Effects of price floors and limits	Retail, gasoline or commodity markets; capacity given	Not modelled	Below-cost laws, price floors, price limiters	Consumer surplus; theory or estimation
This study	Cost-based floors under soft budgets	26 heterogeneous firms; share-target investment with construction lags; endogenous entry and exit	Leverage limits, implicit guarantees, local rescues	Variable cost, own-cost and reference floors with imperfect enforcement	Resource accounting by actor; domestic and global; calibrated to public company data

3. Materials and Methods: The Agent-Based Model

The model is a quarterly, partial-equilibrium agent-based model of a module-manufacturing industry. It starts with 26 firms located in 20 regions; firms enter and exit endogenously. Demand is represented by 4096 fixed consumer draws. Model quarter 80 corresponds to 2023Q1 and quarter 90 to 2025Q3; each run lasts 160 quarters. Quantities and money are measured in normalised model units, not in gigawatts or yuan, because the model is calibrated to ratios (utilisation, margins, price changes and shares of loss-making firms) rather than to levels. As an indicative scale, simulated output in 2023 is 1846 units, against 499 GW of module production in China in that year [48], so one unit of output

corresponds to about 0.27 GW. We report all policy effects relative to the no-policy run; for money values, a useful benchmark is the discounted industry revenue of the no-policy run over the ten-year policy window, 7112 units. Appendix A gives an ODD summary [47] and Appendix B the parameters and notation.

Each quarter proceeds in eight steps: (1) completed construction enters production, while capital retires physically and depreciates on the books; (2) firms observe the common silicon cost and set prices, deciding whether to comply when a price floor is in force; (3) demand is realised, buyers choose and excess demand is rationed; (4) revenue, costs, R&D, interest, taxes and fines are booked; (5) firms service and raise debt, and distressed firms are rescued or exit; (6) surviving firms order new capacity and attempt to innovate; (7) entrants arrive; (8) accounts are audited and local governments update subsidies. All random draws are indexed by seed, quarter, firm and event, so that the same seed produces the same shocks in every policy scenario.

The recorded replication environment uses Python 3.12, NumPy 2.3.5, SciPy 1.17.0, Numba 0.68.0 and pandas 2.2.3. Model settings and random seeds are recorded in the configuration files of the replication package.

3.1. Demand

Potential demand follows a logistic path with a first-order autoregressive (AR(1)) shock, and the mean utility of buying from firm i depends on its price and its product quality q_i:

$$M_t = \frac{M_{\max}}{1 + \exp[-g_M(t - t_0)]} e^{\xi_t}, \qquad \xi_t = \chi\, \xi_{t-1} + \sigma_\xi z_t, \tag{1}$$

$$V_{it} = a - \alpha p_{it} + \beta \ln q_i. \tag{2}$$

Buyer j chooses the option with the highest value of $V_{it} + \sigma\varepsilon_{ij}$, where the outside option is worth $\sigma\varepsilon_{0j}$ and the ε terms are type-I extreme-value draws that stay fixed across quarters and scenarios. Without capacity limits, this gives logit market shares $\exp(V_{it}/\sigma)/[1 + \sum_k \exp(V_{kt}/\sigma)]$. When a firm's demand exceeds its capacity, buyers are served in random order, and those not served choose again among options with spare capacity. Consumer surplus is the realised utility gain of buyers who actually purchase, divided by α.

Because each firm enters the choice set as a distinct option, a larger number of surviving firms raises buyer value even at unchanged prices. Logit demand is equivalent to the demand of a representative consumer with a taste for variety [49], an assumption that is questionable for crystalline-silicon modules, which buyers treat as close substitutes. We therefore report every welfare result also under a variety-neutral specification, in which the mean utility of every firm is reduced by $\sigma \ln(N_t/N_{90})$ from 2025Q3, where N_t is the number of firms selling in quarter t. In the spirit of Ackerberg and Rysman [50], this removes the gain from a larger choice set while keeping buyer heterogeneity, the outside option and rationing; it corresponds to a representative buyer whose valuation does not rise with the number of suppliers. We keep the discrete-choice structure because it accommodates capacity-constrained rationing and an outside option in a natural way. At the calibrated parameters, the firm-level own-price elasticity of unconstrained demand, $(\alpha/\sigma)p_{it}(1 - s_{it})$, averages 3.7 in 2023 and 2.7 in 2025Q3; Section 6.3 shows how the results change when products are more substitutable.

3.2. Costs and Pricing

Unit variable cost is the sum of a common silicon cost and a firm-specific processing cost:

$$c_{it} = m_t P_t + k_{it}, \qquad k_{it} = \max\left\{k_{\min},\, b_{it}\left[1 + \frac{Q_{it}}{\bar{l}K_{i0}}\right]^{-\nu}\right\}. \tag{3}$$

P_t is a quarterly polysilicon price index for 2020Q1–2025Q4 built from InfoLink weekly quotations. Silicon use per watt, m_t, follows engineering data from the IEA Photovoltaic Power Systems Programme (3.1, 2.7, 2.3 and 2.2 g/W in 2020–2023 [51,52]) and is held at its 2023 value afterwards. Its scale is set so that silicon accounted for 47.9% of module unit cost of goods sold in 2021Q4 [53]. After 2025Q4, the silicon price is held at its last observed level. The processing cost k_{it} falls with cumulative output Q_{it} through learning, with exponent ν, and with successful innovation through the cost base b_{it}, down to $k_{\min}$. Firms also pay a fixed operating cost proportional to operating capacity, and book capital depreciates at 2.5% per quarter; capacity under construction incurs neither.

Firms set prices as a markup over unit variable cost, $p_{it} = (1 + \mu_{it})c_{it}$, and adjust the markup to their market-share growth and to the gap between their utilisation u and the target $u^* = 0.85$:

$$\mu_{it} = \mathrm{clip}\left[\mu_{i,t-1} + \eta_s \frac{s_{i,t-1} - s_{i,t-2}}{s_{i,t-2}} + \eta_u(u_{i,t-1} - u^*),\, \mu_{\min},\, \mu_{\max}\right]. \tag{4}$$

Firms with sufficient cash may cut the markup below zero, down to −8%, to defend market share; other firms keep a minimum markup of 1.5%. Losses therefore arise mainly because prices fall below average total cost, which includes fixed operating costs and depreciation, rather than below variable cost.

3.3. Investment, Finance and R&D

Firms extrapolate recent demand growth and target capacity for a market share slightly above their current one:

$$\mathbb{E}_t M_{t+L} = M_t \exp(\phi h_t L), \qquad K^*_{it} = \frac{s_{it}(1+\gamma)\,\mathbb{E}_t M_{t+L}}{u^*}, \qquad I_{it} = \iota \max\{0,\, K^*_{it} - K_{it} - K^{\mathrm{pipe}}_{it}\}. \tag{5}$$

Here h_t is the average log growth of demand over the last four quarters, ϕ the strength of extrapolation, $L = 4$ the construction lag in quarters, K^{pipe} capacity under construction and $\gamma = 0.10$ the share-pressure parameter. Because every firm aims for a share 10% above its current one, aggregate targets are mutually inconsistent—a fallacy of composition that pushes utilisation below target. Investment is paid in full when ordered, less the local subsidy, and comes on stream after L quarters.

Share pressure is a scenario assumption rather than an estimate. It cannot be identified from the available annual moments: the moments that could pin it down concern investment, and γ, the adjustment speed ι and the extrapolation strength ϕ jointly determine its level. The behaviour it represents is, however, visible in the data. In early 2024, the ten largest module suppliers announced shipment targets for the year that summed to about 700 GW, while the same industry tracker projected global module demand of 460–519 GW; these ten suppliers actually shipped about 502 GW in 2024, after 413 GW in 2023 [54,55]. Announced targets thus exceeded what the market could absorb even under the forecasters' own demand projection—the mutual inconsistency that γ captures. Because announced targets combine share ambitions with expectations of market growth and may partly serve as signals, they motivate the mechanism but do not identify the value of γ. Section 6.2 therefore reports all key results for γ between 0 and 0.20.

Firms borrow under a leverage limit that implicit local guarantees relax. With budget softness $S \in [0, 1]$, the loan ceiling is $\Lambda(1 + \omega S)\max(E_{it}, 0)$ and the interest rate is $r(1 - dS)$,

where E_{it} is net worth. The same softness parameter governs rescues (Section 3.4), so a single number summarises how soft budgets are. Desired capacity in Equation (5) does not depend on profits, but realised investment does: a firm can spend only what its cash above a liquidity reserve and its borrowing room allow, and the borrowing room rises with net worth. Profits can therefore raise investment through financing, through survival—each surviving firm pursues its own share target—and through entry, whose rate rises with the industry profit rate (Section 3.4). Section 5.3 examines these channels for the own-cost floor.

R&D spending has a revenue-linked base and a profit-linked component, subject to available cash:

$$\mathrm{RD}_{it} = \min\{\rho_0 \mathrm{Rev}_{it} + \rho_1 \max(0, \pi_{it}),\ \text{available cash}\}, \tag{6}$$

where π_{it} is profit after interest but before R&D and taxes. Innovation succeeds with probability $1 - \exp[-\zeta \mathrm{RD}_{it}/(c_{it}K_{it})]$ and lowers the processing-cost base by a random step. The estimated profit coefficient ρ_1 is close to zero, so R&D moves with revenue rather than profit; Section 6.1 reports a profit-sensitive alternative.

3.4. Entry, Rescue and Local Governments

Potential entrants arrive in each region at a rate that rises with the industry profit rate and the local subsidy rate; they start at the processing-cost frontier plus a random gap and must build capacity before producing. A firm is in distress when its cash or net worth turns negative, and the local government rescues it with probability

$$\Pr(\text{rescue}) = S\left(\frac{\mathrm{Rev}_{it}}{\mathrm{Rev}_{rt}}\right)^{\kappa} \mathbf{1}\{\text{fiscal room}\}, \tag{7}$$

where Rev_{rt} is regional industry revenue and $\kappa = 0.7$, so locally important firms are more likely to be saved. A rescue restores minimum buffers of cash and equity. Firms that are not rescued go bankrupt: 20% of their assets are lost as a real resource cost, their capacity is scrapped and an employment-adjustment cost is incurred per unit of exiting capacity. Each region's fiscal balance receives land revenue (which declines from quarter 60), corporate taxes and fines, and pays for subsidies and rescues. Regions compete for growth: a region raises its investment-subsidy rate when its industry grows more slowly than in other regions and lowers it under fiscal pressure, and subsidy rates are projected so that their cross-regional mean stays fixed. This separates the competitive reallocation of subsidies from their average level.

3.5. Policy Instruments

We model three price floors. The variable cost floor requires $p_{it} \geq c_{it}$. The own-cost floor requires a price at least equal to the firm's audited average total cost of the previous quarter: unit variable cost plus fixed operating cost and depreciation per unit of last quarter's sales. The reference floor requires a price at least equal to the capacity-weighted mean of audited average total cost in the industry. Audited cost excludes R&D, interest and taxes. Enforcement is imperfect. A firm compares the payoff from complying with that from violating the floor, including current contribution profit, the value of future market share over eight quarters, a liquidity value of cash and the expected fine, which equals the detection probability θ times the fine multiple ψ times violating sales. We set $\theta = 0.25$ and $\psi = 1.0$, values chosen to match target violation rates rather than outcomes.

The other instruments are a hard budget, which sets S to zero for rescues, guarantee-relaxed leverage and interest discounts; subsidy withdrawal, immediate or phased in over eight quarters, which sets the subsidy cap to zero; an admission rule, which allows new capacity only if the firm's processing cost is within 10% of the frontier; forced retirement of

15% of operating capacity, taken from the highest-cost whole firms without compensation; and a demand stimulus that raises potential demand by 15% for eight quarters at a real cost of 0.05 per additional unit. Forced retirement represents coordinated capacity reduction and managed exit, and the admission rule represents capacity standards for new plants. We also study combinations of these instruments and the order in which budgets are hardened and floors introduced.

3.6. Welfare

Welfare over a policy window of T quarters starting at t_0 is

$$W = \sum_{t=t_0}^{t_0+T-1} \delta^{t-t_0}\left[\mathrm{CS}_t + \mathrm{Rev}_t - \mathrm{VC}_t - \mathrm{FC}_t - \mathrm{CAPEX}_t - \mathrm{RD}_t - \mathrm{BL}_t - \mathrm{EA}_t - \mathrm{SC}_t - \lambda_g F_t\right] + \delta^T \overline{K}_{t_0+T} - \overline{K}_{t_0}, \quad (8)$$

where VC and FC are variable and fixed operating costs, CAPEX is investment at full cost before subsidies, BL bankruptcy losses, EA employment-adjustment costs, SC the real cost of demand stimulus, F net fiscal spending (subsidies, rescues and stimulus outlays minus taxes) and $\overline{K}$ the book value of capital in use. We set $\delta = 0.99$ per quarter and the marginal cost of public funds $\lambda_g = 0.2$. Subsidies, rescues, taxes, fines, interest and default losses are transfers: they cancel in W and affect it only through λ_g. W decomposes exactly into the surplus of consumers, firm owners, local governments and banks, net of real losses and the cost of public funds, and we verify this identity in every run.

Domestic welfare excludes the surplus of foreign buyers, who are assumed to enjoy the same surplus per unit as domestic buyers. Because published production figures for 2024 conflict, we use two export-share scenarios: 41.7% in 2023 in both, then 37.6% in the low-export scenario and 40.1% in 2024 and 41.7% from 2025 in the high-export scenario. We report the low-export scenario unless stated. The equal-surplus assumption is strong, because export products may differ in product mix, warranty terms and landed prices. Global welfare does not depend on it. If foreign buyers' surplus per unit were a constant fraction of domestic buyers'—lower, for example, because tariffs and freight raise their landed prices—the true domestic effect would lie between our domestic and global estimates; if the difference varied with policy, this bound would not hold.

3.7. Involution and Failed Exit

A quarter is in involution when capacity utilisation is below 75%, at least half of operating firms make operating losses, and firms with at least four consecutive quarters of operating losses hold at least 10% of industry capacity; all three conditions must hold for at least four consecutive quarters. A spell is therefore confirmed in real time only in its fourth consecutive quarter. Operating profit is revenue minus variable and fixed operating costs, R&D and depreciation, which matches the pre-interest operating profit in company reports. Price–cost ratios and R&D are reported as outcomes but are not used in the definition. Clearance requires four consecutive quarters in which utilisation returns to within five percentage points of the same seed's pre-downturn mean (quarters 60–79), fewer than half of firms make operating losses and persistent-loss capacity is below 5%. We measure failed exit by cumulative exit capacity and by the high-cost capacity share: the share of capacity held by firms whose unit cost exceeds 110% of the capacity-weighted median.

4. Data, Calibration and Model Validation

4.1. Data

All calibration targets come from public sources. Financial moments use the 2023–2025 annual reports of six leading listed module producers: LONGi Green Energy, Trina Solar, JA Solar, JinkoSolar, Canadian Solar and Risen Energy. Utilisation is module output divided

by disclosed effective capacity for the four firms that report both. The annual price decline is computed from module revenue and shipments for the three firms where the two can be matched, and then averaged across firms; we use this realised-price measure rather than spot quotations because it is consistent with the firm-level financial moments. R&D is expensed R&D. The silicon price index averages verifiable quotations from InfoLink's free weekly reports by quarter. These data have known limitations: the firms are vertically integrated, so consolidated statements include upstream segments and non-module businesses; average realised prices mix products and destinations; and the sample covers leading survivors, not all producers. We therefore treat the moments as proxies when interpreting the calibration.

The restriction to leading survivors deserves comment, because smaller producers differ in cost structure, access to credit, likelihood of exit and exposure to local support. No public quarterly financial data exist for non-listed module makers, so they cannot be targeted directly. Three features of the design limit how much the sample choice drives the policy results. First, the calibration targets industry-level and cross-firm moments, including the dispersion of gross margins, rather than firm-specific parameters of the leading firms, and the 26 simulated firms are heterogeneous in size, cost and quality. Second, the characteristics in which small producers differ most are state-dependent in the model rather than parameter-dependent: the rescue probability rises with a firm's share of regional revenue (Equation (7)), borrowing is capped by net worth, and entrants start small and must build capacity before producing. Third, Section 6.4 re-runs the key scenarios under parameterisations that make small firms less likely to be rescued and that make pricing more aggressive and credit tighter.

4.2. Defining Involution in the Data

An operational definition is useful only if it classifies the case everyone calls involution as involution. Before running the simulations reported here, we checked candidate definitions against quarterly operating profit of the same six firms in 2021–2025, requiring at least four consecutive quarters of involution in 2024Q1–2025Q4 and none in 2021Q1–2022Q4. A stricter definition that also requires suppressed R&D and a minimum share of zombie firms fails this test, because expensed R&D intensity rose during the downturn. The definition in Section 3.7 passes: on public data it dates an involution spell from 2024Q4 onwards, confirmed in real time from 2025Q3.

4.3. Calibration

We calibrate seven parameters by the simulated method of moments [3], after screening 17 candidate parameters with the Morris [56] method: the unit cost of capital, the speed of investment adjustment, the strength of extrapolation, the dispersion of initial costs, innovation efficiency and the two R&D coefficients. The objective is a diagonal quadratic distance over ten moments with fixed moment scales. We minimise it by differential evolution [57] over 1519 candidate evaluations of 20 seeds each and do not claim convergence to a global minimum; final fits use 200 seeds. The remaining parameters are design assumptions or values from a preliminary calibration that we held fixed (Appendix B). Table 2 and its note report the calibration fit.

Before running the calibration, we fixed the criterion that at least seven of the ten moments must be reproduced, in the sense that the public proxy lies inside the 5–95% band across seeds or the relative error of the model mean is below 25%. The 25% tolerance is an operational choice rather than a statistical test. It reflects the considerable uncertainty of the public proxies: the firms are vertically integrated, their consolidated statements include upstream and non-module businesses, and average realised prices mix products,

destinations and contracts signed at different times. Because a tolerance alone can flatter the fit, Table 2 and its note report the seed-band comparisons and relative errors.

Table 2. Fit of the eight calibration moments that meet the calibration criterion.

Moment	Public Proxy	Model Mean	Seed 5–95% Band	Relative Error (%)	Status
Module utilisation, 2023–25 (%)	66.31	72.30	[68.89, 75.34]	9.04	Within tolerance
Module price decline, 2024 (%)	32.77	24.89	[24.17, 25.73]	24.05	Within tolerance
Net-loss firms, 2024 (%)	66.67	66.98	[33.90, 93.66]	0.47	In band
Net-loss firms, 2025 (%)	83.33	79.33	[56.66, 96.67]	4.81	In band
Mean gross margin (%)	8.72	7.72	[5.46, 9.81]	11.43	In band
Gross-margin dispersion (pp)	3.18	3.21	[2.63, 3.81]	0.87	In band
Expensed R&D/revenue (%)	1.78	1.65	[1.65, 1.66]	7.13	Within tolerance
R&D intensity, 2025/2023 (ratio)	1.31	1.00	[0.98, 1.00]	23.75	Within tolerance; direction not reproduced

Note. In band: the public proxy lies in the 5–95% band across 200 seeds. Within tolerance: the proxy lies outside the band, but the relative error of the model mean is below the pre-registered 25%. The band reflects seed variation only, not parameter or measurement uncertainty. The two remaining targets are the 2025 price decline (public proxy 16.43%, model 4.45%; relative error 72.91%) and capital expenditure/average total assets (9.09%, 4.00%; relative error 55.96%); neither meets the matching criterion. The original acceptance rule additionally required both annual price declines and R&D intensity to meet the criterion. Because the 2025 price decline did not, the full original rule was not met; before policy simulations, we retained the estimate without re-estimation and report its fit here.

For the moments in Table 2, the model reproduces a sharp price fall in 2024, the rise in the share of loss-making firms, the level and dispersion of gross margins and the level of R&D intensity. Two qualifications matter. The model reproduces only three-quarters of the 2024 price decline (24.9% against 32.8%). The R&D response ratio is within tolerance, but its direction is wrong: R&D intensity rose by 31% in the data and stayed flat in the model. Because the model understates the 2024 price fall, it probably also understates how far prices fell below average total cost, and hence how often a floor at own cost would bind. An alternative input-cost specification that adds a non-silicon cost path overshoots the observed decline (36.1%), so that the two specifications bracket the data. The key welfare signs are the same in both (Section 6.1): the own-cost floor lowers welfare by 1.09% and 2.21%, the reference floor raises it by 1.81% and 1.70%, and the hard budget combined with the own-cost floor lowers it by 1.30% and 2.56%. The direction of the main welfare results therefore does not hinge on the specification that understates the price decline.

In a parameter-recovery exercise with three synthetic parameter vectors, the largest recovery error per vector, scaled by the parameter range, was 48.2%, 36.9% and 16.5%: fixed costs, learning and investment parameters partly substitute for one another, so the estimates should not be read as precise structural values. Weak identification of this kind is common in agent-based models calibrated to a small number of moments [58,59]. Our response is not to treat the point estimates as precise but to evaluate the policy conclusions across 30 parameter vectors that fit the data comparably (Section 6.2) and in one-at-a-time checks of individual assumptions (Section 6.1), which ten Broeke et al. [60] recommend as a first step in the sensitivity analysis of agent-based models.

4.4. Baseline Dynamics

Figure 1 shows the no-policy baseline. Utilisation is close to full during the early expansion, declines as capacity outpaces demand and reaches about 71% around 2025Q3.

The ratio of price to average total cost falls to about one when silicon prices collapse and drifts below one afterwards, and the share of loss-making firms jumps to 70–80%. At the policy start in 2025Q3, 82% of seeds are in an involution spell judged ex post and 64.5% are confirmed in real time; in the public data, the spell is identified from 2024Q4 and confirmed in real time from 2025Q3. The downturn therefore starts somewhat earlier in the model, but the state at the policy start matches.

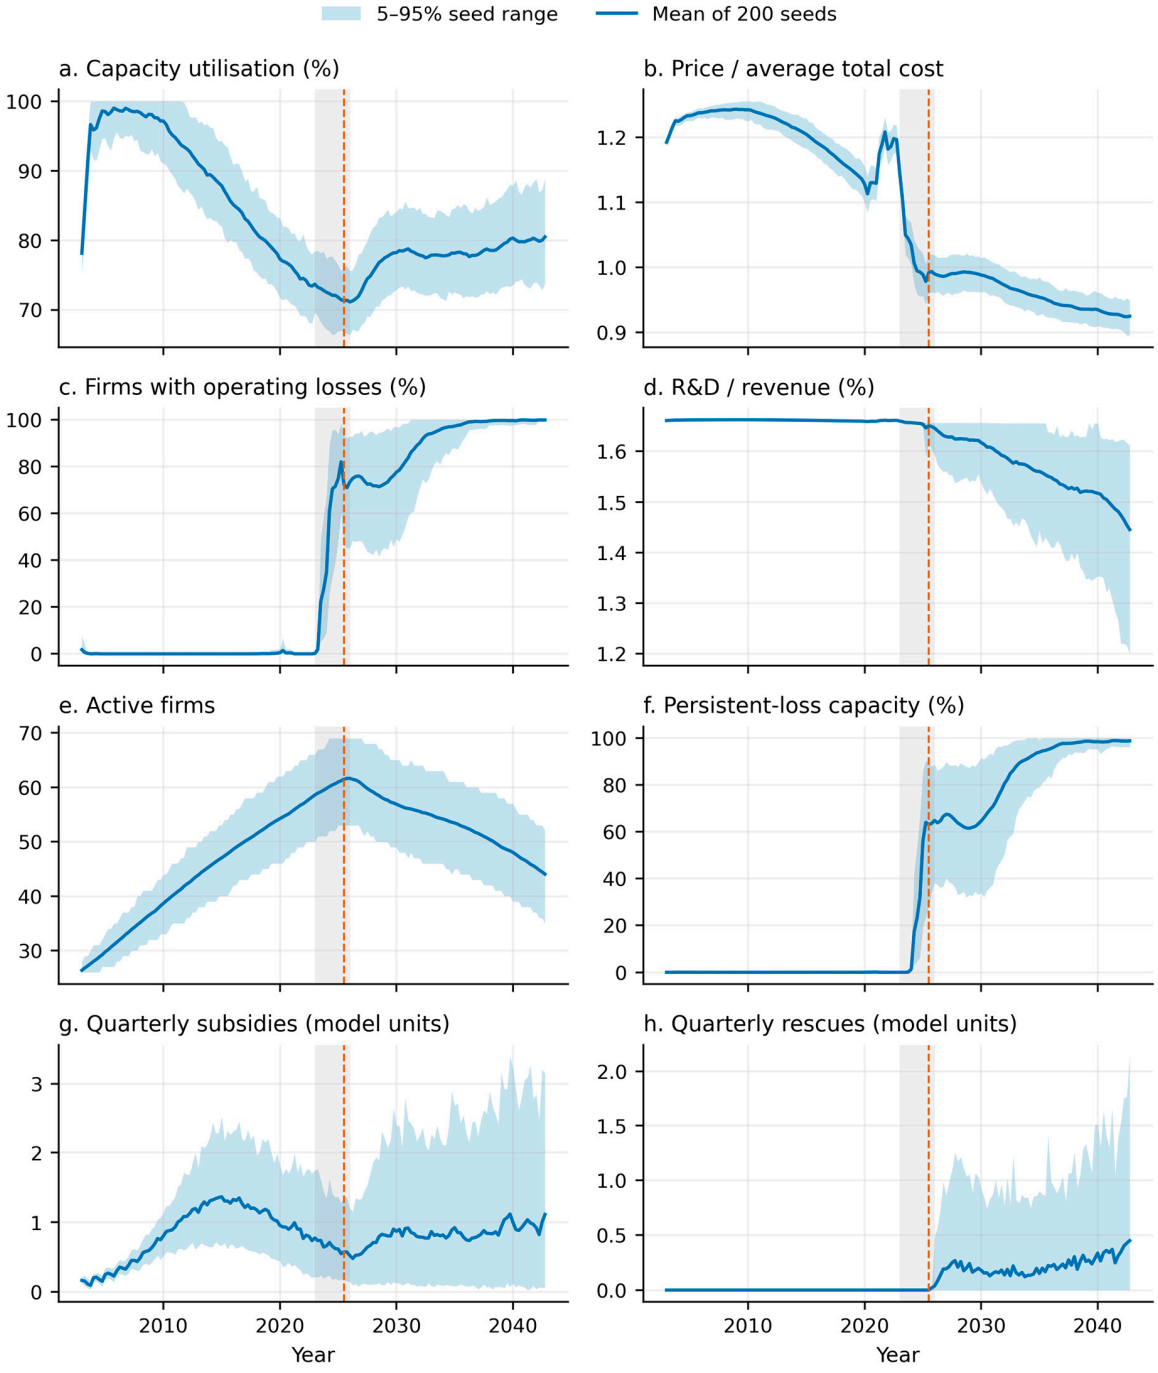


Figure 1. Baseline dynamics without policy. Panels show (**a**) capacity utilisation (%), (**b**) the ratio of price to average total cost, (**c**) the share of firms with operating losses (%), (**d**) R&D expenditure relative to revenue (%), (**e**) the number of active firms, (**f**) the share of persistent-loss capacity (%) and quarterly (**g**) subsidies and (**h**) rescues in model units. Lines show means over 200 seeds and shaded areas the 5–95% seed range. The grey band marks the calibration window (2023–2025) and the dashed line marks the policy start (2025Q3). Years before 2023 are not calibration targets; the observed polysilicon price path enters in 2020Q1.

Without policy, the industry does not return to profitability within the horizon. The share of loss-making firms approaches 100% after about 2033; firms keep exiting slowly and rescues rise, and no seed meets the clearance criterion. The bottleneck is profitability rather than capacity use: after 2028, utilisation is within five percentage points of its pre-downturn mean in 92% of seed-quarters, but persistent-loss capacity never falls below 5%. This persistence reflects share-target investment, credit access and rescues. We therefore read the baseline as a stylised scenario of persistent involution, not as a forecast, and interpret all policy results as differences from it. Two checks address how much the magnitudes depend on this baseline. The welfare and utilisation effects measured over the first eight quarters, which depend little on the long-run baseline, have the same signs as the 40-quarter effects for both binding floors, the hard budget, the admission rule and forced retirement (Appendix C Table A8). And the share-pressure and future-demand checks of Section 6 change the persistence of the baseline itself: the welfare signs of the two binding floors are unchanged for share pressure up to 0.15 and under both demand paths, except that the own-cost floor's effect becomes insignificant under the faster of the two, whereas the admission rule's welfare effect reverses when demand grows faster.

4.5. Supplementary Pattern Checks

Calibration and validation should be distinguished [46]. The checks in this section are not a strict out-of-sample validation. None of the patterns below was a calibration target, but the 2021–2025 operating-profit data of the six firms were used to choose the involution definition (Section 4.2), and the share of firms with operating losses is related to the targeted share of firms with net losses. We therefore read them as supplementary consistency checks and report what the model does and does not reproduce. Three patterns are reproduced. First, the timing of involution: the spell is confirmed in real time in 2025Q3 in the data and in 64.5% of seeds (Section 4.4). Second, the boom: in 2021–2022, the public data show no involution quarter and a share of firms with quarterly operating losses of at most one third in 2021 and zero in 2022; in the model, no seed is in involution in these years and the share of operating-loss firms is below 0.2%. Third, the onset of losses: the quarterly share of firms with operating losses averages 0% in 2023, 58% in 2024 and 96% in 2025 in the data, and 13%, 59% and 75% in the model.

The model does not reproduce the amplitude of the utilisation cycle. In the data, average utilisation of the reporting firms rose from 69% in 2021 to 80–82% in 2022–2023 and fell to 60% in 2024 and 56% in 2025, whereas the simulated average declines gradually from 75% in 2021 to 71% in 2025. The model thus matches the average of the downturn but not its sharpness, a further reason to read the results as comparisons within a stylised scenario.

5. Results

5.1. Design and Inference

All policies start in 2025Q3 (quarter 90), the quarter of the Central Financial and Economic Affairs Commission meeting and the draft Price Law amendment. Unless stated otherwise, initial budget softness is 0.8 and floors are enforced as described in Section 3.5. We report effects over the 40 quarters from 2025Q3 to 2035Q2. Every scenario uses the same 200 seeds as its no-policy counterpart, so each effect is a paired difference with a Student-t confidence interval over seeds. These intervals measure simulation noise for given parameters, not parameter or structural uncertainty, which Section 6 addresses. Relative effects are percentages of the no-policy mean of the same outcome. We report welfare under both logit and variety-neutral demand. Because the model generates many contrasts, we emphasise effects whose sign is stable across the parameter vectors of Section 6.2, and Section 6 reports where alternative specifications reverse them.

5.2. What Sustains Involution?

To see which mechanisms generate involution, we switch seven channels on and off in all 128 combinations and compute their Shapley [61] contributions to three outcomes: a composite index that averages the normalised gaps of the three involution conditions, the number of quarters in involution, and the high-cost capacity share (Figure 2; Appendix C Table A2). Switching a channel off means removing subsidies, holding subsidies at their mean instead of letting regions compete, removing rescues and guarantees, differentiating products, replacing extrapolative by static expectations, shortening the construction lag to one quarter, or removing share pressure.

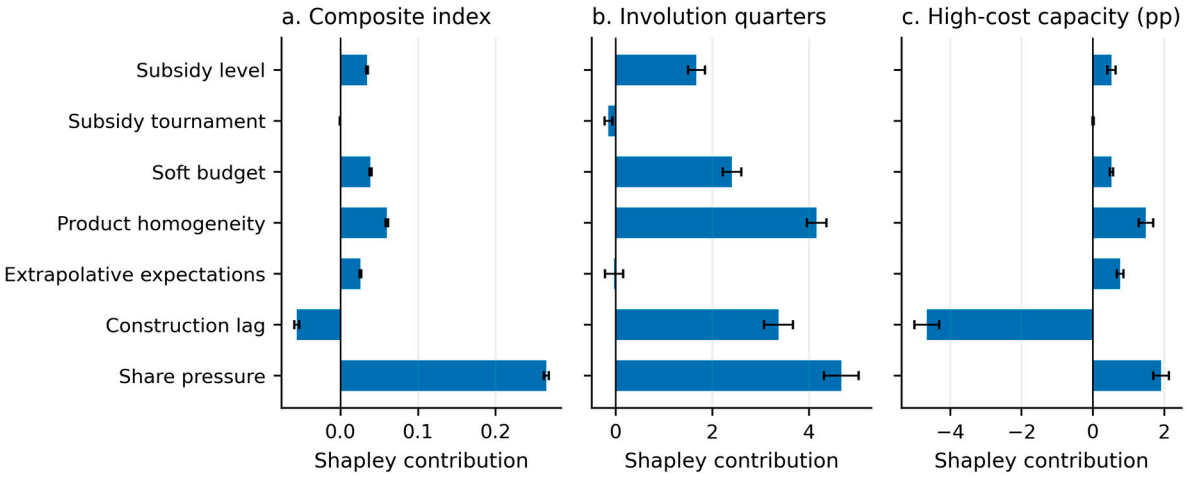


Figure 2. Shapley contributions of the seven channels to (**a**) the composite involution index, (**b**) the number of involution quarters and (**c**) the high-cost capacity share over 40 quarters from 2025Q3, with 95% confidence intervals. Positive values raise the outcome.

Share pressure is the dominant mechanism. It accounts for 72% of the difference in the composite index between all channels on and all off, adds 4.7 quarters of involution and raises the high-cost capacity share by 1.9 percentage points: when every firm targets a larger share than it has, aggregate investment overshoots demand even with correct forecasts. Product homogeneity (4.2 quarters), longer construction lags (3.4 quarters) and soft budgets (2.4 quarters) also lengthen involution, followed by the level of subsidies (1.7 quarters). Holding the mean subsidy fixed, competition among regions contributes little and even shortens involution slightly (−0.15 quarters): what matters is how much is subsidised, not the tournament itself. Higher subsidies and longer lags lower utilisation by 0.48 and 0.64 percentage points, consistent with the evolutionary model of Zhang and Dong [10].

Longer construction lags lower the composite index and the high-cost share but lengthen involution. The two measures capture different things: the composite index averages the normalised gaps of the three conditions and so measures severity, whereas the number of involution quarters counts the quarters in which all three conditions cross their thresholds at once and so measures duration. Capacity under construction incurs neither fixed operating costs nor depreciation. With longer lags, a larger part of new investment is still in the pipeline at any time and its costs start later, so fewer firms make operating losses (−8.5 percentage points) and less capacity is persistently loss-making (−8.9 points). At the same time, capacity ordered before demand slows keeps arriving for longer, so the utilisation shortfall widens slightly and the three conditions hold jointly for more quarters. Longer lags therefore make involution milder but more persistent. The diagnosis points to the returns to expansion; the rest of this section asks whether price floors or capacity instruments act on them.

5.3. The Own-Cost Floor: A Soft Budget Financed by Buyers

How much a floor binds depends on its cost basis. Over the 40 quarters, the variable cost floor binds in 4 of 452,535 firm-quarters: even at the trough, prices stay above unit variable cost, because losses arise from fixed costs and depreciation. This floor therefore has no measurable effect, and we do not discuss it further. The own-cost floor binds in 63% of firm-quarters and the reference floor in 66%; violating sales are 0.3% and 1.1% of total sales.

Table 3 reports the main effects. The own-cost floor deepens involution rather than curing it. Over 40 quarters, it lowers utilisation by 5.8 percentage points, reduces cumulative exit by 2.3%, more than doubles the number of quarters in involution (+112%) and lowers consumer surplus by 1.6%, global welfare by 1.1% and domestic welfare by 0.8%; under variety-neutral demand, the welfare losses are almost identical (1.14% and 0.84%). It raises revenue by 3.4% while output falls slightly, and it raises investment by about 30%. With more capacity and lower sales, utilisation falls, fixed costs per unit rise and firms remain loss-making on an operating basis. The welfare loss equals 3.7% of the no-policy industry's discounted revenue over the same ten years.

Table 3. Effects of policy instruments over 40 quarters (2025Q3–2035Q2).

Instrument	Utilisation (pp)	High-Cost Capacity (pp)	Exit Capacity (%)	Involution Quarters (%)	Consumer Surplus (%)	Global Welfare, Logit (%)	Global Welfare, Variety-Neutral (%)	Domestic Welfare, Logit (%)	Domestic Welfare, Variety-Neutral (%)
Own-cost floor	−5.79	0.38	−2.3	111.7	−1.56	−1.09	−1.14	−0.80	−0.84
Reference floor	−6.50	0.65	−75.2	−52.7	−0.50	1.81	0.48	3.26	1.90
Hard budget	0.78	−0.47	−1.5	−36.0	−0.74	−0.53	−0.01 †	−0.39	0.13
Hard budget + own-cost floor	−4.00	−0.22	−5.3	70.1	−1.93	−1.30	−0.85	−0.90	−0.45
Admission rule	3.79	−0.26	−9.3	−29.6	−0.26	0.53	0.56	1.04	1.06
Forced retirement	3.15	−7.27	63.9	−71.6	−2.26	−2.61	0.37	−2.83	0.20
Subsidy withdrawal	0.18	−0.16	−4.4	−7.6	−0.20	−0.05	0.12	0.05	0.22
Withdrawal + hard budget + admission	5.77	−0.97	−11.8	−68.8	−1.19	0.12	0.72	0.95	1.55
Demand stimulus	−0.45	−0.47	16.2	20.9	3.13	2.84	2.96	2.65	2.78

Note. Paired mean differences from the no-policy run over 200 seeds; initial budget softness 0.8. Percentages are relative to the no-policy mean of the same outcome (utilisation 76.5%, high-cost capacity share 27.1%, 10.9 involution quarters). The utilisation, capacity, involution and consumer-surplus columns refer to the logit specification. Variety-neutral welfare subtracts $\sigma \ln(N_t/N_{90})$ from every firm's mean utility from 2025Q3 (Section 3.1). Domestic welfare uses the low-export scenario. † Not significant at the 5% level; all other entries are significant. Appendix C Table A3 gives 95% confidence intervals and further combinations.

Why do protected margins raise investment when desired capacity does not depend on profits? To examine the channels, we recorded each firm's desired and realised investment. The own-cost floor raises incumbents' desired investment by 40% (199 model units over 40 quarters) and their realised investment by 34% (159 units), whereas the initial investment of entrants rises by only 2.7 units; entrants' later investment is included in the incumbents' figures. A relaxation of financing constraints does not explain the increase: the share of desired investment that firms cannot finance rises under the floor, from 5.2% to 9.1%. The rise in desired investment is consistent with two other forces. More firms survive—1.5

more are active on average—and each pursues its own share target. And market shares are reshuffled more often: because each firm's floor is its own audited cost per unit of last quarter's sales, a firm that loses sales faces a higher floor in the next quarter, whereas a firm that gains sales faces a lower one, and market-share turnover across firms rises by 26%. Firms that gain share invest towards their targets, while firms that lose share keep their capacity. These diagnostics do not separate the channels causally. Switching off entry, for example, changes the baseline fundamentally—only 26 firms remain active on average, against 57, and involution hardly arises—so it cannot isolate the role of entry.

Table 4 decomposes the welfare effects into resource components. Buyer value is consumer surplus plus revenue, which is buyers' gross valuation of what they buy before transfers. The own-cost floor loses welfare on two margins. Higher prices reduce buyer value by 142.5 units, and additional investment (138.2 units) and higher fixed costs (115.3 units) add to the loss; lower variable costs and a larger terminal capital stock offset only part of it. The floor reproduces in miniature the mechanism that created the overcapacity: protected margins sustain capacity.

Table 4. Resource decomposition of welfare effects (40 quarters, model units).

Component	No-Policy Level	Own-Cost Floor	Reference Floor	Hard Budget	Admission Rule	Forced Retirement	Demand Stimulus
Buyer value (consumer surplus + revenue)	31,783.8	−142.5	394.3	−112.9	−40.7	−687.7	1055.6
Variable costs	−5333.3	82.4	2.6	−20.5	22.0	265.4	−193.7
Fixed operating costs	−1500.8	−115.3	−131.1	19.1	61.8	64.9	−57.4
Capital expenditure	−461.7	−138.2	112.8	−5.4	105.3	−198.4	−72.9
R&D	−114.5	−6.7	−11.4	−2.1	−0.9	1.0	−3.2
Bankruptcy and adjustment losses	−60.0	2.9	45.7	−4.4	6.6	−103.8	−9.0
Terminal minus opening capital	−316.4	54.5	14.6	−1.9	−30.1	34.4	−5.0
Cost of public funds and stimulus	−5.0	0.1	6.4	1.6	4.0	−0.9	−34.2
Global welfare	23,992.1	−262.6	433.8	−126.3	128.0	−625.1	680.2

Note. Discounted sums over 2025Q3–2035Q2 in model units; paired means over 200 seeds, initial softness 0.8, logit demand. Each entry is a contribution to welfare, so positive numbers raise welfare. Components may not add up exactly because of rounding.

Do hard budgets turn the floor into a clearing device? They do not. Figure 3 shows the floors' effects by initial budget softness. Under fully soft budgets (softness 1.0), the own-cost floor reduces cumulative exit by 3.6% [1.1, 6.0]; under fully hard budgets (softness 0), it reduces exit by 6.9% [4.3, 9.5]. Exit falls more when budgets are hard, because without rescues the floor becomes the main force keeping weak firms alive. In the main specification, the floor also raises the high-cost capacity share by 0.43 points under soft and 0.32 points under hard budgets, but this effect is not robust across parameter vectors (Section 6.2), whereas the slower exit holds in all of them. Hardening budgets at the same time as introducing the floor lowers utilisation by 4.0 points relative to no policy, lengthens involution by 70% and lowers welfare by 1.3%; hardening reduces the floor's marginal welfare loss by about 30% but does not reverse it. Figure 4 traces the time paths: a hard budget raises utilisation within two years, whereas both cost-based floors keep it near 70% for the whole decade. In effect, a floor at a firm's own full cost guarantees that buyers cover each firm's audited cost: it is a soft budget constraint financed through prices rather than by governments.

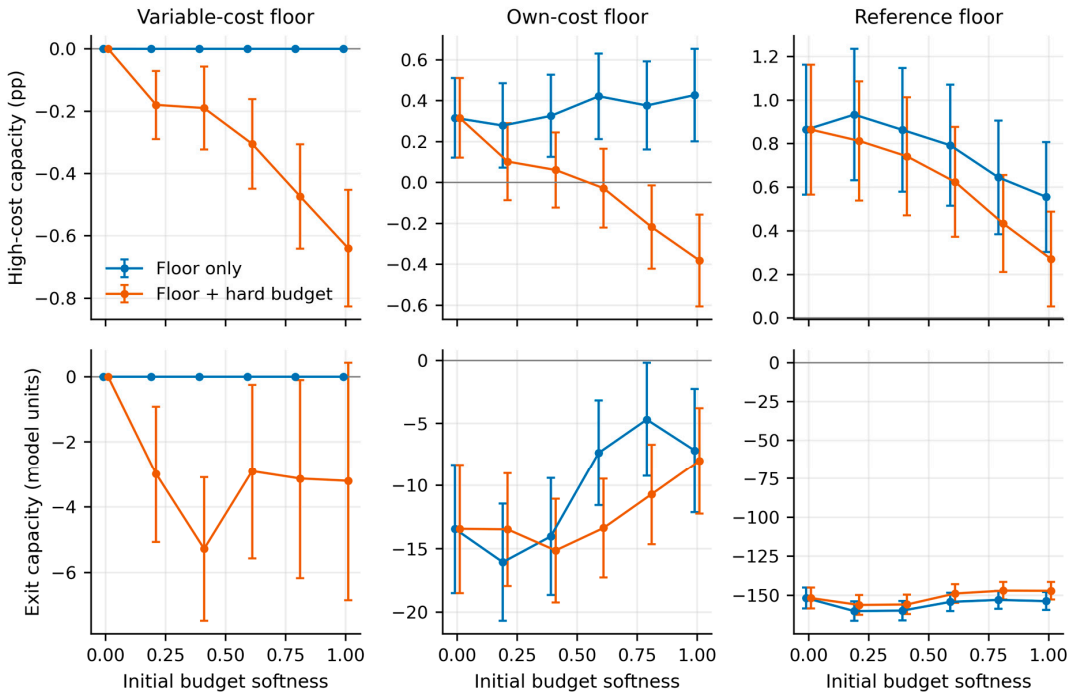


Figure 3. Effects of the three floors on the high-cost capacity share (**top**) and cumulative exit capacity (**bottom**) after 40 quarters, by initial budget softness, with the floor alone (blue) and with budgets hardened at the policy start (orange). For the variable cost floor, which does not bind, the orange line shows the effect of the hard budget alone. Bars are 95% confidence intervals.

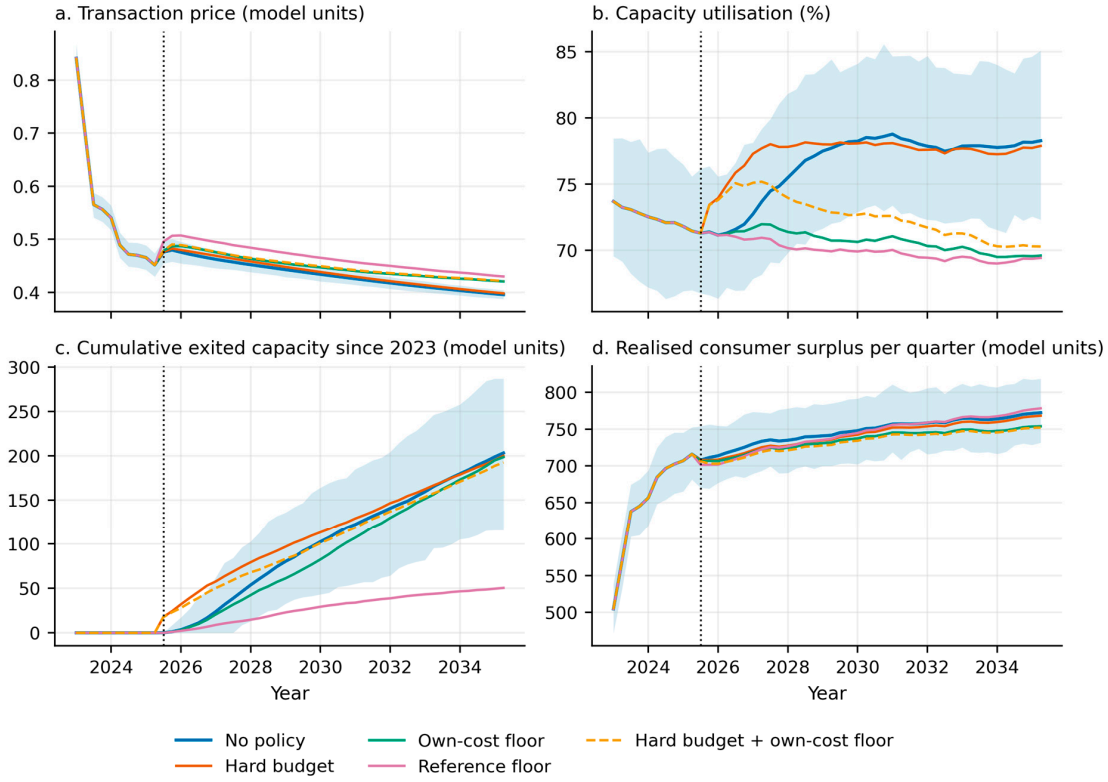


Figure 4. Time paths of (**a**) the transaction price, (**b**) capacity utilisation, (**c**) cumulative exited capacity since 2023 and (**d**) quarterly realised consumer surplus under no policy, a hard budget, the own-cost floor, the reference floor and the hard budget combined with the own-cost floor. Lines are means over 200 seeds; the shaded area is the 5–95% seed range of the no-policy run. The dotted line marks the policy start in 2025Q3.

If a floor is introduced at all, its timing matters to some extent. Hardening budgets four quarters before introducing the own-cost floor yields 0.61% [0.54, 0.68] higher welfare than the reverse order and 0.26% more than introducing both at once (Appendix C Table A4). Because the leading instrument is in force longer, these comparisons reflect both the order and the duration of exposure. The advantage also depends on how R&D responds to profits: when R&D responds strongly ($\rho_1 = 0.15$), hardening first loses by 0.47%, although that variant implies a 2025/2023 R&D-intensity ratio of 0.60 against 1.31 in the data. We therefore treat the sequencing result as conditional. Nor does the floor pay off for consumers later, unless R&D responds strongly to profits. Relative to a hard budget alone, adding the own-cost floor lowers consumer surplus by 0.73% after 8 quarters and by 1.19% after 40. R&D spending rises by 4.9%, but successful innovations barely rise (0.3%), and lower processing costs do not offset higher markups; only with profit-sensitive R&D does the long-run effect turn slightly positive (+0.13%).

5.4. The Industry-Reference Floor: Stabilisation Without Exit

The reference floor works differently, because it binds mainly on efficient firms. Firms with costs at or below the industry median are constrained in 97.5% of their firm-quarters, whereas firms above 110% of the median are constrained in only 17.1%. The floor thus stops low-cost firms from undercutting, which lifts the market price by 7.4% while total sales are unchanged, and high-cost firms gain from the higher price without being constrained. Bankruptcies fall by 78.5%, cumulative exit capacity by 75.2%, and the number of active firms at the end of the window is 28.5% higher than without policy. Utilisation falls by 6.5 points. Involution quarters fall by 53%, but only because margins recover and fewer firms make operating losses, not because capacity leaves the industry.

Global welfare rises by 1.8% and domestic welfare by 3.3% under logit demand, but the source of the gain matters. About 91% of it (394.3 of 433.8 units) is an increase in buyer value at unchanged sales (Table 4). In a logit model, each surviving firm is a distinct option that buyers value, so a larger number of surviving firms raises measured buyer value even when prices rise. Savings on capital spending (112.8 units) and on bankruptcy and adjustment losses (45.7 units) add to the gain, while the larger number of operating plants raises fixed costs by 131.1 units. PV modules are close substitutes, and logit demand is known to overstate the value of additional varieties when products are similar [50]. Under variety-neutral demand, the reference floor's global welfare gain falls to 0.48% [0.45, 0.51] and its domestic gain to 1.90% (Table 3). When products are more substitutable, the gain shrinks further and eventually turns into a loss (Section 6.3). The gain is also unevenly distributed: firm owners gain 708 units and local governments 51 units, whereas banks lose 256 units and consumers 122 units of surplus (Appendix C Table A5).

5.5. Capacity-Side Instruments

Instruments that act on capacity perform differently (Table 3). The admission rule raises utilisation by 3.8 points, shortens involution by 30% and raises welfare by 0.5% globally and 1.0% domestically, with almost identical gains under variety-neutral demand; it gains by preventing investment (105.3 units) and saving fixed costs (61.8 units) at a small cost in buyer value (Table 4). This gain depends on the path of future demand: in the baseline, demand plateaus after 2025, and the rule loses welfare when demand grows as fast as industry projections suggest (Section 6.1). Forced retirement cuts the high-cost share by 7.3 points and involution by 72%, but at the largest measured welfare cost under logit demand (−2.6%): it destroys capacity that buyers value as variety, triggers bankruptcy losses and induces replacement investment of 198.4 units. A hard budget shortens involution by 36% at a welfare cost of 0.5%, almost entirely through lower buyer

value, as early bankruptcies shrink the choice set. Both losses rest largely on the valuation of variety: under variety-neutral demand, forced retirement raises welfare by 0.37% and the hard budget's effect is indistinguishable from zero. Subsidy withdrawal alone has small effects. The package that combines subsidy withdrawal, hard budgets and the admission rule raises utilisation by 5.8 points and shortens involution by 69%, with a small welfare gain under logit demand (0.1% globally, 0.9% domestically) and a larger one under variety-neutral demand (0.7% and 1.6%). Adding the own-cost floor to this package reduces the utilisation gain to 2.0 points and turns the global welfare effect negative (−0.5%).

The demand stimulus raises welfare most (2.8% globally), but it lowers utilisation and lengthens involution by 21%, and its gain depends on an assumed real cost of 0.05 per unit and on counting all additional demand as genuine buyer value. Counting only domestic buyers changes the ranking at the top (Appendix C Figure A1). Instruments that raise prices look better from a domestic perspective, because part of the lost consumer surplus belonged to foreign buyers: the reference floor gains 3.3% of domestic welfare against 1.8% of global welfare and ranks first domestically under logit demand, and the own-cost floor loses 0.8% instead of 1.1%. Because roughly 40% of modules are sold abroad, part of what the domestic industry loses in a price war is a transfer to foreign buyers.

6. Robustness and Parameter Sensitivity

6.1. Alternative Specifications

We recomputed the main contrasts in 38 alternative specifications: alternative behavioural rules (a fixed markup, adaptive expectations, utilisation-based investment without a share target, profit-sensitive R&D, fixed subsidies without a tournament); rescue-only and guarantee-only versions of soft budgets; nine sets of welfare parameters; current-quarter rather than lagged audited cost, and a floor that covers only 70% of fixed costs; stronger enforcement; structural variants (half or 50% more firms, 10 or 30 regions, construction lags of two or six quarters, price sensitivity of demand ±30%, a continued silicon price decline); policy starts in 2024Q1 and 2025Q1 and a real-time trigger on confirmed involution; four independent consumer samples; a non-silicon input-cost path; and two paths of faster future demand growth. Appendix C Table A6 and Figure A2 summarise the results check by check.

Most welfare signs are robust. Hard budgets, the reference floor, forced retirement and the demand stimulus keep the sign of their welfare effects in all 38 checks, as does the welfare loss from combining hard budgets with the own-cost floor. The own-cost floor alone lowers welfare in 35 checks but raises it under profit-sensitive R&D (+0.84%) and with 30% less price-sensitive demand (+0.51%); with 30% more price-sensitive demand, its loss rises to 4.51%. The admission rule's welfare gain depends on future demand. The package of subsidy withdrawal, hard budgets and the admission rule lowers global welfare in 9 checks and domestic welfare in 6. The advantage of hardening budgets first holds in 36 of 38 checks, reversing under profit-sensitive R&D and turning slightly negative (−0.03%) when demand is 30% less price-sensitive.

The future-demand paths require explanation. In the baseline, potential demand follows a logistic path that is close to saturation by 2025, so it grows by only about 4% over the policy window—a plateau consistent with the slowdown of global installations to 9% growth in 2025 [62]. The two alternative paths scale demand to the China Photovoltaic Industry Association's projection of average annual global installations of 725 GW and 870 GW in 2026–2030, against 580 GW in 2025 [62,63]: from 2026Q1 to 2030Q4, potential demand is raised at an annual log rate of 7.3% (CPIA-low) or 13.0% (CPIA-high) and is then held at that level, ending 39% and 84% above the baseline. Both alternative paths therefore imply faster growth than the baseline. This explains why the admission rule's

welfare gain reverses under both (−0.58% and −3.44%): the rule blocks capacity expansion by firms whose costs are not near the frontier, which saves resources when demand plateaus but forgoes supply that buyers value when demand grows. Under the CPIA-high path, the rule lowers output by 2.4% and buyer value by 2041 units, more than the 737 units it saves in capital, fixed and variable costs; under the plateau, output is unchanged. The rule still raises utilisation under both paths (by 6.3 and 6.8 points). The reference floor raises welfare under both paths, and the own-cost floor lowers it under the CPIA-low path (−0.34%) and has no significant welfare effect under the CPIA-high path (−0.02%).

The floor's effects on capacity are less robust. Under soft budgets, both the slower exit and the higher high-cost share keep their sign in 32 of 38 checks. Exit rises significantly under adaptive expectations and with more price-sensitive demand, whether budgets are soft or hard, and, under soft budgets, also with ten regions, a 2024Q1 policy start or a non-silicon cost path. The high-cost share falls under utilisation-based investment and profit-sensitive R&D, two checks that change how firms' investment and R&D respond to margins. In no check does the floor under hard budgets both raise exit and lower the high-cost share.

6.2. Share Pressure and Parameter Uncertainty

Two further checks address the assumptions to which the results are most exposed. First, share pressure γ is the dominant mechanism (Section 5.2) but is fixed rather than estimated (Section 3.3); we re-ran the key policy scenarios at $\gamma = 0$, 0.05, 0.15 and 0.20 without re-estimating other parameters. Second, parameter recovery points to weak identification (Section 4.3). Of the 40 best candidates from the differential-evolution search, 31 match at least seven of the ten moments with 200 seeds, including the 2024 price decline and R&D intensity; we re-ran the key scenarios for the 30 with the lowest loss, which include the main estimate and vary all seven estimated parameters while holding γ at 0.10. Table 5 reports the key contrasts, together with the variety-neutral specification.

With the other parameters held at their estimates, $\gamma = 0.10$ gives the lowest calibration loss of the five values and matches the most moments; this is a conditional comparison, not an estimate of γ. At each value tested between 0 and 0.15, the own-cost floor lowers welfare by 1.0–1.4%, and the reference floor, the admission rule and hardening budgets first raise it; at every value of γ the floor lowers utilisation and slows exit. The increase in involution quarters, however, requires sufficient share pressure: at $\gamma = 0$ and 0.05 the floor shortens involution by 97% and 26%. A plausible reading is that, without strong expansion pressure, guaranteed margins remove the operating losses that the involution definition requires. At $\gamma = 0.20$, which matches only four moments, the floor raises global welfare by 0.50% and a hard budget lowers it by 11.8%. Most welfare effects grow with γ, but the welfare ranking of the twelve policies re-run in these checks stays close to the main ranking (Spearman correlations of 0.78–0.99).

Across the 30 parameter vectors, most conclusions hold without exception. In every vector, the own-cost floor lowers utilisation, lengthens involution and lowers global and domestic welfare, and it slows exit under both soft and hard budgets (significantly in 28 and 30 vectors); the reference floor raises welfare while cutting exit sharply (median −85%); hardening budgets first beats the reverse order; and the admission rule raises welfare and utilisation. Three results are fragile. The floor's effect on the high-cost share has the main sign in only 13 of 30 vectors under soft budgets and 11 under hard budgets. The package of subsidy withdrawal, hard budgets and the admission rule raises global welfare in only 7 vectors and significantly lowers it in 21, although it raises domestic welfare in 29. The small domestic gain from subsidy withdrawal alone has the main sign in only 13 vectors. The main estimate also lies at the edge of the set for several magnitudes: in most vectors,

the own-cost floor's welfare loss (median −0.94%) and the advantage of hardening first (median +0.42%) are smaller, and the reference floor's cut in exit is larger. The median rank correlation between each vector's welfare ranking and the main ranking is 0.99.

Table 5. Key contrasts under alternative share pressure, parameter vectors and demand specification (40 quarters).

Contrast (% of No-Policy Mean)	Main ($\gamma = 0.10$)	$\gamma = 0$	$\gamma = 0.05$	$\gamma = 0.15$	$\gamma = 0.20$	30 Vectors: Median [5–95%]; Same Sign	Variety-Neutral Demand
Own-cost floor: global welfare	−1.09	−1.03	−1.08	−1.37	+0.50	−0.94 [−1.04, −0.76]; 30/30	−1.14
Own-cost floor: involution quarters	+112	−97	−26	+153	+197	+114 [67, 190]; 30/30	+111
Own-cost floor, softness 1.0: exit capacity	−3.6	−39.7	−8.6	−18.0	−21.8	−8.5 [−13.9, −2.4]; 30/30	−3.5
Own-cost floor, softness 1.0: high-cost share	+1.57	−0.96	+1.38	+3.84	+5.81	−0.65 [−2.39, 1.83]; 13/30	+1.59
Own-cost floor, softness 0: exit capacity	−6.9	−12.1	−3.6	−18.1	−33.2	−9.9 [−16.5, −5.2]; 30/30	−7.1
Own-cost floor, softness 0: high-cost share	+1.20	+0.05 †	+1.43	+2.67	+7.37	−0.42 [−1.72, 1.47]; 11/30	+1.23
Reference floor: global welfare	+1.81	+0.98	+1.41	+2.19	+5.16	+1.57 [1.10, 2.63]; 30/30	+0.48
Hard budget: global welfare	−0.53	−0.34	−0.38	−1.18	−11.81	−0.49 [−0.82, −0.36]; 30/30	−0.01 †
Budget first minus floor first (4q): global welfare	+0.61	+0.10	+0.29	+0.62	+0.94	+0.42 [0.30, 0.60]; 30/30	+0.46
Admission rule: global welfare	+0.53	+0.19	+0.35	+0.72	+3.48	+0.29 [0.10, 0.62]; 30/30	+0.56
Forced retirement: global welfare	−2.61	−1.39	−1.98	−3.00	−3.37	−2.39 [−3.13, −2.03]; 30/30	+0.37
Withdrawal + hard budget + admission: global welfare	+0.12	−0.12	+0.06	−0.53	+0.54	−0.21 [−0.38, 0.18]; 7/30	+0.72
Calibration loss	0.99	2.60	1.26	1.30	1.80	0.98–1.06	–

Note. Paired mean differences over 200 seeds as a percentage of the no-policy mean of the same specification; initial softness 0.8 unless stated. The γ columns change only the share-pressure parameter. The vector column gives the median and 5–95% range across the 30 parameter vectors and the number of vectors in which the effect has the sign of the main specification; these ranges describe the set of vectors and are not confidence intervals. Variety-neutral demand subtracts $\sigma \ln(N_t/N_{90})$ from every firm's mean utility from 2025Q3 (Section 3.1). † Not significant at the 5% level; all other entries in the main, γ and variety-neutral columns are significant. Last row: calibration loss under the original ten-moment objective.

6.3. The Value of Variety and Product Substitutability

Because PV modules are close substitutes, two features of demand matter: how much buyers value the number of suppliers, and how easily they switch between suppliers. Table 3 already reports welfare without a valuation of variety. The case against the own-cost floor is unchanged: it still lowers welfare by 1.14% and lengthens involution by 111%, and hardening budgets first still beats the reverse order. The reference floor's gain, however, shrinks from 1.81% to 0.48% [0.45, 0.51], and buyer value accounts for 66% of the smaller gain. Instruments that reduce the number of firms no longer lose welfare: the hard budget's effect becomes indistinguishable from zero (−0.01% [−0.03, 0.02]), forced retirement raises welfare by 0.37% [0.33, 0.40], and the package of subsidy withdrawal, hard budgets and the admission rule raises it by 0.72%. Among the twelve policies re-run, the package and the admission rule now rank second and third in global welfare behind the demand stimulus,

ahead of the reference floor; in domestic welfare, the reference floor falls from first to second, behind the demand stimulus.

To make firms closer substitutes, we also reduce the logit scale σ by a quarter and by a half, which raises the implied firm-level price elasticity in 2023 from 3.7 to 4.9 and 7.3 (Appendix C Table A7). These variants fit the data less well (six of ten moments), mainly because the simulated gross margin falls to 6.4% and 3.3%, against 8.7% in the data. With the other parameters held at their estimates, the observed margins thus favour the calibrated degree of differentiation, although this does not identify differentiation independently of the other parameters. The direction of the results is nevertheless informative. With more substitutable products, the own-cost floor's welfare loss increases to 3.8% and 6.6%, and its effect on utilisation to −9.6 and −12.9 percentage points. Its effect on exit reverses (+30% and +62%), plausibly because buyers switch more easily, so that high-cost firms that must price at their own audited cost lose more of their customers to cheaper rivals. The reference floor's gain falls to 0.17% and then becomes a loss of 3.7%. A 30% higher price coefficient, which raises the elasticity to 4.8, gives the same pattern (own-cost floor −4.51%, reference floor +0.38%). The admission rule raises welfare in all variants. Greater product homogeneity thus strengthens the case against own-cost floors and weakens the case for a reference floor; it reverses only the floor's effect on exit. The measured welfare costs of capacity reduction in the main specification rest largely on the logit valuation of variety, whereas the case against the own-cost floor does not.

6.4. Small-Producer Behaviour

The calibration sample consists of leading survivors (Section 4.1). To check whether the conclusions depend on behaviour that is more typical of smaller producers, we re-ran the key scenarios under four parameterisations without re-estimation: rescues that favour large firms more strongly (importance exponent κ raised from 0.7 to 1.5), more aggressive pricing (minimum markup of zero for all firms and a doubled response of markups to share losses), tighter credit (leverage limit Λ reduced from 1.2 to 0.8), and all three together. Except for the importance exponent, these changes apply to all firms rather than to small firms alone, so they test whether the conclusions survive alternative behavioural parameters rather than establish that the leading firms represent the industry. Each variant reproduces eight of the ten moments, as the main specification does, with calibration losses of 0.99–1.01 against 0.99, so the data do not distinguish them from the main estimate (Appendix C Table A7). In all four, the own-cost floor lowers global welfare (by 0.80–1.11%), lowers utilisation (by 5.6–5.9 points) and lengthens involution (by 70–135%); it slows exit in three variants and leaves it unchanged when rescues favour large firms (−0.7%, not significant). The reference floor raises welfare (1.31–1.87%) while cutting exit by 75–81%, and the admission rule raises welfare (0.50–0.58%).

7. Discussion

7.1. Policy Implications

The cost standard determines what a ban on below-cost sales does. A variable cost standard, the natural reading of a rule against predatory dumping, hardly ever binds in a downturn driven by input deflation and fixed costs. A standard based on each firm's own full cost binds often, but it guarantees each firm the recovery of its own audited cost: it reshuffles market shares, invites further investment and lowers utilisation and welfare in all 30 parameter vectors that fit the data comparably, under variety-neutral demand and more strongly when products are more substitutable. An industry-reference standard stabilises prices and keeps firms alive, but it cuts exit by three quarters or more; most of its measured welfare gain disappears when firm variety is not valued or products are close

substitutes, and coordinated minimum prices raise competition law concerns. None of the three standards clears excess capacity.

Hard budgets are necessary but not sufficient. Removing rescues and guarantees shortens involution by 36% and raises utilisation, at a modest welfare cost that vanishes when variety is not valued. It does not, however, turn a cost-based floor into a clearing device: under hard budgets the floor substitutes for the rescues that have been removed. If floors are used at all, budgets should be hardened first, but this advice depends on how R&D responds to profits, a behavioural parameter that the data pin down only weakly.

Capacity-side instruments address the root. The Shapley decomposition traces involution mainly to share-driven investment, followed by product homogeneity and soft budgets; competition among regions for growth adds little once the mean subsidy is fixed. Consistent with this diagnosis, an admission standard that limits new capacity to near-frontier technology raises utilisation and welfare at every share-pressure value, in every parameter vector, under variety-neutral demand and under the small-producer and substitutability variants, and it raises utilisation under faster demand growth as well. Its welfare gain requires that demand plateaus, as in the baseline: if demand grows as fast as industry projections suggest, blocking non-frontier capacity forgoes valuable supply, so the value of the rule depends on how well regulators anticipate demand. Forced retirement, which represents coordinated capacity reduction and managed exit, achieves the largest reduction in high-cost capacity; its large measured welfare cost reflects the logit value of the firms it removes and becomes a small gain under variety-neutral demand. Combining subsidy withdrawal, hard budgets and an admission standard shortens involution by about two thirds, but its welfare gain is fragile: the global gain disappears in most parameter vectors, and the domestic gain in six alternative specifications. Instruments that change the returns to expansion therefore look more promising than instruments that set prices.

Other capacity-side instruments discussed in policy are not modelled, but the results bear on them. Production quotas would stabilise prices in the short run but share a central weakness of price floors: they freeze market shares and protect high-cost capacity, and they require more information than an admission standard. Mergers and capacity buy-outs could reduce capacity with compensation; modelling them requires an explicit market for assets and acquisition prices, which we leave for future work. Our uncompensated forced retirement is an upper bound on the cost of such schemes to firm owners.

The results also matter for firms, lenders and buyers. For firms, floors protect margins in the short run but induce investment and lower utilisation, so a floor is not a safe basis for expansion plans. For lenders, floors postpone rather than remove credit risk: by keeping indebted firms alive, the reference floor and the admission rule lower bank surplus by 256 and 135 model units (Appendix C Table A5), so lenders should not treat a price floor as credit enhancement. For buyers, floors raise module prices—by 7.4% under the reference floor—and lower consumer surplus under every cost standard, which matters for project developers who sign long-term supply contracts.

Who gains matters. With roughly 40% of modules sold abroad, a price war transfers part of the industry's losses to foreign buyers. Measured on a domestic basis, price-raising instruments look better and price-lowering ones worse. Policy evaluations of involution should therefore state whose welfare they count.

7.2. Limitations

Six limitations stand out. First, the calibration targets come from six leading listed firms; the model reproduces only part of the 2024 price decline and not the amplitude of the utilisation cycle, and parameter recovery shows weak identification. Section 6 shows how the results vary with the parameters, but those ranges are descriptive, not a posterior.

Second, share pressure, the dominant mechanism, is fixed rather than estimated, and the floor's effect on involution duration changes sign when share pressure is weak. Third, the demand parameters are assumed: logit demand treats each firm as a distinct option, and no public transaction-level data exist from which a demand system for utility-scale module purchases could be estimated. The variety-neutral and substitutability checks show that welfare comparisons between instruments that change the number of firms depend on how variety is valued and how close substitutes modules are. Fourth, without policy, the calibrated industry never returns to profitability within the horizon, so the long-run results are comparisons within a stylised scenario of persistent involution, not forecasts. Fifth, the domestic welfare measure assumes that foreign buyers enjoy the same surplus per unit as domestic buyers; global welfare does not depend on this assumption. Sixth, the model is partial equilibrium. Polysilicon prices are exogenous, although they were themselves affected by the 2025 policy campaign: the input path includes the observed rebound through 2025Q4 and is held flat afterwards, and two checks vary it. Because the floors are defined on audited cost, coordinated upstream price increases would pass through into the floors and amplify their price effects. The model also omits trade responses abroad—floors that raise export prices may reduce exposure to anti-dumping measures but also encourage substitution towards foreign capacity—and workers beyond an adjustment cost, and its floors are stylised versions of rules that in practice contain intent conditions, exceptions and enforcement costs.

7.3. Implications for PV Deployment and Environmental Sustainability

The policy ranking also bears on the environmental goals that PV manufacturing serves. Instruments that raise module prices raise the cost of new installations. The reference floor raises the simulated module price by 7.4%. With a global weighted-average installed cost of utility-scale PV of USD 691/kW in 2024 [64] and module prices of roughly USD 0.12/W in that year (the domestic spot average of 0.84 yuan/W at about 7.2 yuan per dollar), modules account for about one sixth of installed cost, so the floor would raise installed costs by roughly 1.2%, before any effect on financing costs. The simulated total sales hardly change, because the outside-option share of the calibrated demand is small; this understates how deployment responds to prices and should not be read as evidence that deployment is insensitive to them. From a deployment perspective, instruments that lower prices without increasing idle capacity are therefore preferable.

Excess capacity has environmental costs of its own. Idle plants embody capital, materials and energy that yield no output, and price wars that bankrupt firms can reduce the R&D on which further efficiency gains depend. A policy that stabilises prices without clearing capacity, such as the reference floor, preserves these costs. These channels point in the same direction as the economic results for the floors, but we do not quantify them: the model contains no emissions, material flows or life-cycle accounting. Linking the industry model to deployment and life-cycle emission models is a natural extension.

8. Conclusions

China's anti-involution policy has turned to prices: a ban on sales below cost is meant to end destructive price wars. Our simulations suggest that such a ban cannot clear an industry whose excess capacity comes from share-driven investment and soft budgets. A ban based on variable cost almost never binds. A ban based on each firm's own full cost lowers utilisation and welfare whether budgets are soft or hard, whether or not variety is valued, and more strongly when products are close substitutes; in the main specification it slows exit and acts as a soft budget constraint financed by buyers. A floor tied to an industry reference cost stabilises prices but keeps firms alive, and most of its measured welfare gain

disappears when the survival of more firms is not valued as additional variety. An admission standard for new capacity, by contrast, raises utilisation and, while demand plateaus, welfare. These are comparisons within a calibrated but stylised scenario of persistent involution. An extension to a demand system estimated for near-homogeneous products, and to endogenous upstream prices, would help settle how capacity instruments compare with an industry-reference floor.

Author Contributions: Conceptualization, C.C. and J.Z.; methodology, C.C.; software, C.C.; formal analysis, C.C.; data curation, C.C.; writing—original draft preparation, C.C.; writing—review and editing, C.C. and J.Z.; visualisation, C.C.; supervision, J.Z. J.Z. also contributed to the interpretation of the results. All authors have read and agreed to the published version of the manuscript.

Funding: This research received no external funding.

Institutional Review Board Statement: Not applicable.

Informed Consent Statement: Not applicable.

Data Availability Statement: The study uses public company reports and public industry data. The replication package, which includes model code, configuration files, random seeds, processed public data, per-seed results and scripts for the tables and figures, has been made available to the reviewers and will be deposited in a public repository upon acceptance.

Conflicts of Interest: The authors declare no conflicts of interest.

Appendix A. ODD Summary

Purpose and patterns. The model evaluates cost-based price floors, budget hardening and capacity instruments, and their welfare incidence, in a manufacturing industry with excess capacity. It is calibrated to the 2023–2025 downturn of China's PV module industry: a sharp price fall in 2024, widespread losses, thin gross margins, utilisation of about two-thirds to three-quarters, and a stable R&D intensity.

Entities, state variables and scales. Firms are characterised by their region, capacity in use, capacity under construction by completion date, processing-cost base, cumulative output, quality, markup, cash, debt, net worth, market share and the number of consecutive loss quarters. Regions hold a fiscal balance and an investment-subsidy rate. A passive bank sector records interest, default losses and equity. Demand is represented by 4096 fixed consumer draws with type-I extreme-value preferences over firms and an outside option. The time step is a quarter; runs last 160 quarters; all quantities are in normalised units.

Process overview and scheduling. See the eight steps in Section 3.

Design concepts. Firms are boundedly rational: they adjust markups adaptively, extrapolate demand growth and target market shares that are mutually inconsistent in aggregate. Local governments compete for growth under fiscal constraints and rescue locally important firms. Overcapacity, losses, zombie survival, exit and welfare emerge from these interactions; involution is never imposed as an objective. Stochastic elements (demand shocks, entry, innovation, rescue, enforcement and rationing) use common random numbers indexed by seed, quarter, firm and event.

Initialisation and input. The model starts with 26 firms with log-normally distributed sizes and heterogeneous costs and quality, and initial capacity equal to initial demand divided by the target utilisation of 0.85. The only external time series is the silicon cost, built from the 2020Q1–2025Q4 polysilicon price index and silicon use per watt, and held flat after 2025Q4; the future-demand checks of Section 6.1 add a demand multiplier.

Submodels. Sections 3.1–3.7.

Verification. Forty-two unit tests cover cash, balance-sheet, fiscal and bank accounting, the conservation of capacity in use and under construction, the welfare decomposition,

market shares and rationing. Without rationing, simulated consumer surplus differs from the analytical log-sum by less than 1%. Every run is audited; the largest accounting residual across all runs is of the order of 10^{-12}. For this revision, the frozen code was re-run in a new environment and reproduced the stored per-seed results to within 10^{-11}.

Appendix B. Parameters and Notation

Table A1. Main parameter values.

Block	Parameter	Symbol	Value	Source
Demand	Maximum potential demand; growth rate; midpoint quarter	$M_{\max}$; g_M; t_0	500; 0.075; 48	Assumption
Demand	AR(1) coefficient; shock standard deviation	χ; σ_ξ	0.65; 0.025	Assumption
Demand	Utility intercept; price coefficient; quality coefficient; logit scale	a; α; β; σ	2.5; 2.0; 0.4; 0.35	Assumption
Costs	Silicon use per watt, 2020–2023 (g/W)	m_t	3.1; 2.7; 2.3; 2.2	IEA PVPS [51,52]
Costs	Silicon share of module unit COGS, 2021Q4	–	47.9%	Jinko Solar [53], InfoLink
Costs	Dispersion of initial costs	–	0.058	Estimated
Costs	Unit cost of capital	–	1.734	Estimated
Costs	Fixed operating cost per unit of capacity	–	0.070	Preliminary calibration
Costs	Learning exponent; lower bound of processing cost	ν; $k_{\min}$	0.058; 0.2	Preliminary calibration; assumption
Costs	Book depreciation; physical retirement (per quarter)	–	2.5%; 0.5%	Assumption
Pricing	Markup response to share growth; to utilisation gap	η_s; η_u	0.045; 0.020	Assumption; preliminary calibration
Pricing	Markup bounds; minimum markup when cash-rich	$\mu_{\min}$–$\mu_{\max}$	1.5–65%; −8%	Assumption
Investment	Target utilisation; share pressure; construction lag	u^*; γ; L	0.85; 0.10; 4 quarters	Assumption
Investment	Extrapolation; adjustment speed	ϕ; ι	0.207; 0.726	Estimated
R&D	Revenue share; profit share; innovation efficiency	ρ_0; ρ_1; ζ	0.0166; 0.00036; 18.79	Estimated
Finance	Leverage limit; guarantee boost; interest rate (quarterly); guarantee discount	Λ; ω; r; d	1.2; 0.8; 2.5%; 0.35	Assumption
Finance	Baseline budget softness; importance exponent	S; κ	0.8; 0.7	Assumption
Government	Mean subsidy rate; maximum subsidy rate	–	0.058; 0.30	Preliminary calibration; assumption
Government	Corporate tax rate; land-revenue decline from quarter 60 (per quarter)	–	20%; 4%	Assumption
Welfare	Discount factor (quarterly); cost of public funds	δ; λ_g	0.99; 0.2	Assumption
Welfare	Bankruptcy loss (share of assets); adjustment cost per unit of exiting capacity	–	20%; 0.1	Assumption
Policy	Detection probability; fine multiple	θ; ψ	0.25; 1.0	Chosen to target violation rates
Policy	Horizon of future-share value; liquidity shadow value	–	8 quarters; 0.25	Assumption
Policy	Admission gap to frontier; forced retirement share; demand stimulus (size, duration, real cost)	–	10%; 15%; 15%, 8 quarters, 0.05	Assumption

Note. "Estimated" refers to the simulated method of moments in Section 4.3; "preliminary calibration" to values from an earlier calibration of the model that were held fixed. All other values are design assumptions, varied in the robustness checks where indicated in Section 6. Other symbols: M_t potential demand; ξ_t demand shock; V_{it} mean utility; q_i quality; ε_{ij} taste draws; c_{it} unit variable cost; P_t polysilicon price index; k_{it} processing cost; b_{it} processing-cost base; Q_{it} cumulative output; μ_{it} markup; s_{it} market share; u utilisation; h_t average log demand growth; K^*_{it} target capacity; K^{pipe}_{it} capacity under construction; I_{it} investment; E_{it} net worth; π_{it} profit before R&D and taxes; N_t number of selling firms; W welfare.

Appendix C. Additional Results

Table A2. Shapley contributions of seven channels (40 quarters from 2025Q3).

Channel	Composite Index	Involution Quarters	High-Cost Capacity (pp)
Subsidy level	0.0341 [0.0329, 0.0352]	1.67 [1.50, 1.85]	0.52 [0.40, 0.64]
Subsidy tournament	−0.0009 [−0.0012, −0.0007]	−0.15 [−0.23, −0.07]	0.00 [−0.02, 0.02]
Soft budget	0.0387 [0.0375, 0.0399]	2.41 [2.21, 2.60]	0.52 [0.48, 0.57]
Product homogeneity	0.0597 [0.0583, 0.0612]	4.15 [3.96, 4.35]	1.48 [1.28, 1.69]
Extrapolative expectations	0.0258 [0.0248, 0.0267]	−0.03 [−0.22, 0.16]	0.77 [0.68, 0.86]
Construction lag	−0.0564 [−0.0595, −0.0533]	3.37 [3.07, 3.67]	−4.66 [−5.00, −4.31]
Share pressure	0.2650 [0.2620, 0.2680]	4.66 [4.30, 5.02]	1.91 [1.69, 2.13]

Note. Shapley contributions over 128 channel combinations × 200 seeds, with 95% confidence intervals. Positive values raise the outcome. Longer construction lags lower the composite index and the high-cost share, because the loss-making share and persistent-loss capacity fall by 8.5 and 8.9 percentage points.

Table A3. Policy effects over 40 quarters with 95% confidence intervals.

Instrument	Utilisation (pp)	High-Cost Capacity (pp)	Exit Capacity (%)	Involution Quarters (%)	Global Welfare (%)	Domestic Welfare (%)	Net Fiscal Spending (Units)
Hard budget	0.78 [0.72, 0.84]	−0.47 [−0.64, −0.31]	−1.5 [−3.0, −0.1]	−36.0 [−41.4, −30.5]	−0.53 [−0.57, −0.48]	−0.39 [−0.44, −0.34]	−8.0 [−8.5, −7.4]
Variable cost floor	0.00 [0.00, 0.00]	0.00 [0.00, 0.00]	0.0 [0.0, 0.0]	0.0 [0.0, 0.0]	0.00 [0.00, 0.00]	0.00 [0.00, 0.00]	0.0 [0.0, 0.0]
Own-cost floor	−5.79 [−6.01, −5.56]	0.38 [0.16, 0.59]	−2.3 [−4.5, −0.1]	111.7 [102.1, 121.4]	−1.09 [−1.16, −1.03]	−0.80 [−0.87, −0.73]	−0.3 [−0.9, 0.4]
Reference floor	−6.50 [−6.72, −6.28]	0.65 [0.38, 0.91]	−75.2 [−78.0, −72.3]	−52.7 [−61.4, −44.1]	1.81 [1.74, 1.88]	3.26 [3.17, 3.35]	−32.2 [−33.0, −31.4]
Hard budget + own-cost floor	−4.00 [−4.18, −3.82]	−0.22 [−0.42, −0.01]	−5.3 [−7.2, −3.3]	70.1 [60.3, 79.9]	−1.30 [−1.36, −1.23]	−0.90 [−0.97, −0.82]	−8.9 [−9.5, −8.3]
Hard budget + reference floor	−5.69 [−5.88, −5.51]	0.43 [0.21, 0.66]	−72.3 [−75.1, −69.5]	−62.2 [−70.8, −53.6]	1.69 [1.62, 1.76]	3.13 [3.04, 3.23]	−34.6 [−35.3, −33.9]
Subsidy withdrawal	0.18 [0.15, 0.21]	−0.16 [−0.24, −0.07]	−4.4 [−5.4, −3.4]	−7.6 [−11.3, −3.8]	−0.05 [−0.07, −0.03]	0.05 [0.02, 0.07]	−25.7 [−26.4, −25.1]
Gradual subsidy withdrawal	0.18 [0.15, 0.21]	−0.13 [−0.20, −0.05]	−2.8 [−3.8, −1.8]	−5.9 [−9.4, −2.4]	−0.04 [−0.06, −0.02]	0.03 [0.01, 0.06]	−23.8 [−24.4, −23.3]
Admission rule	3.79 [3.66, 3.92]	−0.26 [−0.46, −0.06]	−9.3 [−10.7, −7.9]	−29.6 [−34.8, −24.5]	0.53 [0.50, 0.57]	1.04 [1.00, 1.08]	−20.0 [−20.9, −19.2]
Forced retirement	3.15 [3.07, 3.23]	−7.27 [−7.75, −6.80]	63.9 [62.4, 65.5]	−71.6 [−77.4, −65.8]	−2.61 [−2.65, −2.56]	−2.83 [−2.89, −2.77]	4.7 [4.0, 5.4]
Demand stimulus	−0.45 [−0.56, −0.35]	−0.47 [−0.63, −0.31]	16.2 [14.4, 18.0]	20.9 [12.5, 29.4]	2.84 [2.79, 2.88]	2.65 [2.59, 2.71]	30.5 [29.9, 31.1]
Withdrawal + hard budget	0.91 [0.85, 0.98]	−0.53 [−0.71, −0.36]	−4.7 [−6.3, −3.2]	−41.7 [−47.3, −36.2]	−0.54 [−0.59, −0.50]	−0.32 [−0.37, −0.27]	−33.8 [−34.6, −32.9]
Withdrawal + hard budget + admission	5.77 [5.58, 5.95]	−0.97 [−1.19, −0.75]	−11.8 [−14.0, −9.6]	−68.8 [−74.7, −62.9]	0.12 [0.05, 0.19]	0.95 [0.87, 1.02]	−55.1 [−57.0, −53.3]
Withdrawal + hard budget + own-cost floor + admission	1.97 [1.71, 2.24]	−0.99 [−1.27, −0.71]	−17.4 [−19.5, −15.3]	−85.7 [−92.1, −79.3]	−0.51 [−0.60, −0.43]	0.56 [0.47, 0.65]	−69.2 [−70.7, −67.6]

Note. Paired mean differences from the no-policy run over 200 seeds; initial softness 0.8; logit demand. Domestic welfare uses the low-export scenario. Net fiscal spending (subsidies, rescues and stimulus outlays minus taxes) is a discounted sum in model units; its no-policy level is 25.0. Combining the variable cost floor with a hard budget gives exactly the hard-budget results.

Table A4. Sequencing comparisons (40 quarters, own-cost floor, initial softness 0.8).

Comparison	Global Welfare	Domestic Welfare, Low Export	Domestic Welfare, High Export
Budget first (4q) minus floor first	0.61% [0.54, 0.68]	0.73% [0.65, 0.82]	0.76% [0.67, 0.84]
Budget first (4q) minus simultaneous	0.26% [0.22, 0.30]	0.28% [0.22, 0.33]	0.28% [0.23, 0.34]
Budget first (8q) minus floor first	0.81% [0.73, 0.90]	0.97% [0.88, 1.07]	1.00% [0.90, 1.10]
Budget first (8q) minus simultaneous	0.42% [0.37, 0.47]	0.44% [0.38, 0.50]	0.44% [0.38, 0.50]
Gradual minus simultaneous	−0.31% [−0.36, −0.26]	–	–

Note. Paired differences as a percentage of no-policy welfare, with 95% confidence intervals. Positive values favour the first-named arm. In sequences, the leading instrument is in force for 40 quarters and the lagging one for 36 or 32. In the gradual version, softness falls linearly to zero and the detection probability rises from one eighth of its full value to the full value over eight quarters.

Table A5. Surplus by actor over 40 quarters (model units).

Instrument	Consumers	Firm Owners	Local Governments	Banks	Real Losses	Net Fiscal Spending
Hard budget	−181.8	84.3	8.0	−34.0	4.4	−8.0
Own-cost floor	−385.8	155.0	4.2	−39.0	−3.0	−0.3
Reference floor	−122.4	708.3	51.4	−255.6	−45.6	−32.2
Hard budget + own-cost floor	−476.2	220.4	11.4	−67.7	0.5	−8.9
Admission rule	−65.1	297.5	20.0	−135.0	−6.6	−20.0
Forced retirement	−556.4	−173.3	−4.7	214.1	103.8	4.7
Demand stimulus	772.8	−58.3	−30.5	11.3	9.0	30.5

Note. Paired mean differences from no policy; logit demand. Owner and bank surplus include changes in opening and closing stocks of cash, debt and capital, so they should not be read as operating profits. Positive real losses mean larger losses. Global welfare equals consumers + owners + local governments + banks − real losses − 0.2 × net fiscal spending.

Table A6. Agreement with the main results across 38 alternative specifications.

Contrast	Main Effect	Same Direction	Same Direction and Significance	Significantly Opposite
Own-cost floor: global welfare	−1.09% [−1.16, −1.03]	36/38	35/38	2
Reference floor: global welfare	+1.81% [1.74, 1.88]	38/38	38/38	0
Hard budget: global welfare	−0.53% [−0.57, −0.48]	38/38	38/38	0
Admission rule: global welfare	+0.53% [0.50, 0.57]	36/38	36/38	2
Forced retirement: global welfare	−2.61% [−2.65, −2.56]	38/38	38/38	0
Demand stimulus: global welfare	+2.84% [2.79, 2.88]	38/38	38/38	0
Package: global welfare	+0.12% [0.05, 0.19]	28/38	28/38	9
Package: domestic welfare	+0.95% [0.87, 1.02]	32/38	32/38	6
Joint tool: global welfare	−1.30% [−1.36, −1.23]	38/38	38/38	0
Joint tool: domestic welfare	−0.90% [−0.97, −0.82]	37/38	37/38	1
Soft-budget floor: high-cost capacity	+1.57% [0.74, 2.40]	32/38	29/38	2
Soft-budget floor: exit capacity	−3.58% [−6.04, −1.13]	32/38	29/38	5
Hard-budget floor: high-cost capacity	+1.20% [0.46, 1.94]	35/38	31/38	1
Hard-budget floor: exit capacity	−6.88% [−9.48, −4.28]	33/38	32/38	2
Budget first minus reverse, 4q: global welfare	+0.61% [0.54, 0.68]	36/38	36/38	2
Budget first minus reverse, 4q: domestic welfare	+0.73% [0.65, 0.82]	36/38	36/38	1
Joint tool minus hard budget: consumer surplus, 8q	−0.73% [−0.77, −0.68]	38/38	37/38	0

Table A6. *Cont.*

Contrast	Main Effect	Same Direction	Same Direction and Significance	Significantly Opposite
Joint tool minus hard budget: consumer surplus, 40q	−1.19% [−1.25, −1.14]	36/38	36/38	1
Joint tool minus hard budget: R&D, 40q	+4.86% [4.62, 5.09]	38/38	38/38	0

Note. Effects are percentages of the no-policy benchmark of each check, with 95% confidence intervals for the main specification; unless stated otherwise, effects are relative to no policy at initial softness 0.8. A check agrees when the effect has the same sign and the same significance status; unchanged non-significance also counts as agreement. The checks are not independent draws, so agreement rates are not probabilities that a result is true. Soft- and hard-budget floor: the own-cost floor at initial softness 1.0 and 0. Joint tool: the hard budget combined with the own-cost floor. Package: subsidy withdrawal, hard budget and admission rule combined.

Table A7. Key contrasts under small-producer behaviour and higher product substitutability (40 quarters, % of no-policy mean unless stated).

Specification	Moments Reproduced (Loss)	Implied Elasticity, 2023	Own-Cost Floor: Welfare	Own-Cost Floor: Utilisation (pp)	Own-Cost Floor: Exit	Reference Floor: Welfare	Hard Budget: Welfare	Admission Rule: Welfare	Forced Retirement: Welfare
Main	8/10 (0.99)	3.7	−1.09	−5.79	−2.3	+1.81	−0.53	+0.53	−2.61
Rescue favours large firms ($\kappa = 1.5$)	8/10 (0.99)	3.7	−1.11	−5.60	−0.7 †	+1.87	−0.43	+0.55	−2.68
Aggressive pricing ($\mu_{\min} = 0$; η_s doubled)	8/10 (1.01)	3.7	−1.10	−5.85	−4.7	+1.40	−0.46	+0.51	−2.18
Tighter credit ($\Lambda = 0.8$)	8/10 (1.00)	3.7	−0.80	−5.80	−14.2	+1.77	−0.75	+0.58	−2.47
All three small-producer changes	8/10 (1.01)	3.7	−0.97	−5.72	−13.0	+1.31	−0.54	+0.50	−2.16
Logit scale $\sigma \times 0.75$	6/10 (1.11)	4.9	−3.83	−9.59	+30.1	+0.17	−0.42	+0.62	−2.02
Logit scale $\sigma \times 0.50$	6/10 (2.91)	7.3	−6.61	−12.88	+61.6	−3.74	−0.31	+0.73	−1.04
Price coefficient $\alpha + 30\%$	7/10 (1.07)	4.8	−4.51	–	+31.5	+0.38	−0.58	+0.84	−2.56

Note. Paired mean differences over 200 seeds, initial softness 0.8, logit demand, global welfare; parameters changed without re-estimation; except for κ, the small-producer changes apply to all firms. Moments reproduced: number of the ten calibration moments that meet the pre-registered criterion, with the calibration loss in parentheses. Implied elasticity: sales-weighted firm-level own-price elasticity of unconstrained logit demand, $(\alpha/\sigma)p_{it}(1-s_{it})$. The price-coefficient row corresponds to check R6 in Table A6. † Not significant at the 5% level; all other entries are significant.

Table A8. Effects over 8, 20 and 40 quarters.

Instrument	Global Welfare, 8q (%)	Global Welfare, 20q (%)	Global Welfare, 40q (%)	Utilisation, 8q (pp)	Utilisation, 40q (pp)	Exit Capacity, 8q (%)	Exit Capacity, 40q (%)
Own-cost floor	−0.77	−0.67	−1.09	−0.44	−5.79	−17.4	−2.3
Reference floor	+0.60	+1.28	+1.81	−0.81	−6.50	−66.5	−75.2
Hard budget	−1.71	−0.65	−0.53	+3.22	+0.78	+101.0	−1.5
Hard budget + own-cost floor	−1.97	−1.10	−1.30	+2.26	−4.00	+73.0	−5.3
Admission rule	+0.42	+0.55	+0.53	+0.26	+3.79	−9.6	−9.3
Forced retirement	−4.35	−2.78	−2.61	+9.29	+3.15	+328.4	+63.9
Subsidy withdrawal	−0.18	−0.09	−0.05	+0.23	+0.18	+15.9	−4.4
Withdrawal + hard budget + admission	−0.87	+0.05 †	+0.12	+3.60	+5.77	+85.5	−11.8
Demand stimulus	+14.57	+5.68	+2.84	+7.97	−0.45	+21.3	+16.2

Note. Paired mean differences from the no-policy run over 200 seeds, initial softness 0.8, logit demand; each window starts in 2025Q3 and percentages refer to the no-policy mean of the same window. † Not significant at the 5% level; all other entries are significant.

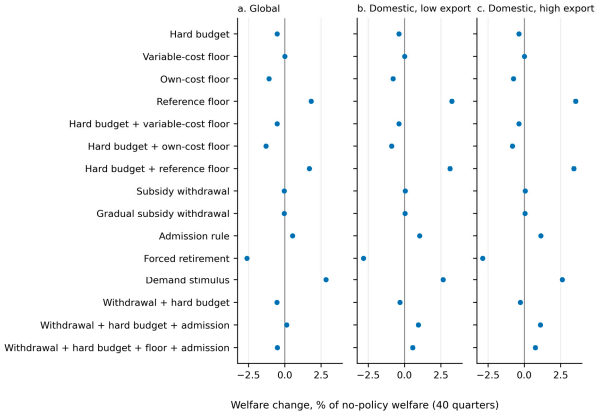


Figure A1. Welfare effects of all instruments after 40 quarters as a percentage of no-policy welfare, for (**a**) global welfare and for domestic welfare in (**b**) the low- and (**c**) the high-export scenarios under logit demand. Bars (mostly hidden by the markers) are 95% confidence intervals.

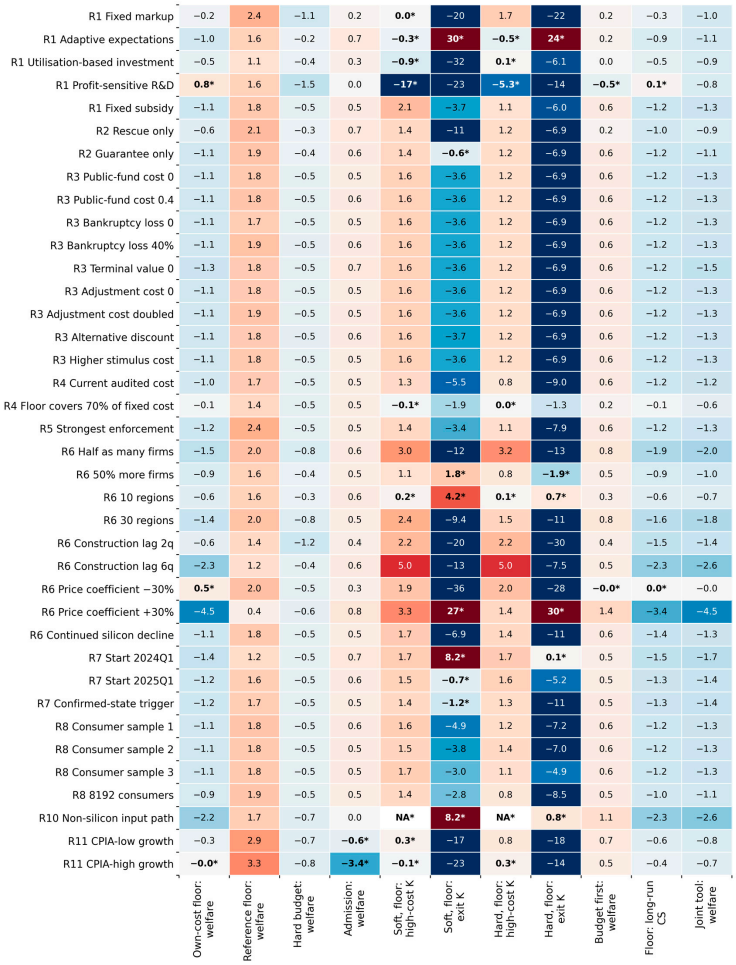


Figure A2. Relative effects (% of each check's no-policy benchmark) of eleven contrasts in the 38 alternative specifications. Rows are grouped by type of check: R1 behavioural rules, R2 soft-budget channels, R3 welfare parameters, R4 cost basis of the floor, R5 enforcement, R6 structure, R7 policy timing, R8 consumer samples, R10 non-silicon input path and R11 future demand. Soft and hard: the own-cost floor at initial softness 1.0 and 0 (K: capacity). Budget first: hardening budgets four quarters before the floor minus the reverse order. Floor: long-run CS: the hard budget combined with the own-cost floor minus the hard budget alone, for consumer surplus after 40 quarters. Joint tool: the hard budget combined with the own-cost floor minus no policy. An asterisk and bold type mark a change in sign or significance relative to the main specification; NA marks undefined percentages. Colours are truncated at the 90th percentile of absolute values.

References

1. Central People's Government of the PRC. Report on the Sixth Meeting of the Central Financial and Economic Affairs Commission. Available online: https://www.gov.cn/yaowen/liebiao/202507/content_7030285.htm (accessed on 13 October 2026). (In Chinese)
2. National Development and Reform Commission. Draft Amendment to the Price Law of the People's Republic of China (Draft for Public Comment). Available online: https://yyglxxbs.ndrc.gov.cn/file-submission/20250724170553766368.pdf (accessed on 13 October 2026). (In Chinese)
3. McFadden, D. A Method of Simulated Moments for Estimation of Discrete Response Models Without Numerical Integration. *Econometrica* **1989**, *57*, 995–1026. [CrossRef] [PubMed]
4. Carranza, J.E.; Clark, R.; Houde, J.-F. Price Controls and Market Structure: Evidence from Gasoline Retail Markets. *J. Ind. Econ.* **2015**, *63*, 152–198. [CrossRef]
5. Allain, M.-L.; Chambolle, C. Anti-Competitive Effects of Resale-Below-Cost Laws. *Int. J. Ind. Organ.* **2011**, *29*, 373–385. [CrossRef]
6. Lin, J.Y. The Wave Phenomenon and the Reconstruction of Macroeconomic Theory for Developing Countries. *Econ. Res. J.* **2007**, *1*, 126–131. (In Chinese)
7. Lin, J.Y.; Wu, H.; Xing, Y. The "Wave Phenomenon" and the Formation Mechanism of Excess Capacity. *Econ. Res. J.* **2010**, *10*, 4–19. (In Chinese)
8. Geng, Q.; Jiang, F.; Fu, T. Policy Subsidies, Overcapacity and China's Economic Fluctuations: An Empirical Test of an RBC Model with Capacity Utilisation. *China Ind. Econ.* **2011**, *5*, 27–36. (In Chinese)
9. Jiang, F.; Geng, Q.; Lü, D.; Li, X. Regional Competition, Institutional Distortions and the Formation Mechanism of Overcapacity. *China Ind. Econ.* **2012**, *6*, 44–56. (In Chinese)
10. Zhang, Q.; Dong, Y. Incremental Process Innovation, Capacity Construction Cycles and Overcapacity: A Simulation Analysis Based on a Neo-Schumpeterian Evolutionary Model. *Economist* **2014**, *8*, 33–42. (In Chinese)
11. Zhang, Q.; Zhao, J.; Zhou, D. Can the Cancellation of Government Subsidies Alleviate the Phenomenon of Overcapacity in the Photovoltaic Module Industry? From a Dynamic Perspective. *Environ. Dev. Sustain.* **2024**, *26*, 6419–6441. [CrossRef] [PubMed]
12. He, X.-Z.; Westerhoff, F.H. Commodity Markets, Price Limiters and Speculative Price Dynamics. *J. Econ. Dyn. Control* **2005**, *29*, 1577–1596. [CrossRef]
13. Hernández, C.E.; Cantillo-Cleves, S. A Toolkit for Setting and Evaluating Price Floors. *J. Public Econ.* **2024**, *232*, 105084. [CrossRef]
14. Kornai, J. The Soft Budget Constraint. *Kyklos* **1986**, *39*, 3–30. [CrossRef]
15. Kornai, J.; Maskin, E.; Roland, G. Understanding the Soft Budget Constraint. *J. Econ. Lit.* **2003**, *41*, 1095–1136. [CrossRef]
16. Caballero, R.J.; Hoshi, T.; Kashyap, A.K. Zombie Lending and Depressed Restructuring in Japan. *Am. Econ. Rev.* **2008**, *98*, 1943–1977. [CrossRef]
17. Shen, G.; Chen, B. Zombie Firms and over-Capacity in Chinese Manufacturing. *China Econ. Rev.* **2017**, *44*, 327–342. [CrossRef]
18. State Council of the PRC. Regulations on Fair Competition Review (State Council Decree No. 783). Available online: https://www.gov.cn/zhengce/content/202406/content_6957049.htm (accessed on 13 October 2026). (In Chinese)
19. NPC Standing Committee. Anti-Unfair Competition Law of the People's Republic of China (Revised 27 June 2025). Available online: https://flk.npc.gov.cn/law-search/search/flfgDetails?bbbs=ff8081819715552b40197b1016efc5437 (accessed on 13 October 2026). (In Chinese)
20. Li, H.; Zhou, L.-A. Political Turnover and Economic Performance: The Incentive Role of Personnel Control in China. *J. Public Econ.* **2005**, *89*, 1743–1762. [CrossRef]
21. Corwin, S.; Johnson, T.L. The Role of Local Governments in the Development of China's Solar Photovoltaic Industry. *Energy Policy* **2019**, *130*, 283–293. [CrossRef]
22. Yu, S.; Lu, T.; Hu, X.; Liu, L.; Wei, Y.-M. Determinants of Overcapacity in China's Renewable Energy Industry: Evidence from Wind, Photovoltaic, and Biomass Energy Enterprises. *Energy Econ.* **2021**, *97*, 105056. [CrossRef]
23. Chen, J.; Wang, T. Government Subsidies, R&D Expenditures and Overcapacity: Empirical Analysis in Photovoltaic Companies. *Chin. Manag. Stud.* **2023**, *17*, 343–364. [CrossRef]
24. Ge, P.; Zhu, R.; Chen, Y.; Huang, X. Selective Industrial Policy and Overcapacity: Evidence from a Quasi-Natural Experiment in China. *Econ. Syst.* **2024**, *48*, 101191. [CrossRef]
25. Liu, B.; Lai, J. Can Breaking Administrative Monopoly Reduce Overcapacity? A Quasi-Natural Experiment Based on the Fair Competition Review System. *J. Financ. Econ.* **2021**, *47*, 34–47. (In Chinese) [CrossRef]
26. Barwick, P.J.; Kalouptsidi, M.; Zahur, N.B. *China's Industrial Policy: An Empirical Evaluation*; NBER Working Paper No. 26075; National Bureau of Economic Research: Cambridge, MA, USA, 2019.
27. Barwick, P.J.; Kalouptsidi, M.; Zahur, N.B. Industrial Policy: Lessons from Shipbuilding. *J. Econ. Perspect.* **2024**, *38*, 55–80. [CrossRef]
28. Goldfeld, S.M.; Quandt, R.E. Budget Constraints, Bailouts, and the Firm Under Central Planning. *J. Comp. Econ.* **1988**, *12*, 502–520. [CrossRef]
29. Zhang, C.; Chen, Y.; Zhou, H. Zombie Firms and Soft Budget Constraints in the Chinese Stock Market. *Asian Econ. J.* **2020**, *34*, 51–77. [CrossRef]

30. Chang, Q.; Zhou, Y.; Liu, G.; Wang, D.; Zhang, X. How Does Government Intervention Affect the Formation of Zombie Firms? *Econ. Model.* **2021**, *94*, 768–779. [CrossRef]
31. Dai, Y.; Li, X.; Liu, D.; Lu, J. Throwing Good Money After Bad: Zombie Lending and the Supply Chain Contagion of Firm Exit. *J. Econ. Behav. Organ.* **2021**, *189*, 379–402. [CrossRef]
32. Geng, Y.; Liu, W.; Wu, Y. How Do Zombie Firms Affect China's Industrial Upgrading? *Econ. Model.* **2021**, *97*, 79–94. [CrossRef]
33. Mao, Q.; Xu, J. Zombie Firms, Misallocation and Manufacturing Capacity Utilization Rate: Evidence from China. *Econ. Transit. Inst. Change* **2024**, *32*, 641–682. [CrossRef]
34. Fan, L.; Tang, T.; Yang, G. Administrative Reform and the Disposal of Zombie Enterprises: Evidence from China. *Small Bus. Econ.* **2023**, *61*, 1845–1869. [CrossRef] [PubMed]
35. Ma, X.; Tang, T.; Cai, G. Non-State Shareholder Governance and the de-Zombification of State-Owned Enterprises: Evidence from Mixed Boards of Listed SOEs. *J. Financ. Res.* **2021**, *3*, 95–113. (In Chinese)
36. Wang, X.; Qiao, B. "To Be or Not to Be": The Effect of Rising House Prices on Firm Zombification. *J. East China Norm. Univ. Humanit. Soc. Sci.* **2021**, *53*, 152–164. (In Chinese) [CrossRef]
37. Xiang, H.; Xiang, W. Soft Budget Constraint, Strategic Delegation and Overcapacity of Chinese SOE. *iBusiness* **2023**, *15*, 71–83. [CrossRef]
38. Bortolomiol, S.; Lurkin, V.; Bierlaire, M. Price-Based Regulation of Oligopolistic Markets Under Discrete Choice Models of Demand. *Transportation* **2022**, *49*, 1441–1463. [CrossRef] [PubMed]
39. Dosi, G.; Fagiolo, G.; Roventini, A. Schumpeter Meeting Keynes: A Policy-Friendly Model of Endogenous Growth and Business Cycles. *J. Econ. Dyn. Control* **2010**, *34*, 1748–1767. [CrossRef]
40. Dosi, G.; Fagiolo, G.; Napoletano, M.; Roventini, A. Income Distribution, Credit and Fiscal Policies in an Agent-Based Keynesian Model. *J. Econ. Dyn. Control* **2013**, *37*, 1598–1625. [CrossRef]
41. Caiani, A.; Godin, A.; Caverzasi, E.; Gallegati, M.; Kinsella, S.; Stiglitz, J.E. Agent Based-Stock Flow Consistent Macroeconomics: Towards a Benchmark Model. *J. Econ. Dyn. Control* **2016**, *69*, 375–408. [CrossRef]
42. Zhang, X.; Zhu, Q.; Li, X.; Pan, Y. The Impact of Government Subsidy on Photovoltaic Enterprises Independent Innovation Based on the Evolutionary Game Theory. *Energy* **2023**, *285*, 129385. [CrossRef]
43. Chen, Z.; Cheung, K.C.K.; Qi, X. Subsidy Policies and Operational Strategies for Multiple Competing Photovoltaic Supply Chains. *Flex. Serv. Manuf. J.* **2021**, *33*, 914–955. [CrossRef]
44. Wang, Y.; Wang, D.; Shi, X. Exploring the Dilemma of Overcapacity Governance in China's Coal Industry: A Tripartite Evolutionary Game Model. *Resour. Policy* **2021**, *71*, 102000. [CrossRef]
45. Zuo, R.; Yang, M.; Ding, Y.; Wang, S.; Jin, Y.; Wu, J. Navigating the Labyrinth of Involution: LLM-Assisted Smart Agent-Based Simulation of Social Resource Competition. *Humanit. Soc. Sci. Commun.* 2026; *in press*. [CrossRef]
46. Fagiolo, G.; Moneta, A.; Windrum, P. A Critical Guide to Empirical Validation of Agent-Based Models in Economics: Methodologies, Procedures, and Open Problems. *Comput. Econ.* **2007**, *30*, 195–226. [CrossRef]
47. Grimm, V.; Railsback, S.F.; Vincenot, C.E.; Berger, U.; Gallagher, C.; DeAngelis, D.L.; Edmonds, B.; Ge, J.; Giske, J.; Groeneveld, J.; et al. The ODD Protocol for Describing Agent-Based and Other Simulation Models: A Second Update to Improve Clarity, Replication, and Structural Realism. *J. Artif. Soc. Soc. Simul.* **2020**, *23*, 7. [CrossRef]
48. JA Solar Technology Co., Ltd. 2023 Annual Report. Available online: https://pdf.dfcfw.com/pdf/H2_AN202404291631677933_1.pdf (accessed on 13 October 2026). (In Chinese)
49. Anderson, S.P.; de Palma, A.; Thisse, J.-F. A Representative Consumer Theory of the Logit Model. *Int. Econ. Rev.* **1988**, *29*, 461–466. [CrossRef]
50. Ackerberg, D.A.; Rysman, M. Unobserved Product Differentiation in Discrete-Choice Models: Estimating Price Elasticities and Welfare Effects. *RAND J. Econ.* **2005**, *36*, 771–788.
51. IEA PVPS. Trends in Photovoltaic Applications. 2021. Available online: https://iea-pvps.org/wp-content/uploads/2022/01/IEA-PVPS-Trends-report-2021-4.pdf (accessed on 13 October 2026).
52. IEA PVPS. Trends in Photovoltaic Applications. 2024 (Japanese Edition). Available online: https://iea-pvps.org/wp-content/uploads/2025/12/Trends-2024-Japanese.pdf (accessed on 13 October 2026).
53. Jinko Solar Co., Ltd. 2021 Annual Report. Available online: https://pdf.dfcfw.com/pdf/H2_AN202204221560930246_1.pdf (accessed on 13 October 2026). (In Chinese)
54. InfoLink Consulting. 2023 Global Module Shipment Ranking. Available online: https://www.infolink-group.com/energy-article/solar-topic-2023-global-module-shipment-ranking (accessed on 13 October 2026).
55. InfoLink Consulting. 2024 Global Module Shipment Ranking: Significant Gaps Between Tiers. Available online: https://www.infolink-group.com/energy-article/solar-topic-2024-global-module-shipment-ranking-significant-gaps-between-tiers (accessed on 13 October 2026).
56. Morris, M.D. Factorial Sampling Plans for Preliminary Computational Experiments. *Technometrics* **1991**, *33*, 161–174. [CrossRef]
57. Storn, R.; Price, K. Differential Evolution—A Simple and Efficient Heuristic for Global Optimization over Continuous Spaces. *J. Glob. Optim.* **1997**, *11*, 341–359. [CrossRef]
58. Grazzini, J.; Richiardi, M.G.; Tsionas, M. Bayesian Estimation of Agent-Based Models. *J. Econ. Dyn. Control* **2017**, *77*, 26–47. [CrossRef]

59. Platt, D. A Comparison of Economic Agent-Based Model Calibration Methods. *J. Econ. Dyn. Control* **2020**, *113*, 103859. [CrossRef]
60. ten Broeke, G.; van Voorn, G.; Ligtenberg, A. Which Sensitivity Analysis Method Should I Use for My Agent-Based Model? *J. Artif. Soc. Soc. Simul.* **2016**, *19*, 5. [CrossRef]
61. Shapley, L.S. A Value for n-Person Games. In *Contributions to the Theory of Games*; Kuhn, H.W., Tucker, A.W., Eds.; Princeton University Press: Princeton, NJ, USA, 1953; Volume 2, pp. 307–318.
62. LONGi Green Energy Technology Co., Ltd. 2025 Annual Report. Available online: https://pdf.dfcfw.com/pdf/H2_AN202604281821711124_1.pdf (accessed on 13 October 2026). (In Chinese)
63. Jinko Solar Co., Ltd. 2025 Annual Report. Available online: https://pdf.dfcfw.com/pdf/H2_AN202604291821762409_1.pdf (accessed on 13 October 2026). (In Chinese)
64. IRENA. Renewable Power Generation Costs in 2024. Available online: https://www.irena.org/Publications/2025/Jun/Renewable-Power-Generation-Costs-in-2024 (accessed on 13 October 2026).

Disclaimer/Publisher's Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.