

Who Pays for Geo-economic Policy?

What company accounts, tariff studies and carbon-border contracts establish

Indiconomics Research Paper 04 · 28 September 2026

Abstract

A policy's named payer or recipient is an incomplete guide to its economic consequences. This paper compares three instruments through their commercial relationships: Indian production-linked incentives, US tariffs and European carbon-border cost clauses. Its new accounting reconstruction concerns two subsidiaries in one Indian manufacturing group. It distinguishes subsidy recognition, customer benefits and cash received, identifying a substantial customer-linked component without treating it as an estimate of final incidence. Tariff research supplies a different kind of evidence: estimated price changes with specific denominators and counterfactuals. Public carbon-border templates identify proposed reimbursement rights, but no completed payment. The comparison rejects three shortcuts: equating a gross incentive with the manufacturer's benefit, interpreting a retail price percentage as a firm's share of a tariff bill, and treating a contractual recovery right as cash recovered. It does not rank the instruments or estimate their aggregate welfare effects. The practical implication is to match each policy claim to the records that can establish it, while evaluating who benefits separately from whether the policy creates additional activity.

1 The recipient's name is the beginning of the investigation

An incentive of ₹180 crore can appear in a manufacturer's accounts without adding ₹180 crore to the revenue it reports from that incentive. Dixon Electro Appliances Private Limited provides a concrete example. For FY2025-26, Dixon's consolidated report disclosed that gross telecom incentive alongside ₹58.20 crore of other operating revenue after customer benefits. The difference is ₹121.80 crore: 67.67% of the gross amount. This is an accounting comparison, not a claim that the government paid the wrong company. [DIXON26, Note 59(c)]

The customer relationship changes what the headline amount can tell us. A reader interested in the manufacturer's direct benefit needs to follow its relationship with the customer. A reader interested in industrial development needs to go further, asking whether sharing the incentive helps secure orders, reduce prices or support investment. Those are different inquiries. The first can be advanced by a financial disclosure; the second needs evidence about behaviour and an appropriate comparison with what would otherwise have happened.

The distinction matters whenever governments try to change commercial decisions by changing prices or paying firms. Tariffs can protect producers while raising their customers' input costs. An incentive can improve a supplier's ability to win an order while leaving a large part of the direct financial benefit with the buyer. An importer can seek reimbursement for carbon-related expenses through a sales agreement. None of these possibilities makes the policy effective or ineffective by definition. Each changes the route through which its consequences must be investigated.

The question is how far public records let us trace a policy cost or benefit through commercial relationships beyond the named payer or recipient. The paper answers it through selected cases with deliberately different evidence. The main new empirical work reconstructs matched disclosures for two Dixon subsidiaries. Published tariff studies provide estimates of price adjustment. Public templates from a steel trader show how firms propose to allocate carbon-border costs. The latter sources are reused from the accompanying carbon-border investigation and are not counted as new transactions.

The comparison stops where each source stops. The manufacturer disclosures establish accounting relationships and, for one incentive cohort, cash receipt. The tariff studies estimate changes relative to specified alternatives. The templates describe rights that would take effect if agreed. Combining them reveals which conclusions survive a change in the evidence being used. It does not create a common numerical measure that can be averaged across an Indian grant, a US tariff and a European contract.

The cases were chosen for what they let us observe, not sampled to estimate how common these mechanisms are. In particular, the two Indian subsidiaries belong to one group. Their disclosures cannot establish an industry-wide sharing rate. Their value is more specific: they expose commercial links that disappear when analysis stops at the government's headline amount. Confidential customer contracts, payment histories and a new causal estimate of industrial policy remain outside its evidence.

2 Four questions that require different evidence

The economic incidence of a policy concerns its eventual effects on people and firms after prices, quantities and other decisions adjust. A legal obligation identifies who must comply with a rule. An accounting entry records an amount under a reporting framework. A cash record establishes a payment. These concepts can inform one another, but they answer different questions.

Consider a manufacturer that recognises a government incentive and a related obligation to a buyer. The accounts may identify the size of both amounts. They do not necessarily disclose when the government pays, when the buyer is paid, or whether product prices would have been different without the scheme. Even a complete payment record would not settle the last question. The contract might have transferred a benefit that the buyer would otherwise have obtained through a lower negotiated price.

Conversely, a carefully identified price response need not reveal which invoice carried a tax. A consumer price can respond to an imported input through several firms, while the customs remittance takes place at one border transaction. A lower percentage change at the retail stage may reflect a larger price base, domestic content or distribution services. It cannot automatically be interpreted as the fraction of a fixed bill absorbed earlier in the chain.

The inquiry therefore keeps four questions separate. What obligation or entitlement is documented? What associated commercial transfer is recognised or agreed? What cash has actually changed hands? What change in economic outcomes can be attributed to the policy? These are dimensions of evidence rather than a universal sequence. A price study can identify a response without observing every payment. A signed agreement can establish a legal right without proving that the right was exercised.

There is substantial prior research behind these distinctions. Goolsbee's historical investigation of investment incentives links tax changes to capital-equipment prices, demonstrating a route through which benefits can reach suppliers. Its US evidence and identification do not establish the outcome of Indian production incentives. They locate this paper's contribution appropriately: the general possibility of benefit shifting is established economics; the additional work here is a matched reconstruction of specific disclosures and their comparison with other forms of incidence evidence. [GOOLSBEE98]

The empirical method is consequently conservative about aggregation. Monetary amounts are compared only when entity, scheme, period and accounting basis match. Percentage price effects retain the price object and counterfactual used by their source. Contractual provisions retain their conditions. Unknown observations stay unknown. An absent payment date is not recorded as a late payment, and an unobserved customer transfer is not entered as zero.

An overstated manufacturer benefit can distort expectations of profitability. An unsupported financing-cost estimate can turn an administrative process into an allegation of financial distress. A price percentage mistaken for an incidence share can assign a burden to the wrong party. All three errors can survive impeccable arithmetic if the underlying measures are mismatched.

3 The customer's share of the recognised subsidy

The Dixon evidence provides two distinct comparisons. In the telecom case, the relevant amounts concern one year's incentive. In the mobile case, the figures are cumulative balances across performance years. Table 1 preserves that distinction. Its receipt row is another view of the telecom incentive, not a third grant.

Table 1. Matched subsidy disclosures and what the arithmetic establishes

Entity and basis	Matched amounts	Reconstructed comparison	Interpretation boundary
Dixon Electro Appliances; telecom PLI, FY2025-26	Gross incentive ₹180 crore; other operating revenue net of customer benefits ₹58.20 crore	Customer-benefit component ₹121.80 crore, or 67.67% of gross incentive	Accounting benefit-sharing component; not customer cash paid, net profit or final consumer benefit
Same entity and incentive cohort	85% received; ₹27 crore remaining receivable	₹153 crore received against that year's incentive	Other cohorts' receipts are excluded; this is not total incentive cash received in the year
Padget Electronics; mobile PLI, cumulative balances at 31 March 2026	Receivable ₹1,110.06 crore; corresponding customer liability ₹726.34 crore	Linked liability is 65.43% of receivable; balance difference ₹383.72 crore	Multiple performance years; neither annual income nor a measured funding requirement

Sources: Dixon FY2025-26 Note 59(c); Padget FY2025-26 Note 44(a), also disclosed in Dixon Note 59(a). Calculations use the reported lakh amounts before conversion to crore. The subsidiaries are related; rows must not be pooled into a sample average.

The telecom disclosure also says that the prior year's incentive was received. Its amount is not supplied in the same passage. That detail prevents a tempting mistake: the derived ₹153 crore is a receipt against the current incentive cohort, not all incentive cash received during FY2025-26. The explicit 85% statement supports the calculation. Merely subtracting a closing receivable from recognised revenue would not have been enough. [DIXON26, Note 59(c)]

The distinction between an accounting benefit and cash is especially visible at Padget. Its mobile overperformance receivable awaited formal determination and disbursement at the reporting date. Management supported its assessment with legal advice; the auditor drew attention to the matter without modifying the opinion. The accompanying customer liability relates to the same incentives. These statements identify the company's accounting position, not an independent award by the programme administrator or a guarantee of collection. They also do not say that the balances remained unsettled in September 2026. [PADGET26, auditor's report and Note 44(a)]

The scheme's general rules help explain that separation. The 2020 guidelines for electronics production-linked incentives (PLI) provide for unused annual incentives to be allocated to eligible overperformers in the same target segment, subject to excess performance and the available pool. Verification, assessment and payment approval remain separate administrative steps. The identified amendments change the sales definitions, performance window and some claim procedures. Recon-

structuring Padget's own entitlement would therefore require its applicable approval terms, ceilings and claim history, which this investigation does not possess. An applicant's filing deadline is also not a deadline by which the government must pay it. [MEITY20, paragraphs 11.1 and 12; MEITY21A; MEITY21B; MEITY23]

Accounting recognition has its own test. Indian Accounting Standard (Ind AS) 20 requires reasonable assurance that the conditions will be met and the grant received; recognition generally follows the periods of the related costs rather than the arrival of cash. That explains how the two dates can differ. It does not certify that a particular company's assessment is correct. The standard's option to deduct an income-related grant from a related expense also does not, by itself, establish the appropriate treatment of benefits shared with customers. We retain the issuer's description of its reported operating revenue rather than turn a general presentation rule into a company-specific compliance judgment. [ICAI20, paragraphs 7-9, 12, 16 and 29-31]

The customer-linked liability rules out treating the entire matched receivable as available to the manufacturer without a corresponding obligation. Yet subtracting the liability does not produce an estimate of economic profit. Profit also depends on production costs, other contractual terms and the business that would have existed without the incentive. The balance difference is informative precisely when it is left at its actual accounting scope.

Nor is the customer necessarily the final beneficiary. A customer may use lower effective procurement costs to reduce its selling price, expand output, improve its margin or accept a different commercial arrangement. The disclosure cannot discriminate among these outcomes. Establishing one link in a chain is useful evidence, but it does not establish every subsequent link. A study of final consumer benefits would need prices, quantities and a comparison that separates the incentive from simultaneous changes in demand and costs.

An outstanding receivable raises a separate financing question. Suppose the manufacturer must settle its obligation to the customer before the government pays. It may then need to fund the interval. Suppose instead that customer payment is conditional on government receipt. The same closing balances could coexist with a quite different cash requirement. Both are logically possible here because the relevant settlement terms are not disclosed in the passages used. Assigning a borrowing rate to the balance difference would disguise that missing timing evidence as precision.

The appropriate cash reconciliation makes the missing information explicit. For one consistently defined incentive account:

$$\begin{aligned} \text{Closing receivable} &= \text{opening receivable} + \text{recognised claims} \\ &\quad - \text{receipts} + \text{other movements.} \end{aligned}$$

To solve for receipts, all terms must cover the same scheme, entity and cohort. Other movements can include reclassification, reassessment or other changes that the source must identify; they cannot simply be set to zero. A matching customer-payable account needs its own reconciliation. Netting the two closing balances neither supplies their payment dates nor identifies who provided funds in the meantime.

We attempted a wider bridge using Padget's separate statements. It does not close. The government-authorities receivable and annual grant-income lines have different scopes, and the grant disclosures include deferred-income accounting. A reported roll-forward of deferred grant income reconciles a balance in the accounts, not a series of cash receipts. The cash-flow statement does not supply the missing dedicated PLI receipt series. The usable result is therefore the matched overperformance pair, together with the separately explicit telecom receipt statement. [PADGET26, Notes 8, 21, 24 and 44; cash-flow statement]

The matched disclosures support a narrower finding. We can show that recognised incentives have substantial customer-linked components in these disclosures. We can identify one current-year incentive cohort's partial cash settlement. We cannot use the same evidence to measure late-payment costs, a financing constraint or the programme's employment additionality. A more ambitious conclusion would require more evidence; the accounting result does not disappear because those other questions remain open.

The strongest rival interpretation is that the sharing is part of how the policy is supposed to work. A supplier that passes benefits to its buyer may win orders or sustain production it would otherwise lose. Customers' bargaining power may also shape the arrangement. The observed split alone cannot distinguish successful industrial support from an ineffective transfer. To do so, a study would need to connect the commercial terms to additional output, investment or learning, and to the public cost of producing those outcomes.

Calling the subsidy "wasted" because the manufacturer retains a smaller accounting component would assume that receipt by the named firm is the policy objective. No demonstrated measure of social benefit supports that judgment here.


4 Tariffs change several prices at once

A tariff investigation has to follow a different trail. An importer can remit a duty while foreign suppliers, domestic competitors and downstream buyers adjust their prices and quantities. To locate the resulting economic burden, observing the remittance is insufficient. The researcher needs to estimate how outcomes differ from those without the tariff, and to be explicit about which price is being measured.

The US International Trade Commission's (USITC) corrected 2023 retrospective provides a useful example. In its calibrated Section 232 model, the average 2018–2021 price effect for covered steel imports at delivery was +22.73%, compared with +0.74% for domestically produced steel and +2.39% for the US market average. The corresponding aluminium effects were +7.97%, +0.87% and +1.55%. Figure 1 displays the different price objects on a common scale. They are modelled changes relative to a simulated no-Section-232 outcome, not observed inflation rates or shares of the same tax bill. [USITC23, Tables 5.2–5.3]

The border price is only one price in the market

USITC modelled Section 232 effects • published averages for 2018–2021



Different price bases: these percentages are not shares of the same tariff bill.
 Source: USITC Publication 5405, corrected May 2023, Tables 5.2–5.3 (pp. 125–126).
 Calibrated static model; source table reconstruction, not a rerun of the model.

Figure 1. Published 2018–2021 average modelled effects. Source: USITC Publication 5405, corrected May 2023, Tables 5.2–5.3. The figure reconstructs published values; the underlying model was not rerun.

A market-average price combines products from different sources. Only part of that market may face the covered tariff, while substitution changes the competitive position of other suppliers. A domestic producer's price can rise by a smaller percentage than the delivered price of the covered import and still materially change its business. Its revenue depends on quantity as well as price, and its profit also depends on costs. Reading the domestic price change as its share of tariff revenue would skip every one of those relationships.

The model and its underlying empirical evidence also require separate treatment. USITC's domestic and downstream results come from static partial-equilibrium models. Import-supply assumptions are calibrated to approximate the report's border evidence; agreement between the delivered-price results and that evidence is therefore not an independent replication. The framework helps trace selected sectoral consequences under its assumptions. It does not establish economy-wide welfare or the value of dynamic security benefits. [USITC23, Chapters 5–6 and Appendices F–G]

That limitation cuts in both directions. A government cannot establish success simply by citing a protected industry's expansion. A critic cannot establish total failure simply by citing higher input costs in selected downstream industries. Gross production values are not value added, profit or social surplus. Adding selected winners and subtracting selected losers can double count transactions and leave important changes outside the calculation. A welfare assessment needs a more complete treatment of resources, responses and policy objectives.

Nor is border pass-through a universal constant. Amiti, Redding and Weinstein's January 2020 working paper found a substantial contrast between its pooled near-full result and steel-input pass-through approaching one-half after a year. Its data run through October 2019. The later USITC exercise differs in period, coverage and specification. This paper has not decomposed those differences and does not explain

the disagreement by selecting whichever estimate best fits a preferred account of tariff policy. [ARW20, Figure 2 and discussion]

Neither study supplies a pass-through rate valid across sectors and horizons. A new estimate intended to settle the steel comparison would need a common product definition and period, followed by explicit changes in specification. Until then, both findings remain tied to their actual designs. That is a substantive boundary on inference rather than a reason to disregard all border-price evidence.

The September 2026 revision of Amiti, Heise and Weinstein's working paper illustrates another distinction. Its uniform 10% tariff exercise gives a twelve-month price effect of 2.639 percentage points for covered non-oil consumer goods, with a reported 95% bootstrap interval of 1.45–4.38. The components are 1.694 direct, 0.175 through strategic pricing responses and 0.770 through input costs. Their indirect sum is 0.945, or 35.81% of the illustrated total. The estimation sample extends from January 2022 to February 2026; the retail groups cover 20% of the overall consumption basket. This is an attributed relative-price estimate and hypothetical policy exercise, not an estimate of aggregate inflation produced by this paper. [AHW26, Table 4 and accompanying notes]

It would be wrong to subtract that retail percentage from the tariff rate and call the remainder a firm's absorbed share. The tariff applies at one stage, while the covered retail prices include other goods and services. Input-cost changes and competitors' responses can also affect the retail result. To allocate a particular monetary bill, the numerator and denominator would need to refer to that same bill. A smaller price percentage on a larger or differently composed base does not supply that matching.

This is the tariff counterpart of the subsidy problem. The arithmetic becomes economically interpretable only after the object is defined. A gross grant, a customer-linked accounting component and a cash receipt are not interchangeable amounts. Neither are a foreign supplier's price, a delivered import price and a consumer price. The appropriate comparison follows the commercial chain while retaining each measure's denominator.

5 A reimbursement clause shows the allocation a firm seeks

Carbon-border contracts provide evidence at an earlier point in a transaction. In two public templates updated on 9 June 2026, Marubeni-Itochu Steel Europe GmbH proposes ways to recover costs associated with the EU Carbon Border Adjustment Mechanism (CBAM). The customer-facing template permits actual costs to be passed downstream without a surcharge. The supplier-facing template makes reimbursement conditional on unsuccessful customer recovery and requires evidence if requested. Both take effect on signature. The public blank documents establish proposed terms, not executed transactions, incurred costs or completed reimbursements. [MISE-C; MISE-S]

The conditional supplier route matters. A reader who examined only the list of recoverable costs could conclude that a supplier would necessarily bear every listed expense. Reading the paired documents changes that interpretation: the proposed

allocation depends on customer recovery. Yet it would be equally mistaken to conclude that customers in fact paid the costs. The documents reveal what the trader seeks to arrange, while the outcome depends on agreement, actual expenses and settlement.

This evidence adds something that a statutory label alone cannot provide. It identifies a commercial route through which one firm's policy-related expense might become another firm's payment. It also shows why the immediate counterparty matters. "Exporter" and "importer" can be useful categories, but the commercial result may turn on which party has agreed to reimburse which costs, under which conditions. A policy-level statement about statutory responsibility cannot settle those contractual details.

The next decisive observations would be straightforward in concept: a signed agreement, a qualifying invoice and a settlement record. They would show whether the proposed allocation operated in a particular transaction. Evidence about changes in product prices, volumes or supplier selection would then be needed to investigate wider incidence. The burden of proof rises with the claim. An analysis of published templates does not need to pretend that those confidential records are available, but it must stop before claiming a realised transfer.

Timing could matter even where ultimate reimbursement is certain. A firm that pays an expense before recovering it may face an interval of financing. Another contract might place payment dates close together or provide advance funding. The cost allocation alone does not identify either arrangement. A paper interested in financing would need the relevant dates and funding terms, just as it would for the Indian subsidy claim. The analogy is about the missing evidence, not a claim that the two instruments generate the same cash exposure.

There is also a difference between allocating a compliance expense and changing incentives to reduce it. A firm may be able to document costs precisely without being able to control the underlying emissions or verification process. Another party may have the technical ability to lower them. Whether a reimbursement arrangement rewards that effort depends on its terms and the wider commercial relationship. The templates examined here do not establish the resulting investment response. They make that response a concrete question for transaction-level research.

The carbon-border case therefore contributes contractual evidence to the comparison, rather than an additional estimate to place beside the tariff coefficients or subsidy ratios. Pooling proposed rights with subsidy ratios and tariff coefficients would conceal their different evidentiary status.

6 What the comparison changes for policy evaluation

The cases expose three shortcuts with different failures. The subsidy disclosure contains a customer-linked component that a gross-recipient description misses. The tariff studies use different price bases, so their percentages cannot be treated as divisions of one monetary amount. The contract templates describe conditional rights, so they cannot establish payments. These findings share a question about

commercial allocation, but they do not share an estimand: there is no single underlying effect that all three measurements recover.

Table 2. One question, different observations

Evidence	Defensible finding	Evidence needed for the stronger claim
Matched subsidiary disclosures	Recognised incentive, linked customer component and one cohort's explicit receipt	Matched payment histories for complete cash allocation; prices and a counterfactual for ultimate incidence
Tariff price studies and calibrated models	Source-specific responses relative to a stated comparison	Harmonised data and identification for a new common estimate; broader modelling for welfare
Public carbon-border templates	Proposed conditional recovery rights	Signed terms, invoices and settlements for realised reimbursement

This is an evidence comparison, not a ranking of policy quality, legal force or administrative maturity. Detailed entity, period and source locators are retained in the replication notes.

The first implication is for measurement. A programme dashboard organised only around the named recipient can accurately describe public disbursement while leaving commercial sharing invisible. A financial statement can disclose sharing while leaving actual settlement unclear. Neither source should be made to answer a question it was not constructed to measure. Linking their scopes is more useful than declaring one number the definitive policy benefit.

For a subsidy evaluation, the minimum useful extension would be a scheme-specific bridge between recognised claims, determinations and receipts, with a separately identified customer-sharing component where material. Such a bridge would not require public disclosure of every commercial detail. Aggregated or appropriately protected reporting could preserve confidentiality while distinguishing a stock of claims from a year's payment flow. The relevant reporting design would still have to consider costs and what the programme can reliably collect.

A second implication concerns evaluation targets. If the policy aims to encourage additional productive activity, the recipient's accounting share is an intermediate outcome. A large share might accompany no change in behaviour. A smaller share might help secure an otherwise unviable order. Evaluation should specify the intended activity and test it against a defensible alternative, including the opportunity cost of public funds. The disclosures in this paper cannot supply that causal comparison, but they show why a test based solely on the gross amount received would be incomplete.

Financing relief is also an intermediate mechanism rather than a sufficient explanation of productivity gains. Research on Hungarian firms by Goel and co-authors distinguishes subsidy recipients' outcomes from the incremental response of credit-

constrained firms; the latter comparison does not produce a general guarantee of lasting efficiency gains. The setting is different from Indian PLI. Its relevance is to the hypothesis: relaxing a financing constraint and improving productive performance are separate propositions requiring separate evidence. [GOEL21]

For tariffs, the analogous improvement is to publish the price object alongside the headline estimate. “Pass-through” should specify the stage, product coverage, period and counterfactual. A result for imports need not describe retail goods, and a retail result need not describe the whole consumption basket. Where a model turns empirical estimates into downstream effects, the dependence should be visible. It is particularly misleading to present the model’s agreement with its own calibration target as a second independent confirmation.

For carbon-border arrangements, research should distinguish proposed terms from adopted terms and adopted terms from collected reimbursements. A public contract collection can map plausible allocations and identify clauses worth investigating. A transaction sample can test whether those allocations occur, while retaining the possibility that observed contracts are selected. A representative claim about exporters would require a sampling design that reaches beyond firms willing or able to publish their documents.

These proposals are not interchangeable. Adding payment dates does not fix a price study’s counterfactual. Adding a causal regression does not reveal whether a particular reimbursement clause was signed. Each missing link requires the type of evidence suited to it. A broad demand for “more data” is less useful than stating which observation would change the conclusion and why.

The comparison also makes room for a serious defence of the policies. Customer benefits may be intended; a recoverable claim may be inexpensive to finance; tariff protection may pursue security benefits outside a static sector model; carbon-cost recovery may allow trade to continue under a new compliance obligation. None of these possibilities should be assumed to settle the evaluation either. They are rival mechanisms to investigate, not reasons to treat every observed transfer as a success.

An effective research agenda would follow the same relationships over time. For the subsidy cases, later matching disclosures could establish settlement and clarify the customer obligation. For tariffs, harmonising periods and product definitions could help determine how much of the apparent disagreement reflects design rather than economic change. For carbon-border costs, signed contracts and invoices could turn a map of proposed rights into evidence of actual allocation. Each extension would add an identifiable result rather than merely increase the number of documents cited.

7 Follow the relationship as well as the amount

A matched record of government settlement and customer settlement for the same incentive cohort would most directly extend the company-level finding. It would show how the recognised incentive and customer obligation became cash allocation,

preserving the distinction between receipt by the manufacturer and payment to the buyer.

The next question would be harder: what changed in production, prices or investment because the policy existed? Settlement records would not answer it. But they would establish a commercial transfer that could help interpret a later causal investigation. For these disclosures, following the same incentive through both relationships is the next useful observation.

References and source versions

Citation keys in the paper resolve to the sources below. Historical balances are dated observations, not current settlement claims.

- DIXON26** Dixon Technologies (India) Limited. *Integrated Annual Report 2025–26*, consolidated financial statements, Notes 59(a)–(c); year ended 31 March 2026. Official issuer PDF. Relevant Note 59 pages: PDF 436–437, printed 499–500. Audit emphasis: PDF 336, printed 399.
- PADGET26** Padget Electronics Private Limited. *Audited financial statements for FY 2025–26*, year ended 31 March 2026. Official issuer-hosted PDF. PDF 1, 15–16, 39, 45, 47, 63; Note 44(a) identifies mobile PLI. Scan/OCR checked against original page images.
- MEITY20** Ministry of Electronics and Information Technology. *Guidelines for the Production Linked Incentive Scheme for Large Scale Electronics Manufacturing*, 1 June 2020. Signed official PDF. Paragraph 11.1, PDF 21–22; paragraphs 12.1–12.9, PDF 22–23. Read with the identified amendments below; no company-specific award or exhaustive administrative history is asserted.
- MEITY21A** MeitY. Addendum amending the PLI guidelines, 11 March 2021. Official implementing-agency copy, PDF 1–3.
- MEITY21B** MeitY. Notification extending PLI LSEM tenure, 23 September 2021. Official PDF, PDF 1, 3–4. First-round applicants’ performance-year choices cannot be assigned to a specific company without its applicable records.
- MEITY23** MeitY. Addendum amending paragraph 12.2, 28 March 2023. Official implementing-agency copy, PDF 1–2. The identified exception concerns certain deferred claims, not an established payment due date for the company.
- ICAI20** Institute of Chartered Accountants of India. *Compendium of Indian Accounting Standards, 2025–26*, Ind AS 20: Accounting for Government Grants and Disclosure of Government Assistance. Official PDF. Paragraphs 7–9, 12, 16, 29–31; PDF 4–6, 8. Official ICAI compilation used after the MCA route was inaccessible. No company-compliance opinion is supplied.

USITC23 United States International Trade Commission. *Economic Impact of Section 232 and 301 Tariffs on U.S. Industries*. Publication 5405, March 2023, corrected May 2023. Official PDF. Tables 5.2-5.3, pp 125-126; Chapters 5-6; Appendices F-G. Published calibrated-model output is reconstructed, not re-estimated.

ARW20 Amiti, Mary; Stephen J. Redding; David E. Weinstein. *Who's Paying for the US Tariffs? A Longer-Term Perspective*. NBER Working Paper 26610, January 2020. Working-paper PDF. Discussion PDF 5 and Figure 2 PDF 8. Published counterpart: *AEA Papers and Proceedings* 110, 541-546, DOI 10.1257/pandp.20201018. Working version used; no byte-equivalence with the published version is claimed.

AHW26 Amiti, Mary; Sebastian Heise; David E. Weinstein. *The Anatomy of Tariff Pass-Through into Consumer Prices*. Federal Reserve Bank of New York Staff Report 1201, August 2026, revised September 2026. Original PDF; DOI 10.59576/sr.1201. Table 4 PDF 20/printed 19 and related notes; sample/methods PDF 9-10, 19. A working paper; no new microdata replication here.

GOOLSBEE98 Goolsbee, Austan. "Investment Tax Incentives, Prices, and the Supply of Capital Goods." *Quarterly Journal of Economics* 113(1), February 1998, 121-148, DOI 10.1162/003355398555540. Author-hosted manuscript, dated June 1997, is the accessible version read. Used for the qualitative mechanism and research positioning; no numerical estimate transplanted to PLI and no replication claim.

GOEL21 Goel, Tirupam; Péter Lang; Eszter Balogh; Ádám Banai; Martin Stancsics; Előd Takáts; Álmós Telegdy. *Credit Constrained Firms and Government Subsidies: Evidence from a European Union Program*. BIS Working Paper 984, December 2021. Original PDF. Introduction PDF 3-6 and conclusion PDF 26. Distinguish incremental effects for constrained firms from effects of subsidies generally.

MISE-C Marubeni-Itochu Steel Europe GmbH. *Supplementary Agreement on Cost Adjustment under the EU CBAM Regulation*, MISE as seller to customer, updated 9 June 2026. Public template, §§ 2-4, PDF 1-2. Blank counterparties/signatures; reused from A12.

MISE-S Marubeni-Itochu Steel Europe GmbH. Same title, MISE as buyer from supplier, updated 9 June 2026. Public template, §§ 2-6, PDF 1-2. Upstream reimbursement is conditional on unsuccessful customer recovery; reused from A12.