

Industrial Policy for Semiconductors in the AI Hardware Era: Text Based Measurement and Cross Country Evidence

YE, Kai ^{1*}

¹ Hezhou University, China

* YE, Kai is the corresponding author, E-mail: yktaikongyin@gmail.com

Abstract: As demand for AI hardware increases, semiconductor industry policies, through subsidies, tax credits, and targeted financing, are expanding, leading to uneven distribution of policy support documents and strategic frameworks, as well as long implementation delays. This paper proposes a replicable measurement process to identify semiconductor-related interventions in global trade alerts through iterative validation. Each indicator is categorized according to value chain objectives and policy tools, providing a dataset and evaluation framework for semiconductor research.

Keywords: Semiconductor Industrial Policy, AI Hardware, Text Mining, Subsidy Quantification, International Spillovers.

Disciplines: Natural Language Processing.

Subjects: Text Mining.

DOI: <https://doi.org/10.70393/6a6374616d.333730>

ARK: <https://n2t.net/ark:/40704/JCTAM.v3n1a06>

1 INTRODUCTION

In recent years, subsidies, tax credits, and directed finance have returned to the center of industrial policy debates, often justified by resilience, security, and technological leadership. It is still unclear whether this resurgence should be read as a short run response to successive shocks, or as a more durable reorientation toward state shaped industrial allocation. Part of the difficulty is that the relevant sectors are dynamically organized. Policies that look modest in the present may, over a longer horizon, shift learning trajectories, entry conditions, and the geography of innovation, even when headline trade indicators move slowly.^[1]

Semiconductors are an especially consequential setting for this discussion. They sit upstream of a wide range of downstream activities, and the rise of AI hardware demand makes supply constraints more salient and potentially more macro relevant.^[2] At the same time, semiconductors are a dual use technology, so policy rationales that mix efficiency, security, and strategic rivalry are common, and the same instrument can plausibly operate through multiple channels. Under these conditions, it seems risky to assume either that industrial policy is generally effective, or that it is mostly wasteful. Effects may exist only to some extent, differ across value chain stages, and vary with pre existing capabilities, while further research is needed to clarify the mechanisms that dominate in practice.

A recurring source of disagreement in the industrial

policy literature is not only identification strategy, but also what is actually being measured. Firm based approaches that infer support from financial statements and financing conditions can be disciplined in accounting terms, yet they often face limited coverage and fragile comparability across jurisdictions, and they rely on counterfactual benchmarks that are themselves altered by policy.^[3] Text and event based approaches can offer broader cross country coverage, but they face a different set of issues, including strategic wording, incomplete documentation, and the gap between policy announcements and economically meaningful implementation. In semiconductors, these problems become sharper. Industry boundaries are porous, supply chains are segmented, and many interventions are framed under broad labels such as advanced manufacturing or digital infrastructure. Coding systems such as HS classifications can be too coarse, missing key upstream and complementary segments, and policy databases may contain entries with incomplete industry tags.^[4] As a result, measurement is not a technical detail that can be solved once and then ignored. It shapes both the construction of treatments and the plausibility of later causal claims.

This paper pursues three linked objectives. First, it constructs a value chain aware record of semiconductor industrial policy using a transparent text based identification and validation pipeline, while treating measurement uncertainty as an empirical constraint rather than a background detail. Second, it evaluates associations between policy exposure and outcomes that are salient in the AI hardware era, including capacity proxies, trade and sourcing

patterns, innovation indicators, and price related measures, using designs that emphasize timing diagnostics and sensitivity to alternative exposure definitions.^[5] Third, it addresses the possibility of cross border spillovers by defining foreign exposure measures that reflect supply chain linkages, since purely domestic evaluations can confound efficiency gains with international reallocation. The aim is not to close the debate over industrial policy, but to narrow the set of plausible narratives by making both measurement and inference contestable and replicable.

2 DATA AND MEASUREMENT

This paper constructs a cross-country policy dataset primarily from Global Trade Alert, which provides systematic coverage of commercial and industrial interventions. The choice is pragmatic rather than ideal, because some forms of support may be under-recorded when implemented through opaque credit channels or subnational programs, so an observed “absence” of policy in the dataset should not be interpreted as a definitive absence of support in reality.^[6] Our baseline window is 2010–2025, consistent with the reference text-based industrial policy pipeline. We keep the construction modular so that later extensions are feasible, while treating end-of-sample measurement as potentially noisier because of reporting lags and retrospective updates.

2.1 IDENTIFICATION OF SEMICONDUCTOR-RELEVANT POLICIES

GTA sometimes provides HS codes, yet HS-based identification is likely to be incomplete for semiconductors. Codes can be too broad, they can miss upstream inputs and equipment, and a meaningful fraction of records lacks HS codes altogether. This matters because systematic omission of upstream segments would mechanically bias both the measured level of policy intensity and its value-chain composition, which then contaminates downstream inference. We identify semiconductor-relevant measures using a combined approach: HS codes where available, and a domain dictionary applied to policy descriptions to create a candidate set. The dictionary is refined iteratively because early versions may miss policies framed under broad “advanced manufacturing” language, while later expansions can pull in adjacent ICT measures whose link to semiconductor capability is weak. These refinements are treated as substantive design choices rather than cosmetic tuning, since they shift treatment assignment and timing.^[7]

2.2 MANUAL VERIFICATION AND CONTEXTUAL CHECKS

Candidate policies are manually verified for semiconductor relevance, and contextual information is collected from policy documents and other corroborating sources. This step cannot fully remove subjectivity, yet it reduces mechanical errors that arise when brief summaries

omit implementation details, and it aligns with the reference study’s emphasis on validation and local expertise. Each verified policy is mapped to its primary value-chain target, distinguishing inputs, design and manufacture, and assembly packaging and testing, consistent with the reference study’s organization. Instrument types are coded when documentation allows, separating direct spending, tax incentives, preferential lending, equity injections, guarantees, and procurement-related measures. Some interventions plausibly span multiple stages; in such cases we record secondary targets for robustness checks rather than forcing an artificially crisp label.^[8]

Because policy descriptions are short and often strategic in wording, relevance screening and value chain targeting inevitably involve judgment. To reduce idiosyncratic coding, we implement a structured reliability procedure. A stratified sample of candidate records is double coded for semiconductor relevance, primary value chain target, and instrument type. Disagreements are logged by category, then resolved through a documented rule set that prioritizes explicit eligibility language and implementable clauses over aspirational statements. Records that remain ambiguous after reconciliation are retained with an ambiguity flag and are reclassified in sensitivity checks. This procedure does not remove subjectivity, yet it makes the direction and scope of judgment visible, which matters for later inference when timing and intensity depend on classification boundaries.^[9]

2.3 QUANTIFYING POLICY INTENSITY AND DATA-QUALITY DISCIPLINE

We attempt bottom-up quantification by extracting monetary commitments from policy text and source documents, then converting to a common currency and recording the implied horizon. Where policies report ceilings instead of realized disbursements, we retain both a ceiling measure and an implementation-status flag when feasible.^[10] Tax-based instruments are coded as tax equivalents only when a defensible base can be inferred; otherwise, we apply conservative qualitative intensity scores mainly for sensitivity analysis. Because the dataset is assembled from heterogeneous textual records, the measurement process itself raises data-engineering issues that are not always foregrounded in policy evaluations, including reproducible extraction, consistent transformations, and the handling of rare categories. In this respect, applied semiconductor manufacturing studies that emphasize automated ETL, systematic quality control, and class-imbalance handling provide a useful methodological reminder, even if the empirical objects differ. Yin (2025) stresses that data pipelines can silently accumulate bias when imbalance and missingness are treated casually, which motivates our versioned ETL scripts, explicit missingness reporting, and stratified sampling for validation rather than ad hoc cleaning.^[11]

A practical risk is double counting because policy

packages can appear multiple times with amendments or re-phrasings. We address this by assigning package identifiers and coding incremental changes only when an amendment plausibly alters effective support. Measurement remains uneven across countries, and the reference study notes that support in semiconductors is difficult to quantify and that some jurisdictions may be systematically under-captured due to limited visibility of subnational financial support. We treat this as a central limitation, and later results are interpreted with this asymmetry in mind.

2.4 OUTCOME LINKAGE AND MECHANISM-RELEVANT PROXIES

The policy indices are merged to outcome panels on capacity proxies, trade flows, innovation indicators, and price measures. The choice of outcomes is guided by the idea that policy may affect not only final output, but also operational constraints and bottlenecks along the semiconductor production system, which can be especially salient under AI-hardware-driven demand surges. Recent manufacturing oriented studies provide a useful reminder that operational constraints such as downtime, yield, congestion, and logistics can shape observed capacity and trade outcomes even when headline investment expands. These studies do not identify cross country policy effects, yet they discipline how we interpret mixed macro patterns by clarifying which margins may move first and which margins may remain hidden in aggregate proxies.^[12] Predictive maintenance and fault-severity classification, as studied in Yin (2025) highlight how equipment reliability and downtime can be treated as operational constraints that translate into capacity outcomes rather than merely plant-level accounting totals. Bottleneck detection methods, such as the active-period approach in Yin (2025) suggest that capacity constraints may emerge from localized congestion patterns, which implies that country-level capacity proxies can conceal within-facility constraints. Defect prediction and explainable AutoML, as in Yin (2025), underscore that yield and quality can shift even when nominal capacity expands, which motivates our use of innovation and price-based measures as partial cross-checks rather than relying on a single capacity series. Finally, wafer logistics and scheduling under uncertainty, as discussed in Yin (2025), remind us that policy targeting one segment of the chain may be diluted if logistics frictions bind elsewhere.^[13]

Outcome definitions, sources, and comparability discipline

Outcome selection follows a mechanism oriented logic, while recognizing that no single proxy fully captures semiconductor performance. Capacity outcomes include measures that approximate production scale and expansion, while acknowledging that nominal capacity may diverge from effective output when yield, uptime, or qualification bind. Trade outcomes include flows and sourcing patterns that reflect cross border dependence and re routing, while recognizing that third country transit can blur origin

interpretation. Innovation outcomes include patent based and related indicators, while recognizing cross country differences in patenting strategy and time to commercialization. Price related measures are used as partial cross checks for scarcity and constraint pressure, while recognizing that contract structure and product mix can shift measured prices. For each outcome, the paper reports source coverage, frequency, and missingness patterns, and it repeats baseline estimates on a restricted sample where documentation and outcome coverage are jointly stronger, so that the core patterns are not driven by extreme differences in observability.^[14]

These studies do not “prove” policy effects, and their evidence is not directly a substitute for cross-country identification. They mainly serve to discipline how we think about channels and how we interpret null or mixed results: a weak export response, for instance, could reflect the binding of equipment uptime, yield, or logistics rather than the absence of policy impact.^[15]

This paper provides a replicable pipeline that converts policy texts into value-chain-aware, partially quantified measures of semiconductor industrial policy.^[16] The pipeline balances recall and precision through dictionary-based identification, manual validation, and conservative coding choices. At the same time, it does not claim definitive measurement. Uneven transparency, reporting lags, and unobserved financial support imply that measured policy intensity may diverge from true support, and further research is needed to triangulate these indices with firm-level accounts or model-based inference.^[17]

Methodologically, it is also possible that tools developed in adjacent domains could improve parts of the pipeline, although transferability is not guaranteed.^[18] For example, uncertainty modeling strategies used in generative diffusion approaches for financial volatility, as in Chen (2025), may offer ideas for sensitivity analysis when policy intensity is only partially observed. Supervised anomaly-detection practice in fraud analytics, as in Chen (2025), is conceptually aligned with detecting duplicate or inconsistent records in large text corpora, though domain shifts can introduce new failure modes.^[20] Even incentive-mechanism perspectives, such as Pang (2020), provide a cautionary lens: reporting incentives and organizational behavior can shape what becomes observable in administrative and textual records. Considering these factors, the chapter positions measurement not as a solved prelim step, but as a continuing source of uncertainty that must remain visible when causal claims are later discussed.^[21]

3 EXPERIMENTAL DESIGN AND DATA ANALYSIS

The methodological stance is intentionally cautious. Industrial policy in semiconductors is rarely a clean external shock. It can be a reaction to anticipated shortages, a response

to geopolitical risk, or an instrument for longer-run technological positioning, and these motives may coexist within the same policy package. Considering the above factors, the goal here is not to force a single deterministic causal narrative, but to develop a research design that is sufficiently transparent to be challenged, replicated, and stress-tested, while still being informative to some extent about the direction and timing of potential effects. A further motivation is that policy “intensity” is not naturally observed. Even when policy texts disclose budgets, execution can be delayed, reallocated, or absorbed by pre-existing investment plans. This leads us to treat measurement error, uneven observability across jurisdictions, and implementation frictions as central empirical objects, rather than as residual inconveniences. In practical terms, the chapter combines event-style dynamics with continuous-exposure panel designs, and then uses heterogeneity analysis and robustness checks to narrow the set of plausible interpretations. Further research is needed to resolve the remaining ambiguity, yet a disciplined methodology can still shift the debate from assertion to evidence.^[19]

3.1 EMPIRICAL UNITS, CONSTRUCTED

TREATMENTS, AND THE LOGIC OF EXPOSURE

The unit of analysis is a country-year panel. For each country and year, we observe a set of semiconductor-related policy records that have been validated for relevance and coded for monetary magnitude when documentation permits, as well as coded for value-chain targets and instrument types. The analysis constructs a continuous exposure measure at the country-year level by aggregating documented magnitudes, and it also constructs composition measures that capture how policy attention is distributed across value-chain segments, including inputs and equipment, design and manufacture, and assembly, packaging, and testing.^[22]

The reason to carry both intensity and composition is conceptual rather than cosmetic. A given level of fiscal support may operate very differently if it is aimed at fabrication capacity expansion than if it is aimed at equipment access or materials, and the same aggregate number can hide different bottleneck-relaxation strategies. At the same time, segmentation is not always well-defined in policy texts. In early iterations, we found that broad “strategic technology” packages were difficult to map to a single stage, and naive classification created an artificial sense of precision. This motivated an approach that assigns a primary target while recording secondary tags for sensitivity analysis, so that borderline cases can be reclassified and results can be checked for fragility. A parallel difficulty concerns de-duplication and policy packages. A single legislative act can generate multiple entries, while amendments can blur whether a record represents new support or a continuation. We addressed this through a package reconciliation procedure that is conservative by design, partly because the cost of double counting is not symmetric.^[23] It can create mechanical spikes that are later mistaken for substantive

policy surges, which then contaminates event timing and distorts inference.

Event-based analysis requires a notion of when a policy episode begins. We define a policy surge year using transparent rules that combine changes in exposure with the documented start of major programs. However, the chapter treats the surge year as an empirical construct rather than as a natural fact, because reporting lags, cross-year budgeting, and staged implementation can shift the economically relevant moment away from the announcement date.^[24]

Anticipation behavior is a second reason for caution. Semiconductor investments are multi-year and often depend on expectations about future support, export conditions, and demand. Firms may accelerate site selection and equipment ordering during legislative debate, not only after enactment. This means that pre-treatment periods can already contain policy-related adjustments, which complicates a naive reading of event-study coefficients. The analysis therefore places substantial weight on dynamic diagnostics. It tests for pre-trends, it examines alternative windows, and it treats early movements as informative signals rather than as mere violations to be ignored.^[25]

3.2 OUTCOMES, PREPROCESSING, AND TRIANGULATION STRATEGY

The outcome set is intentionally multi-dimensional. We track capacity proxies, trade flows and sourcing patterns, innovation indicators, and price-related measures when available. The motivation is not to increase the number of results, but to reduce the risk that any single proxy dominates interpretation. Capacity measures can miss yield and uptime, trade flows can be re-routed through third countries, and innovation indicators can be noisy in short windows. Considering these limitations, triangulation becomes a methodological necessity rather than a stylistic preference. Preprocessing aims to be rule-based and replicable. Variables that are highly skewed are transformed in a consistent way. Extreme values are examined through winsorization sensitivity rather than removed ad hoc. Missingness is treated as potentially non-random. During exploratory analysis, we observed that missing monetary coding is more likely when policy instruments are indirect or when documentation is fragmented, which suggests that missingness can correlate with the institutional features that also predict outcomes. For that reason, the baseline does not rely on aggressive imputation. Instead, the chapter reports coverage and performs checks that restrict samples to comparably documented settings, while acknowledging that such restrictions trade generality for measurement comparability.^[26]

The event-study design estimates how outcomes evolve around policy surge years relative to a baseline period, controlling for country fixed effects and year fixed effects. The main role of this design is diagnostic and descriptive. It shows whether outcomes begin to diverge prior to the surge,

and it indicates whether observed changes are short-lived adjustments or persistent shifts. Interpretation is intentionally layered. A post-surge increase in a capacity proxy could reflect new investment induced by policy, but it could also reflect demand-driven private investment that policy merely accompanies. A post-surge change in trade patterns could reflect import substitution, but it could also be driven by export controls or supply disruptions that coincide with policy activation. The event-study does not mechanically separate these narratives, so the chapter treats it as a structured way to evaluate which narratives are consistent with timing and dynamics, rather than as a final causal proof. To avoid over-reading the pattern, we run alternative timing rules, alternative baselines, and placebo exercises that shift the surge year. We also compare the dynamics across outcome families. A response that appears only in one proxy but not in related proxies is treated with skepticism, while a coherent pattern across capacity, trade, and price measures is treated as more plausible, even though further research is needed to rule out all confounding.^[27]

3.3 EXPERIMENTAL DESIGN II: CONTINUOUS-EXPOSURE PANEL DIFFERENCE-IN-DIFFERENCES

Because semiconductor policy often arrives in waves and varies substantially in magnitude, the chapter complements event studies with a continuous-exposure panel design. The core idea is to relate within-country changes in outcomes to within-country changes in exposure, while absorbing time-invariant country characteristics and common global shocks through fixed effects. This design has an advantage. It uses more of the variation in the data than a binary treatment indicator, and it can be more informative when policy intensity rises gradually rather than abruptly. Yet it also has a vulnerability. Exposure can increase for reasons that are correlated with outcomes, such as anticipated supply chain risk, domestic political cycles, or global demand surges. This makes endogeneity a central concern rather than an afterthought. The chapter responds by incorporating controls that proxy for global cycles and demand conditions when feasible, by conducting sensitivity analysis to alternative exposure definitions, and by using lag structures to test whether exposure changes precede outcome changes in a way that is not easily explained by contemporaneous shocks. The estimated coefficients are therefore interpreted as average associations consistent with a policy channel, to some extent, rather than as immutable structural effects. Where estimates are sensitive to specification choices, that sensitivity is reported explicitly and used to refine the range of defensible claims.

A key methodological choice is to treat heterogeneity as an object of study. We test whether the relationship between policy exposure and outcomes varies with value-chain targeting, with baseline capability proxies, and with supply-chain positioning. This approach is motivated by the

observation that the same policy instrument can operate differently depending on complementary constraints. An equipment-focused support package may matter more where maintenance, uptime, or tool availability is binding, while a fabrication-focused package may matter more where permitting, workforce, and utilities are the limiting factors. However, heterogeneity models can be misleading if they simply re-label confounding as “mechanisms.” For that reason, interaction results are interpreted alongside descriptive evidence on policy composition and alongside robustness checks. Competing explanations are considered. For example, stronger associations in advanced economies could reflect better measurement rather than stronger true effects, while weaker associations in low-transparency settings could reflect missing support rather than policy ineffectiveness. The chapter does not claim to resolve these ambiguities fully, but it uses them to discipline how much weight is placed on any single interaction pattern.

The robustness plan is designed around three empirical risks. The first is treatment mismeasurement, including incomplete monetary coding and inconsistent documentation. The second is policy endogeneity, including reactive adoption during shocks. The third is outcome proxy limitations, including re-routing in trade and short-run noise in innovation.

To address these risks, the analysis proceeds in layers. It compares monetary exposure with count-based exposure, it varies event definitions, it restricts samples to subsets with higher documentation comparability, and it tests alternative outcome families. Placebo tests shift treatment timing and examine whether “effects” appear when they should not. Sensitivity checks adjust transformations and outlier handling to ensure that a small set of observations is not driving the results. The chapter treats robustness not as a ritual, but as a way to map which claims are stable across plausible coding choices, and which claims remain too fragile to support strong policy conclusions.

All processing steps are documented, including dictionary versions, de-duplication rules, and aggregation conventions. This is not merely a technical preference. In a domain where policy texts are strategically written and where measurement is contestable, reproducibility is part of the substantive argument, because it allows other researchers to challenge the treatment definition rather than to debate conclusions in the abstract. The research process was not smooth. Early pipeline versions over-emphasized narrow keywords and missed broad strategic packages. Later expansions improved recall but increased false positives, especially for adjacent electronics measures. We adjusted the dictionary iteratively and introduced validation steps that record disagreements rather than hiding them. These adjustments changed the distribution of exposure and, in some cases, shifted the timing of apparent surges. Instead of treating this as an embarrassment, the chapter treats it as evidence that measurement choices matter and that results should be presented with appropriate modesty.

3.4 CROSS BORDER SPILLOVERS AND FOREIGN

EXPOSURE DESIGN

A domestic policy evaluation can be misleading if policy induced changes are transmitted through supply chains and international reallocation. To make spillovers empirically testable, the paper defines foreign exposure measures that map partner country policy intensity into domestic vulnerability and opportunity channels. The baseline foreign exposure for country c in year t is constructed as a weighted sum of partner policy intensities, where weights reflect pre policy trade linkages in mechanism relevant inputs, equipment, and intermediate categories. This construction follows a simple logic: if a country is strongly dependent on upstream imports from a partner that intensifies semiconductor policy, domestic outcomes may change through price pressure, supply availability, or re routed trade even without domestic policy changes. A parallel channel arises when partner policy expands downstream capacity that competes in third markets, which may shift exports, prices, and sourcing patterns in other countries.

4 EMPIRICAL RESULTS AND INTERPRETATION

Considering the intrinsic endogeneity of industrial policy in semiconductors, the chapter does not treat any single coefficient as a decisive structural object. Instead, it treats the results as evidence that is informative to some extent, while also being exposed to alternative explanations such as demand shocks, export controls, measurement asymmetries, and anticipatory investment behavior. This leads us to interpret the empirical patterns with disciplined caution, and to keep the link between empirical facts and policy narratives explicit rather than assumed. The results are presented with explicit links between narrative claims and empirical objects. Table 1 summarizes policy exposure distributions, coding coverage, and missingness by country group. Figure 1 visualizes the time profile of exposure and the concentration of documented support. Figure 2 reports the main event study dynamics for a capacity proxy and a trade outcome, with pre trend diagnostics shown explicitly. Table 2 reports continuous exposure panel estimates under a baseline control set, while Table 3 reports robustness to count based exposure, alternative surge timing rules, and restricted samples with stronger documentation. Table 4 reports value chain targeting heterogeneity, and Table 5 reports spillover estimates using foreign exposure measures under alternative weighting schemes. This structure is intended to make disagreements about measurement and timing concrete, since many debates in this domain hinge on treatment construction rather than only on regression specification.

4.1 BASELINE PATTERNS AND FIRST-PASS

DESCRIPTIVE EVIDENCE

The baseline descriptive evidence shows that the policy exposure measures are highly skewed across countries and over time. A small number of jurisdictions account for a large share of documented policy intensity, and the timing of exposure increases is clustered around major legislative and planning moments rather than being smoothly distributed. This clustering matters because it suggests that policy “treatment” is episodic, which makes dynamic designs particularly relevant, yet it also raises the possibility that the same external shock that motivates policy could simultaneously shape outcomes.

When outcomes are plotted around periods of exposure increases, the patterns are not uniform. Capacity proxies tend to rise in some treated jurisdictions, yet the timing is not always immediate, and in several cases the visible adjustment begins prior to the official policy surge year. This pre-movement is consistent with the possibility of anticipation, but it is also consistent with the alternative possibility that both policy and capacity expansion are jointly driven by an underlying demand shock. Trade outcomes show even more variation. In some cases, the post-surge period coincides with a relative increase in domestic production proxies and a change in import composition, while in other cases trade patterns appear dominated by broader supply chain re-routing that is difficult to attribute to a specific domestic policy package. Two interpretive frictions emerged during the descriptive phase. First, in countries where monetary coding coverage is lower, the exposure measure may understate true support, making descriptive comparisons potentially misleading. Second, several outcomes, especially innovation and prices, appear more volatile in short windows than anticipated, which required extending observation windows and examining alternative smoothing conventions. These adjustments did not eliminate volatility, yet they clarified that some apparent “effects” in short windows were likely fragile.

4.2 EVENT-STUDY RESULTS AND TIMING

DIAGNOSTICS

The event-study estimates provide two types of information. They indicate whether the outcome paths for high-exposure countries begin to diverge before the policy surge year, and they reveal the shape of the post-surge adjustment, whether it resembles a temporary jump, a gradual trend break, or a pattern that fades. Across the main capacity-related outcomes, the event-study patterns are broadly consistent with a gradual post-surge adjustment in several jurisdictions, with the largest movements often appearing two to four years after the surge. Such lags are plausible given permitting, construction, tool installation, and qualification requirements in semiconductor production. However, the event-study also reveals that in a nontrivial subset of cases the pre-surge coefficients are not statistically close to zero, suggesting that simple parallel-trends assumptions are at least partially strained. This does not invalidate the analysis, but it does shift the burden toward cautious interpretation and toward additional robustness checks that test alternative

timing rules and windows. For trade-related outcomes, the event-study patterns are more mixed. Some jurisdictions show post-surge reductions in import dependence or changes in sourcing that align with import substitution narratives. Yet in other jurisdictions, the trade response is weak or even moves in the opposite direction, which could reflect that domestic policy expands downstream capacity that still relies on imported inputs, or that export controls and geopolitical frictions reshape trade flows in ways that overwhelm domestic policy signals. Considering these factors, trade event studies are interpreted as evidence about timing and directionality rather than as decisive proof of import substitution.

Innovation outcomes provide a different picture. Patent-based measures and related innovation proxies show slower and noisier responses, with patterns that are sensitive to window choice. One plausible interpretation is that innovation responds through longer learning and coordination cycles. Another is that innovation indicators are poorly aligned with the specific margin affected by policy, especially when policy is targeted at manufacturing scale and reliability rather than at upstream R and D. A third possibility is that innovation responses are present but masked by measurement noise and cross-country differences in patenting strategies. The event-study evidence does not cleanly separate these stories, which is why it is presented alongside panel results and heterogeneity analysis.

4.3 PANEL ESTIMATES WITH CONTINUOUS

EXPOSURE AND THEIR STABILITY

The panel estimates relate within-country outcome changes to within-country exposure changes while controlling for country and year fixed effects. In several specifications, exposure is positively associated with capacity proxies and with certain innovation proxies, though the magnitudes vary with the control set and with alternative exposure definitions. The sensitivity itself is informative. It suggests that some of the variation in exposure coincides with macro and geopolitical shocks that also influence outcomes, which is consistent with the underlying concern about endogeneity. When the exposure measure is replaced by a count-based measure, some capacity associations remain while others attenuate. This pattern can be read in two ways. It might imply that monetary coding carries signal beyond simple counts, perhaps reflecting genuine differences in scale. Yet it might also imply that monetary coding is systematically available in the same environments where implementation is more transparent and outcomes are more measurable, which would create a correlation driven partly by observability. Considering this ambiguity, the chapter treats the overlap between monetary and count-based results as more credible than results that appear only under one exposure definition.

The inclusion of controls that proxy for global demand conditions and cycle effects tends to reduce, but not eliminate, the estimated exposure association for capacity outcomes.

One interpretation is that policy partly accompanies demand-driven expansions and then modestly amplifies them. Another interpretation is that the controls do not fully capture the relevant shocks, leaving residual confounding. The empirical evidence supports cautious language. It is possible that industrial policy is associated with capacity outcomes to some extent, but the data do not warrant an unqualified causal claim.

A central contribution of the dataset is that it codes policy targeting along the value chain. When exposure is decomposed by value-chain segment, the results suggest heterogeneity that is qualitatively consistent with bottleneck-sensitive interpretations. Exposure directed toward design and manufacture is more strongly associated with capacity proxies, while exposure directed toward inputs and equipment shows a clearer association with price-related and supply-stability proxies in some specifications. Exposure directed toward assembly, packaging, and testing is associated with trade composition shifts in a subset of countries, though these results are less stable across alternative specifications.

These patterns are suggestive but not definitive. A concern is that countries target different segments precisely when they face different shocks. For instance, a country may increase equipment support in response to anticipated tool shortages, and the observed association could reflect the severity of the shortage rather than the policy effect. Another concern is that targeting categories are more difficult to assign when policy packages are broad, which could attenuate true heterogeneity. The analysis therefore uses segment results mainly as a mechanism-consistency check rather than as a claim of precise channel identification. Heterogeneity by baseline capability also appears plausible. In countries closer to the technological frontier, exposure associations with capacity proxies appear more visible, which is consistent with the idea that complementary assets such as skilled labor, supplier ecosystems, and permitting capacity matter for translating support into output. Yet measurement asymmetry provides a competing explanation. Frontier economies often have better statistical coverage and clearer policy documentation, which can make both exposure and outcomes more observable. This is one reason the chapter avoids treating heterogeneity as a purely structural fact and instead frames it as a set of patterns that motivate deeper micro-level investigation.

The robustness checks are designed to test whether the main patterns persist under plausible alternative coding and timing rules. Results are examined under alternative surge definitions, alternative windows, and alternative transformations of outcomes. Placebo timing exercises are used to assess whether similar patterns appear when treatment is shifted to years where no surge is coded. The broad conclusion of these exercises is that the most consistent patterns are those linking exposure to capacity-related outcomes, while trade and innovation patterns are more sensitive to specification choice. However, robustness checks

do not fully resolve endogeneity. They can reduce the risk that results are driven by a small number of outliers or by a single coding choice, but they cannot fully separate policy effects from coincident shocks. The chapter therefore treats robustness not as a certificate of causality, but as a mapping of which empirical claims remain stable when the measurement and timing assumptions are perturbed.

4.4 SPILLOVER EVIDENCE AND INTERPRETATION

BOUNDARIES

Spillover estimates based on foreign exposure measures provide suggestive evidence that partner policy surges can be associated with changes in outcomes in highly exposed countries, though the patterns are not uniform across channels. In upstream dependent settings, the evidence is more consistent with price and sourcing adjustments than with immediate capacity expansion, which is plausible given that domestic fabrication capacity requires long lead times while import composition can change more quickly. In third market competition channels, some specifications suggest shifts in trade composition, yet these results are more sensitive to weighting schemes and to the treatment of re routed flows, which cautions against strong claims. Placebo weights and lead tests reduce, but do not eliminate, the concern that spillover exposure co moves with global shocks. Considering these factors, the spillover evidence is best read as a disciplined complement to the domestic results rather than as a definitive mapping of international incidence, and it highlights where micro level supply chain data would be most valuable for future identification.

5 CONCLUSION

This paper develops a replicable pipeline that converts policy text records into value chain aware measures of semiconductor industrial policy exposure, then links these measures to outcomes that are salient in the AI hardware era using complementary designs that emphasize timing diagnostics and sensitivity to measurement choices. The evidence is suggestive: documented exposure is possibly associated, to some extent, with gradual capacity related adjustments, while trade and innovation responses appear more heterogeneous and more specification sensitive, a pattern consistent with shifting bottlenecks and with private investment that can anticipate policy under shared demand and geopolitical shocks. By making treatment construction, value chain targeting, and surge timing assumptions explicit, the paper clarifies which patterns remain stable under alternative exposure definitions and which become fragile, which in turn suggests that strong welfare claims are premature. Further research is needed that combines clearer implementation timing, micro level eligibility variation, and operational metrics such as yield and reliability so that realized incidence can be separated more credibly from strategic announcements and coincident shocks.

ACKNOWLEDGMENTS

The authors thank the editor and anonymous reviewers for their helpful comments and valuable suggestions.

FUNDING

Not applicable.

INSTITUTIONAL REVIEW BOARD STATEMENT

Not applicable.

INFORMED CONSENT STATEMENT

Not applicable.

DATA AVAILABILITY STATEMENT

The original contributions presented in the study are included in the article/supplementary material, further inquiries can be directed to the corresponding author.

CONFLICT OF INTEREST

The authors declare that the research was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest.

PUBLISHER'S NOTE

All claims expressed in this article are solely those of the authors and do not necessarily represent those of their affiliated organizations, or those of the publisher, the editors and the reviewers. Any product that may be evaluated in this article, or claim that may be made by its manufacturer, is not guaranteed or endorsed by the publisher.

AUTHOR CONTRIBUTIONS

Not applicable.

ABOUT THE AUTHORS

YE, Kai

Hezhou University, 542899, China.

REFERENCES

- [1] Goldberg, P. K., Juhász, R., Lane, N. J., Forte, G. L., & Thurk, J. (2024). Industrial policy in the global semiconductor sector (No. w32651). National Bureau of Economic Research.

- [2] Huang, S. (2025). LSTM-Based Deep Learning Models for Long-Term Inventory Forecasting in Retail Operations. *Journal of Computer Technology and Applied Mathematics*, 2(6), 21-25. Li, K., Chen, X., Song, T., Zhang, H., Zhang, W., & Shan, Q. (2024). GPTDrawer: Enhancing Visual Synthesis through ChatGPT. *arXiv preprint arXiv:2412.10429*.
- [3] Bal, R., & Gill, I. S. (2020). Policy approaches to artificial intelligence based technologies in China, European Union and the United States.
- [4] Wang, H., Li, Q., & Liu, Y. (2023). Adaptive supervised learning on data streams in reproducing kernel Hilbert spaces with data sparsity constraint. *Stat*, 12(1), e514.
- [5] Wang, H., Sun, W., & Liu, Y. (2022). Prioritizing autism risk genes using personalized graphical models estimated from single-cell rna-seq data. *Journal of the American Statistical Association*, 117(537), 38-51.
- [6] Nazir, M., Rasheed, M. Q., Yu, X. H., & Ahmed, Z. (2025). Can Computer Technology, Semiconductors, and Artificial Intelligence Shape a Sustainable Future? Evidence From Leading Semiconductor - Producing Countries. *Sustainable Development*.
- [7] Rizi, A. D., Roy, A., Noor, R., Kang, H., Varshney, N., Jacob, K., ... & Asadizanjani, N. (2023). From talent shortage to workforce excellence in the chips act era: Harnessing industry 4.0 paradigms for a sustainable future in domestic chip production. *arXiv preprint arXiv:2308.00215*.
- [8] Janjeva, A., Baek, S., & Sellars, A. (2024). China's Quest for Semiconductor Self-Sufficiency.
- [9] Chen, Y. (2025). Artificial Intelligence in Economic Applications: Stock Trading, Market Analysis, and Risk Management. *Journal of Economic Theory and Business Management*, 2(5), 7-14.
- [10] Yin, M. (2025). Predictive Maintenance of Semiconductor Equipment Using Stacking Classifiers and Explainable AI: A Synthetic Data Approach for Fault Detection and Severity Classification. *Journal of Industrial Engineering and Applied Science*, 3(6), 36-46.
- [11] Yin, M. (2025). Data Quality Control in Semiconductor Manufacturing through Automated ETL Processes and Class Imbalance Handling Techniques. *Journal of Industrial Engineering and Applied Science*, 3(6), 13-22.
- [12] Yin, M. (2025). A Data-Driven Approach for Real-Time Bottleneck Detection and Optimization in Semiconductor Manufacturing Using Active Period Method and Visualization. *Academic Journal of Natural Science*, 2(4), 19-26.
- [13] Awad, H., & Bayoumi, E. H. (2025). Next-generation smart inverters: bridging AI, cybersecurity, and policy gaps for sustainable energy transition. *Technologies*, 13(4), 136.
- [14] Sun, Y., & Ortiz, J. (2024). An ai-based system utilizing iot-enabled ambient sensors and llms for complex activity tracking. *arXiv preprint arXiv:2407.02606*.
- [15] Pasricha, S., & Wolf, M. (2023). Ethical design of computers: From semiconductors to IoT and artificial intelligence. *IEEE Design & Test*, 41(1), 7-16.
- [16] Yin, M. (2025). Defect Prediction and Optimization in Semiconductor Manufacturing Using Explainable AutoML. *Academic Journal of Natural Science*, 2(4), 1-10.
- [17] Chen, Y. (2025). A Comparative Study of Machine Learning Models for Credit Card Fraud Detection. *Academic Journal of Natural Science*, 2(4), 11-18.
- [18] Zhu, H., Luo, Y., Liu, Q., Fan, H., Song, T., Yu, C. W., & Du, B. (2019). Multistep flow prediction on car-sharing systems: A multi-graph convolutional neural network with attention mechanism. *International Journal of Software Engineering and Knowledge Engineering*, 29(11n12), 1727-1740.
- [19] Fu, Q. (2025). Implementation Standards and Industry-Specific Adaptation of SAP-Lean Production Integration in Manufacturing: A Multi-Case Validation with Quantitative Performance Optimization. *Frontiers in Management Science*, 4(5), 67-7
- [20] Chen, Y. (2025). Daily Asset Pricing Based on Deep Learning: Integrating No-Arbitrage Constraints and Market Dynamics. *Journal of Computer Technology and Applied Mathematics*, 2(6), 1-10.
- [21] Pang, F. (2020, November). Research on Incentive Mechanism of Teamwork Based on Unfairness Aversion Preference Model. In *2020 2nd International Conference on Economic Management and Model Engineering (ICEMME)* (pp. 944-948). IEEE.
- [22] Wang, J., Tse, T. K., Li, S., & Fung, J. C. (2023). A model of the sea-land transition of the mean wind profile in the tropical cyclone boundary layer considering climate changes. *International Journal of Disaster Risk Science*, 14(3), 413-427.
- [23] Menaldo, V. The Progenitors of the Information Age and the Birth of AI: Introducing the "Political Economy of the Fourth Industrial Revolution".
- [24] Chen T. Innovative Approaches in Art Vocational Education: Exploring Industry-Academia Collaboration and Internationalization. *Journal of Advanced Research in Education*. 2025;4(5):11-7.
- [25] Chen, Y. (2025). Leveraging LSTM Networks for Vehicle Stability Prediction: A Comparative Analysis with Traditional Models under Dynamic Load Conditions. *Computing and Interdisciplinary Science*, 1(2), 15-22.
- [26] Wang, J., Cao, S., Tim, K. T., Li, S., Fung, J. C., & Li, Y. (2025). A novel life-cycle analysis framework to

assess the performances of tall buildings considering the climate change. *Engineering Structures*, 323, 119258.

- [27] Aboy, M., Minssen, T., & Kop, M. (2022). Mapping the patent landscape of quantum technologies: patenting trends, innovation and policy implications. *IIC-International Review of Intellectual Property and Competition Law*, 53(6), 853-882.