SUPPLY CHAIN RESEARCH · EDITORIAL FEATURE

The Nearshoring Arithmetic

Why the Model That Justifies a Relocation Is Not the Model That Decides It

A landed-cost comparison takes an afternoon and reaches the investment committee. The variables that determine whether the move works take months to establish and are mostly absent from it. In February 2026 the tariff assumption underneath every 2025 relocation model was held to have never been lawful.

Independent, vendor-neutral research for supply chain and technology leaders
Supply Chain Research | supplychainresearch.com | 2026

Contents

Why the Model That Justifies a Relocation Is Not the Model That Decides It

01 Executive summary

A relocation decision arrives at the investment committee as a landed-cost comparison. The current site costs a stated amount per unit, the proposed site costs less, and the difference multiplied by annual volume produces a saving that justifies the capital. The comparison is built from four numbers that are knowable in an afternoon: the hourly wage rate at each location, the freight cost per unit, the tariff rate applying on the day the model was built, and duty and brokerage. It is a coherent artefact, it is easy to review, and it is almost never the analysis that determines whether the relocation succeeds.

The variables that decide the outcome are absent from it, and they are absent for a structural reason: each requires site-specific investigation over months rather than a published rate. Yield loss during the production ramp, the depth of the second and third-tier supplier ecosystem near the new site, engineering support travel that continues for years, the cost of running both sites during qualification, working capital tied up against a line not yet proven, and the variance rather than the average of logistics reliability. Alongside these sits the assumption that has proved least stable of all. Through 2025, relocation models were anchored on an emergency-powers tariff schedule. In February 2026 the Supreme Court held by six to three that the statute never authorized tariffs at all, the administration substituted a different authority with a one hundred and fifty day statutory limit, and between one hundred and sixty-six and one hundred and seventy-nine billion dollars already collected became the subject of refund claims. We say honestly that nearshoring is frequently the right answer, for reasons of tariff exposure, lead time, intellectual property risk, and resilience that a unit-cost model does not capture at all. The argument here is that the arithmetic used to justify it is the wrong arithmetic, and that firms reaching the right conclusion are frequently doing so despite their model rather than because of it.

8.6%
of Mexico's record 2024 foreign direct investment was new greenfield investment, the lowest proportion since 1993

6-3
the Supreme Court majority holding in February 2026 that the emergency statute does not authorize tariffs

~1%
of Mexican export growth since 2018 explained by direct transshipment, according to Federal Reserve research

The argument in brief

- **The model is built from what is easy to obtain.** Wage rate, freight, tariff, and duty are published. Ramp yield, ecosystem depth, and transfer cost require months of site investigation, so they are omitted, and they are what decides the outcome.
- **The tariff assumption was never stable.** Every 2025 relocation model anchored on a schedule that a court held in February 2026 was not lawfully imposed. A ten-year model cannot rest on a rate that changed twice inside a week.
- **The headline investment number contains almost no new capacity.** Mexico's record 2024 inflow was 77.9 percent reinvested profits from firms already present. Greenfield was 8.6 percent, the lowest share in three decades.

- **The reallocation is real and it is not mostly transshipment.** Independent estimates put direct rerouting at a small fraction of the shift. The convenient dismissal is as wrong as the promotional claim.
- **Separate the four decisions.** Tariff hedge, lead time, intellectual property risk, and unit cost are different arguments with different evidence. Compressing them into one payback figure hides which one is carrying the decision.

02 The model that justifies, and the model that decides

The structural claim of this article can be stated before any evidence is assembled, and most executives who have actually moved a production line recognise it immediately. There are two analyses in every relocation decision: the one that gets built and reviewed, and the one that determines whether the move works. They overlap far less than anyone involved assumes. Figure 1 sets them side by side.

Figure 1. The relocation decision and the model used to justify it. A landed-cost comparison is buildable in an afternoon from four knowable numbers. The variables that determine whether a relocation succeeds are mostly absent from it.

Consider why the landed-cost comparison has the shape it does. It is assembled from inputs that exist in published, retrievable form. Hourly wage rates are available from statistical agencies and from the firms that provide relocation services. Freight costs per container are quoted daily. Tariff rates are in the schedule. Duty and brokerage are calculable. An analyst with a spreadsheet can produce a defensible comparison in a day, and that comparison has the two properties an investment process rewards: it is quantitative and it is reviewable by people who have never visited either site.

The variables in the right-hand column have the opposite properties. Nobody publishes the yield curve for a new line during its first eighteen months, because it depends on the specific process, the specific workforce, and the specific product. The depth of the second and third-tier supplier base near a candidate site is discoverable only by going there and walking it. Engineering support travel is a function of how much tacit knowledge lives in the current site's workforce, which nobody has ever measured. Each of these requires months of investigation, produces a range rather than a number, and cannot be reviewed by anyone who has not done the work. So the model that reaches the committee contains the first column, and the second column becomes a set of problems the operation discovers after the capital is committed.

This is not an argument that relocation decisions are usually wrong. Many are right, and the sections on fairness later in this article set out why. It is an argument that the decision and its justification have come apart. A firm that relocates for sound strategic reasons and builds a landed-cost model to satisfy its capital process has produced a document that does not describe its actual reasoning, and it will therefore not be able to tell afterward whether the reasoning held. A firm that relocates because the landed-cost model said so has made a decision on an analysis that omits most of what matters. Both are common, and the second is the more dangerous.

The remainder of this article does four things. It examines what the trade data actually shows, because the strategic case is frequently argued from figures that are either promotional or dismissive and rarely accurate. It examines what happened to the tariff assumption in February 2026, which is the clearest

available demonstration that the model's most important input was never stable. It sets out the six omitted costs in enough detail to be actionable. And it proposes a way of underwriting the decision that separates the four distinct arguments usually compressed into one number.

There is a further reason the two analyses diverge that is organizational rather than analytical. The landed-cost model is typically built by finance or by a strategy function, because it is a financial artefact and those functions own the capital process. The variables in the second column are known to manufacturing engineering, quality, and the plant management who have run transfers before, and those groups are frequently consulted after the case is approved rather than while it is being built. The knowledge exists inside the firm; it simply sits in a different function from the one holding the pen.

The corrective is procedural and costs almost nothing. A relocation case should require sign-off from whoever will actually run the transfer before it reaches the committee, with that person asked specifically what the ramp curve looked like the last time and what the new site will lack that the current one has. Firms that institute this find that the case changes materially, usually by extending the timeline and raising the transition cost, and that the resulting document is one the operation recognises rather than one it inherits.

03 What the trade data actually shows

Before examining the models, it is worth establishing what has actually happened in trade flows, because the public discussion is dominated by two claims that cannot both be true and are both overstated. Figure 2 shows the measured position.

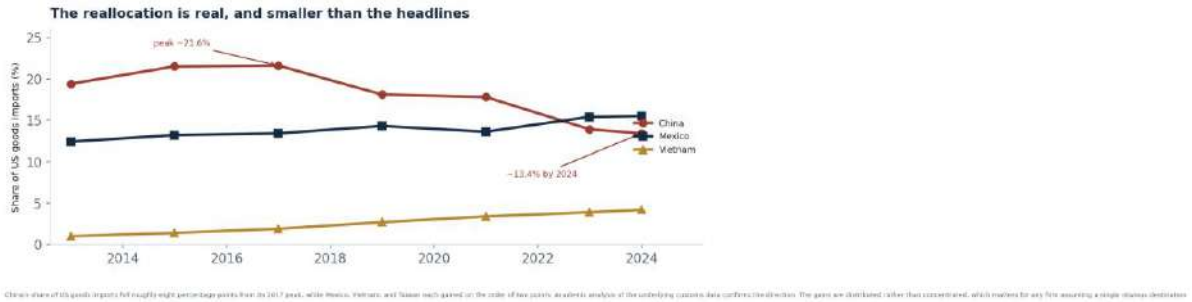

Figure 2. China's share of US goods imports fell roughly eight percentage points from its 2017 peak, while Mexico, Vietnam, and Taiwan each gained on the order of two points. The gains are distributed rather than concentrated.

The measured shift is substantial and it is not in dispute. Academic analysis of United States customs data by researchers publishing through the Centre for Economic Policy Research documents China's share of US goods imports falling from a peak of roughly twenty-one percent in 2017 to about thirteen percent by 2024, with Vietnam, Mexico, and Taiwan each gaining on the order of two percentage points. Mexico became the largest single source of US goods imports, at roughly fifteen and a half percent in 2024. This is a real reallocation of trade, visible in official statistics, and any firm assuming that nothing has changed is mistaken.

Two features of the pattern matter for a firm building a relocation case. The first is that the gains are distributed rather than concentrated: no single destination absorbed the shift, which means the assumption that one obvious location is emerging as the replacement is not supported. Mexico, Vietnam, Taiwan, India, and others each captured a portion, and the firm's own decision has to be made on its own product, process, and supply base rather than on a general direction of travel. The second is the timing. The largest movements coincided with tariff changes rather than with any underlying cost convergence, which is a warning about durability: a shift driven by policy is a shift that policy can reverse, and the following sections examine what happened when it did.

It is also worth noting what the data does not show. It does not show a general return of manufacturing to the United States at a scale visible in the import statistics, because the share loss went overwhelmingly to other exporting countries rather than to domestic production. United States manufacturing construction spending did rise sharply, by roughly fifty percent in 2022 and sixty-two percent in 2023, but it fell by around five percent in 2025 and was on track for a further decline in 2026 according to Census Bureau series. The construction boom was real, it was concentrated in semiconductors and batteries under specific federal programmes, and it has turned. A firm building a relocation case on the premise of a broad and continuing reshoring wave is building on a series that has already reversed.

One further feature of the reallocation deserves attention because it complicates the planning assumption most firms make. The shift has been uneven across product categories rather than uniform across trade. Categories with simple assembly, low capital intensity, and portable process knowledge moved readily; categories requiring deep supplier ecosystems, specialised tooling, or long-qualified processes moved slowly or not at all. A firm reasoning from the aggregate share statistics to its own category may therefore be reasoning from a population dominated by products unlike its own.

The practical instruction is to establish what has happened in the specific product category rather than in trade overall, which is available from the same customs data at a more granular level. Where the category has moved substantially, the firm has evidence that the move is feasible and can seek out the operations that made it. Where it has not moved despite tariff pressure, that is informative in the opposite direction: it suggests the barriers are structural rather than inertial, and the firm should establish what those barriers are before assuming it can overcome them.

04 The transshipment question

The most common objection to the reallocation data is that it is largely cosmetic: that Chinese goods are being routed through Mexico and Vietnam with minimal processing, so the country-of-origin statistics have changed while the actual supply chain has not. The objection is worth taking seriously and the evidence does not support it in the strong form. Figure 3 assembles the estimates.

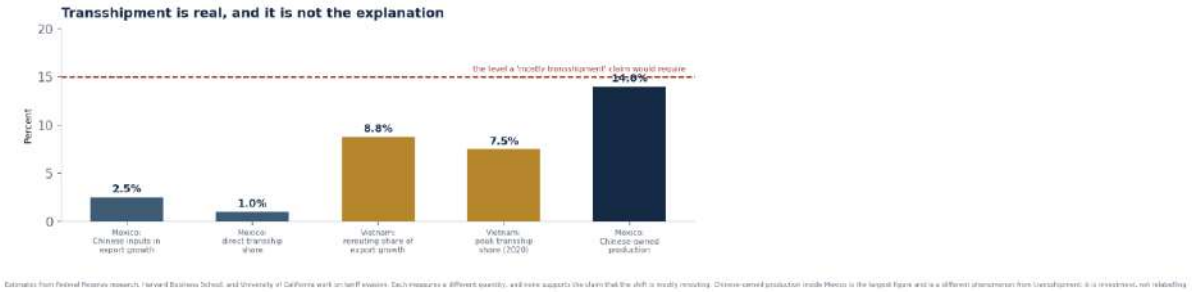

Figure 3. Estimates from Federal Reserve research, Harvard Business School, and University of California work on tariff evasion. Each measures a different quantity, and none supports the claim that the shift is mostly rerouting.

Federal Reserve research published in 2025 examined the Mexican case and concluded that direct transshipment explains only about one percent of Mexican export growth to the United States since the trade war began, with Chinese inputs accounting for a further two to three percent of recent export growth. Separately, Chinese-owned production located inside Mexico accounts for a materially larger figure, around fourteen percent, but that is a different phenomenon: it is Chinese firms investing in Mexican production capacity, which involves real plant, real employment, and real value added within Mexico, and it is not relabelling.

The Vietnamese case has been studied by researchers at Harvard Business School, who estimated that around nine percent of Vietnam's export growth to the United States between 2018 and 2021 was attributable to rerouting rather than to genuine production. Work at the University of California examining tariff evasion through product-level tracking put the peak transshipment share of Vietnam's US exports at around seven and a half percent in 2020, with the Mexican figure never exceeding one and a half percent. These estimates use different methods and measure slightly different quantities, and they converge on the same conclusion: rerouting is real, it is detectable, and it is a modest fraction rather than the explanation.

There is a countervailing signal that deserves reporting alongside these figures. A commercial trade-intelligence firm reported that suspect rerouting transactions rose by around seventy-six percent in the first ten months of 2025, from roughly one hundred million to one hundred and eighty-eight and a half million. That figure comes from a vendor selling trade-intelligence services and its methodology is proprietary, so it is an interested estimate rather than a measurement, and it is flagged as such. It is nonetheless directionally consistent with the observation that enforcement attention increased sharply in the same period, and with the introduction of penalty rates specifically targeting transshipment, including a framework applying a substantially higher rate to goods determined to have been rerouted rather than produced.

For a firm building a relocation case, the practical implication is narrower than the general debate. If the plan involves establishing production in a country that has attracted rerouting scrutiny, the origin determination is a live compliance question rather than a settled one, and the substantial-transformation and tariff-shift rules governing it are technical and enforced. A relocation that produces a product whose origin could be challenged has acquired a regulatory exposure that no landed-cost model contains, and the penalty differential between a compliant origin and a determined transshipment is large enough to invalidate the economics of the move entirely.

There is a second-order consequence of the enforcement environment that relocation planners should price explicitly. Where origin determinations become contested, the administrative burden of proving origin falls on the importer, and that burden is substantial: bills of materials, production records, and supplier declarations traced through multiple tiers, maintained continuously rather than assembled on demand. Firms that relocate into jurisdictions under scrutiny acquire a documentation obligation that did not exist at the prior site, and staffing it is a recurring cost that no landed-cost model contains.

The obligation is also asymmetric in its consequences. A firm whose origin claim is upheld gains nothing beyond the rate it expected; a firm whose claim is rejected faces retroactive duty, penalties, and potential exclusion of goods, which can arrive years after the production occurred. That asymmetry means the origin question should be settled before the site decision rather than after, ideally through a binding ruling where the jurisdiction offers one. Firms that treat origin as a compliance detail to be resolved during implementation have deferred the one question that can invalidate the entire economic case.

05 A record FDI number with no factories in it

The single most misread statistic in this field is Mexican foreign direct investment, which has been reported as evidence of a nearshoring boom while its composition says something close to the opposite. Figure 4 decomposes it.

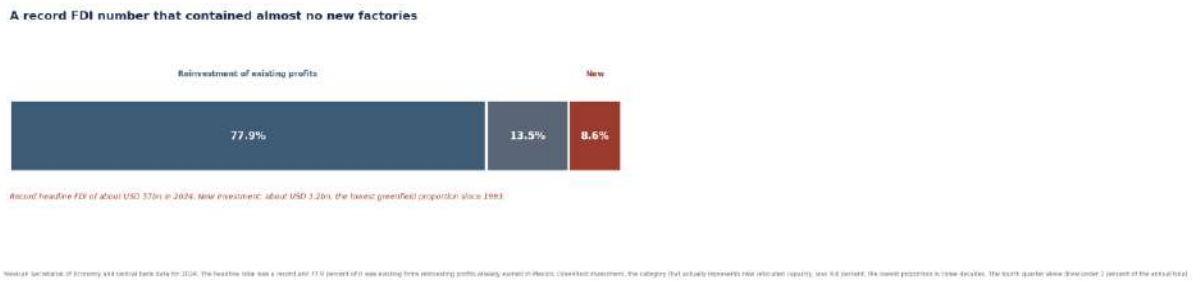

Figure 4. Mexico's record 2024 inflow was 77.9 percent reinvestment of profits by firms already present. Greenfield investment, the category representing truly new relocated capacity, was 8.6 percent, the lowest proportion since 1993.

Mexican official data for 2024 recorded total foreign direct investment of approximately thirty-seven billion dollars, a record in nominal terms and widely reported as confirmation that nearshoring was arriving. The composition tells a different story. Roughly seventy-eight percent of the total, about twenty-nine billion dollars, was reinvestment of profits by companies already operating in Mexico. Intercompany accounts made up a further portion. New investment, the category that captures a firm establishing capacity it did not previously have, was approximately eight and six tenths percent, around three point two billion dollars, and the lowest proportion of the total recorded since 1993.

The distinction matters enormously for anyone reasoning about relocation. Reinvested profits are what an established operation does with money it has already earned in the country: expanding a line, upgrading equipment, adding a shift. It is a sign that existing operations are performing, which is plainly positive, and it is not evidence that new firms are relocating. Greenfield investment is the measure of new arrivals, and it was at a three-decade low as a share of the total in the same year the nearshoring narrative peaked. A firm reading the headline number as market validation for its own relocation has read a number that does not measure the thing it needs.

The quarterly pattern reinforces the point. The fourth quarter of 2024 drew under two percent of the annual total, roughly six hundred and seventy-six million dollars, the weakest quarterly inflow in many years and coinciding with the United States election and the uncertainty about trade policy that accompanied it. Investment decisions of this kind are made months in advance, and a collapse of that magnitude in a single quarter indicates that firms were deferring commitments rather than accelerating them. Subsequent data showed manufacturing-specific inflows recovering through 2025, but the composition question persists and any firm citing Mexican FDI as evidence should cite the greenfield component rather than the headline.

There is a broader methodological lesson here that recurs throughout this article. Headline aggregates in trade and investment data are frequently composed of categories that behave differently and mean different things, and the composition is available in the same publications that report the total. The headline is what circulates; the composition is what informs a decision. A firm whose relocation case

cites any aggregate figure should establish what the aggregate contains before relying on it, and in this instance the exercise takes about ten minutes and changes the interpretation completely.

06 February 2026: the assumption becomes void

Every relocation model contains a tariff assumption, because the whole point of the exercise is comparing landed cost across jurisdictions and tariffs are a component of landed cost. Through 2025 that assumption was drawn from an emergency-powers schedule. What happened to it is the clearest available demonstration that the model's most consequential input was never a stable number. Figure 5 sets out the sequence.

Figure 5. A ten-year capital model requires a tariff assumption. In February 2026 the Supreme Court held by six to three that the statute relied upon does not authorize tariffs, and a different authority with a 150-day limit was substituted within days.

On 20 February 2026 the Supreme Court decided the consolidated cases challenging tariffs imposed under the International Emergency Economic Powers Act, holding by six votes to three that the statute does not authorize the President to impose tariffs. The Court applied the major questions doctrine, reasoning that it operates with particular force where the asserted delegation involves the core congressional power over revenue. The decision invalidated the legal basis for a schedule under which somewhere between one hundred and sixty-six and one hundred and seventy-nine billion dollars had already been collected, and importers began filing refund claims in substantial numbers.

The administration substituted a different authority within days. A surcharge was imposed under Section 122 of the Trade Act of 1974, announced initially at ten percent and rising, effective 24 February 2026, and that provision carries a statutory limit of one hundred and fifty days. Other authorities were unaffected: Section 232 measures on steel, aluminium, and copper remained in place, as did Section 301 measures relating to China. The net effect for a firm modelling landed cost was that the applicable rate changed twice within a week, the replacement carried an expiry date measured in months, and the underlying question of which authority can lawfully impose what remained contested.

The implication for capital modelling is not that tariffs will fall or rise, which nobody can forecast, but that the tariff input to a ten-year model is qualitatively different from the other inputs and should be treated differently. Wage rates move gradually and predictably enough to model with a growth assumption. Freight rates are volatile but mean-reverting and hedgeable. A tariff rate is a policy variable that can change by executive action overnight, can be litigated, and can be held retroactively to have never been lawfully imposed. Modelling it as a point estimate for a decade is not conservative or aggressive; it is a category error about what kind of number it is.

What follows in practice is a change in how the input is handled. The tariff assumption should be modelled as a range spanning plausible policy outcomes, including the case in which the differential the relocation was built to capture disappears entirely. The firm should establish what happens to the return

if the tariff differential goes to zero, because that scenario has now occurred once within the modelling horizon of most current relocation programmes. And where the relocation is justified primarily by the tariff differential rather than by cost, lead time, or risk, the firm should recognise that it is taking a position on trade policy rather than making an operational decision, and should size the commitment accordingly. That is a defensible thing to do knowingly and an uncomfortable thing to discover afterward.

07 The six costs the comparison omits

Turning from the external environment to the model itself, the omissions are consistent across the relocation cases this research examined, and they are specific enough to be checked. Figure 6 lists them.

Figure 6. Six costs that decide the outcome and are absent from the comparison. None is exotic; each requires site-specific investigation rather than a published rate, which is why the model omits them.

Yield loss during the ramp is the largest and the most consistently underestimated. A new line staffed by a new workforce running an established process does not achieve the yield of the mature site immediately, and the gap can persist for many months. The cost is scrap, rework, expedited freight to cover shortfalls, and management attention, and it accrues at exactly the point when the business case assumed savings would begin. Firms that have run transfers can generally quote their own historical ramp curves, which makes this the easiest of the six to estimate properly, and the estimate is frequently the difference between a three-year payback and a six-year one.

Ecosystem depth is the least visible and the most likely to be decisive. A mature manufacturing location has an accumulated base of second and third-tier suppliers: the toolmaker who can turn a fixture in two days, the plating shop, the calibration service, the specialist maintenance contractor. These do not appear in any cost model, they are taken for granted at the existing site because they have always been there, and they may not exist near the new one. Their absence converts routine problems into expedited international shipments and multi-week delays. Establishing whether they exist requires walking the local supplier base, and firms that skip this step discover the gap in the first quarter of operation.

The remaining four are more tractable and equally absent. Engineering support travel continues for years and reflects how much undocumented process knowledge lives in the existing workforce, which is always more than the documentation suggests. Dual running means paying for both sites throughout qualification, which for a regulated or safety-critical product can run to many quarters. Working capital rises because inventory must be held against a line not yet proven, and that inventory is financed. And logistics reliability should be modelled on its variance rather than its average: a route with a good average transit time and a long tail imposes buffer inventory and expediting costs that the average conceals entirely. A model that includes the average and excludes the tail has understated the cost of the less reliable route.

A practical way to test whether a case has addressed these is to ask what the model assumes about the first eighteen months specifically. A relocation model that shows savings beginning in month one has almost certainly not incorporated the ramp, because no transfer achieves target yield immediately. A

model that shows savings beginning at a stated month, with a stated yield curve leading to it and a stated cost of the shortfall in the interim, has been built by someone who has done this before. The shape of the first two years is the single most diagnostic feature of a relocation business case and it takes a reviewer thirty seconds to check.

The same test applies to the transition period cost line. If the model contains no line for dual running, engineering travel, or expedited freight during qualification, those costs have not been forgotten by the operation, which will incur them regardless; they have simply been excluded from the document on which the capital decision rests. The gap will surface later as budget variance attributed to execution problems rather than to a business case that never included them, which is how a structurally optimistic model perpetuates itself across successive projects.

08 The wage rate is not the labour cost

The single input carrying most of the benefit in a relocation model deserves separate treatment, because it is both the most cited figure in the field and the one with the weakest provenance. Figure 7 shows the problem.

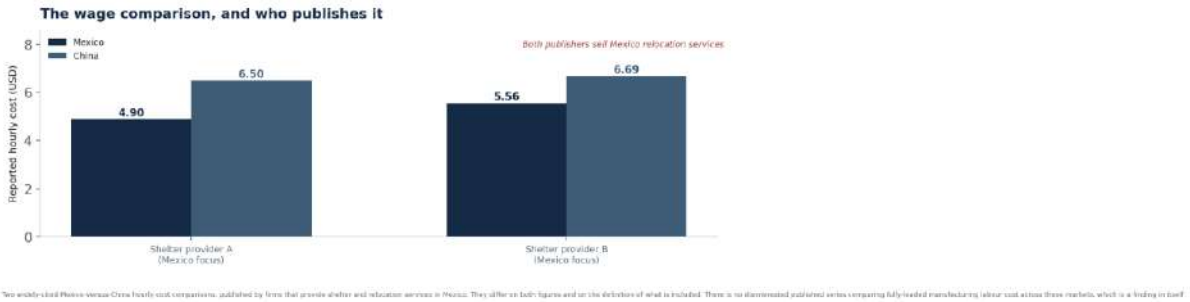

Figure 7. Two widely-cited Mexico-versus-China hourly cost comparisons, published by firms that provide relocation services in Mexico. They differ on both figures and on what is included.

The comparisons in circulation come predominantly from shelter-service providers, which are firms that establish and operate manufacturing facilities in Mexico on behalf of foreign companies. One such provider cites a Mexican rate of around four dollars ninety per hour against a Chinese rate of six dollars fifty, described as twenty-five percent lower. Another, using payroll data from its own operations, gives an entry-level Mexican figure of five dollars fifty-six fully fringed against a Chinese figure of six dollars sixty-nine described as fully loaded. Both are plausible, they differ on both sides of the comparison, and both are published by firms whose business is persuading companies to manufacture in Mexico. No disinterested published series comparing fully-loaded manufacturing labour cost across these markets was located, which is itself worth stating as a finding.

The deeper problem is that hourly rate is the wrong unit even when the number is accurate. What determines cost is labour input per unit of output, which combines the rate with productivity, and productivity differs across locations for reasons including equipment, process maturity, workforce experience, absenteeism, and turnover. A location with a rate forty percent lower and productivity thirty percent lower delivers a much smaller advantage than the rate comparison implies, and in some processes delivers none. Productivity-adjusted unit labour cost is the correct measure, comparisons of it are scarce, and the ones in circulation come from the same interested sources as the rate comparisons.

Turnover deserves specific attention because it is where relocation models most often go wrong in the destination market. Manufacturing labour markets in the principal nearshoring destinations have experienced high turnover in periods of rapid industrial growth, and where several plants compete for the same workforce in the same industrial park the rate can be severe. High turnover imposes recruitment cost, training cost, a productivity ramp for every replacement, elevated error rates during that ramp, and supervisory load, none of which appears in an hourly rate. A firm modelling the destination on the rate alone has modelled the cheapest possible version of an operation it has not yet run.

The constructive response is to build the labour input from the firm's own experience rather than from published comparisons. Most companies contemplating a transfer already operate multiple sites and can observe how labour cost per unit differs across them and why. That internal comparison is better evidence than any external series, because it holds the product and process constant and because the firm can interrogate it. Where the firm has no comparable internal reference, the honest position is that the labour saving is an estimate with a wide band, and the model should say so rather than presenting a rate differential as though it were the answer.

Energy and industrial input costs deserve mention alongside labour, because they have become a larger differentiator than they were when most relocation frameworks were designed. Industrial electricity pricing, grid reliability, and the availability of interconnection capacity vary widely across candidate locations and have moved substantially in recent years. For energy-intensive processes these can exceed the labour differential in magnitude, and unlike labour they are frequently subject to long lead times for capacity provision that constrain the schedule as well as the cost.

The practical instruction is that a candidate site assessment should establish delivered industrial energy cost, contracted reliability, and the lead time for any additional capacity the operation will require, at the same level of specificity applied to labour. Where the process is energy-intensive this should carry equal weight in site selection. Firms that have discovered an interconnection queue measured in years after committing to a site have learned that some inputs cannot be procured on the schedule a capital plan assumes, and that no amount of labour saving compensates for a line that cannot be powered.

09 Announcements, hires, and completions

A firm building a relocation case will encounter figures purporting to show how many companies have already done it and how many jobs have resulted. Those figures require careful handling, because the gap between what is announced and what is completed is one of the widest in this field. Figure 8 shows it.

Figure 8. Indexed presentation of figures published by a reshoring advocacy organization that also sells total-cost advisory services. The organization compiles announcements from press reporting rather than from payroll data.

The most widely cited reshoring statistics come from an advocacy organization that compiles announcements from press reporting and publishes annual totals. Its figures indicate roughly two million jobs announced since 2010, with around two hundred and forty-four thousand announced in 2024. The organization also sells total-cost-of-ownership advisory services and exists to promote reshoring, which makes it an interested source, and its methodology is announcement-scraping rather than employment measurement. Its own published material is candid about the limitation, noting that of the roughly two million announced positions only about one point seven million appear to have been filled, and that a small proportion of companies stating reshoring plans complete them.

A related qualification concerns what kind of jobs the announcements represent. The same organization reports that the large majority of recently reshored positions fall into higher-technology categories, which means low headcount per dollar of investment. That is consistent with what the construction data shows: the capital went into semiconductor and battery facilities under specific federal incentive programmes, which are capital-intensive and employ relatively few people per dollar invested. A firm reasoning about labour markets from these announcements is reasoning from a sample dominated by a small number of extremely large, incentive-driven, capital-intensive projects that resemble its own situation hardly at all.

The announcement-to-completion gap has a practical use for a firm building its own case, which is as a calibration on optimism. If the visible population of relocation announcements completes at a low rate, and the announcing firms had access to the same analysis this article is describing, then the base rate for completion is a warning about the reliability of relocation planning generally rather than about any particular firm's competence. The disciplined response is to ask what the firm's own plan assumes about schedule and to compare that against the observed distribution of how long these projects actually take, which is considerably longer than announcements imply. The semiconductor facilities announced at the start of this decade are illustrative: the interval from announcement to full staffing is measured in years rather than quarters.

10 What the filings record

Announcements are promotional documents. The reliable record of what happened to relocation and capacity decisions sits in securities filings, where firms describe closures, write-offs, and restructuring under a legal obligation to be accurate, and where the language is correspondingly unadorned.

Two examples illustrate the class of evidence available. A bearing manufacturer disclosed in its annual report the closure of a North Carolina plant it had acquired in early 2023 and closed in 2025, affecting roughly sixty employees, with cumulative pretax costs of approximately four point nine million dollars associated with the action. A paper products company disclosed idlings and associated write-offs at multiple facilities. Neither disclosure is dramatic and neither made significant news, which is exactly why they are useful: they record capacity decisions that were made, did not work, and were reversed, in the routine language of a filing rather than the language of a case study.

Larger and more visible programmes show the same pattern of timeline extension rather than reversal. A major semiconductor facility announced early this decade has seen its production timeline extended by several years. Tracking of announced clean-energy and electric-vehicle manufacturing projects by an environmental advocacy organization, which is an interested source and is flagged as such, identified dozens of projects cancelled, closed, or scaled back since 2022, with a substantial concentration of those cancellations occurring in the most recent year, representing tens of thousands of announced jobs and tens of billions of dollars of announced investment. One vehicle manufacturer cancelled a multi-billion-dollar plant expansion.

The reason to assemble these rather than to rely on case studies is the selection problem. Vendor and consultancy case studies describe relocations that worked, because firms that had a difficult experience do not participate in case studies and the advisers who supported them do not publicise the engagement. Filings capture the other tail, because disclosure obligations do not permit a firm to omit a plant closure. Neither source is a sample, and this article claims no rate. What the filings establish is that capacity decisions of exactly the kind under discussion are reversed at material cost by competent firms with professional advisers, which is the standard a business case exercise needs: evidence that the failure mode is real enough to model, rather than a probability that nobody can compute.

A related class of evidence sits in the timelines rather than the reversals, and it is more common. Large facilities announced with a stated production date are routinely reported later with the date extended, and the extensions are disclosed matter-of-factly in filings and investor communications rather than presented as setbacks. A firm assembling a relocation schedule should treat these as the relevant base rate for its own planning, because the announcing companies were large, well-advised, and financially capable, and their schedules still moved by years rather than months.

The mechanism behind schedule extension is worth naming because it is the same across cases: construction and equipment installation proceed roughly to plan, and the qualification, staffing, and yield-ramp phases take substantially longer than assumed. This maps directly onto the omitted costs catalogued earlier. The delay is not a construction problem but a start-up problem, and start-up is precisely the phase the landed-cost model treats as instantaneous. A firm that budgets realistically for

start-up will produce a longer schedule than its peers announce and a more accurate one than they achieve.

11 The fairness case: the strategic logic holds

This article has been critical of the analysis supporting relocation decisions, and a reader who concluded that nearshoring is generally unwise would have taken away something the evidence does not support and this article does not intend.

The strongest point is lead time, and it is not an argument the landed-cost model can make. Production in northern Mexico reaches United States markets in a day or two by road, against three or more weeks by ocean from East Asia. That compression changes what the business can do: it reduces the inventory required to cover replenishment, it shortens the cycle from design change to delivered product, it permits smaller and more frequent orders, and it reduces the exposure to a demand forecast made months in advance. Those are real economic benefits, they are substantial for products with short life cycles or volatile demand, and a unit-cost comparison captures none of them.

The second point is that tariff volatility, which this article has treated as a reason to distrust the model, is simultaneously a reason to consider regional production. If the tariff regime governing a long ocean supply chain can change by executive action, and the February 2026 sequence demonstrates that it can change in both directions and be litigated, then a firm with production inside the destination market or under a regional trade agreement has reduced its exposure to that variable regardless of which way it moves. The correct conclusion from tariff instability is not that relocation is unwise but that relocation justified by a specific tariff differential is fragile while relocation justified by reducing tariff exposure generally is robust. Those are different arguments and firms routinely conflate them.

A third point concerns the underlying investment position, which is stronger than the year-to-year noise suggests. Foreign direct investment into United States manufacturing has grown substantially over the past fifteen years in stock terms, and Mexican manufacturing-specific inflows continued rising through 2025 even as the composition question discussed earlier persisted. Peso movements have kept Mexican effective labour cost inflation moderate in dollar terms. Firms that established regional capacity have in many cases found it valuable for reasons they did not model, including proximity for engineering support and easier management oversight, which are the mirror image of the omitted costs this article has catalogued.

A fourth point deserves emphasis because it constrains the whole argument. Intellectual property risk, geopolitical concentration, and single-country dependency are genuine strategic exposures that a firm may reasonably decide to reduce at a cost, and the fact that a decision cannot be justified on unit cost does not make it wrong. A board that decides it is unwilling to concentrate critical production in one jurisdiction has made a risk judgment, and risk judgments do not require a positive net present value to be correct. The fair synthesis is that the strategic case for regional production is frequently sound, that the arithmetic offered in support of it is frequently not, and that the remedy is to argue the actual case rather than to dress a strategic judgment in a cost model that does not survive examination.

A fifth point in fairness concerns the counterfactual most relocation critiques ignore, which is the cost of staying. A production base concentrated in a single distant jurisdiction carries an exposure that does not appear in any current cost line and becomes visible only when it is realised: a port closure, an export

restriction, a sudden tariff, a regulatory action against a supplier, or a geopolitical event. Firms that experienced the disruptions of the early part of this decade have direct evidence of what that exposure costs when it materialises, and evidence of that kind is a legitimate input to a capital decision even though it cannot be expressed as an expected value.

The honest framing is therefore not that relocation must clear a cost hurdle but that concentration and dispersion each carry costs, one visible and one contingent, and the firm is choosing between them rather than choosing whether to spend. Stated that way, the decision becomes a question about risk appetite and the shape of the exposure the firm is willing to hold, which is a question a board is well equipped to answer. Stated as a payback calculation, it becomes a question about a forecast, which is a question nobody can answer well.

12 Separate the four decisions

The constructive principle that follows is to stop compressing four distinct arguments into a single payback figure, because they rest on different evidence, carry different risks, and would be evaluated differently if they were separated.

The first is the unit-cost argument: that producing in the new location is cheaper per unit at steady state. This is the argument the model actually addresses, it is the weakest of the four once the omitted costs are included, and it is the one most sensitive to the labour, productivity, and ramp assumptions examined earlier. A firm should state what it believes the steady-state unit-cost differential to be, with a band, after including transfer costs amortised over a stated period, and should be prepared for that band to be wide.

The second is the tariff-exposure argument, which has two forms that must be distinguished. The narrow form is that a specific current tariff differential makes the new location cheaper, and that form is fragile for the reasons Figure 5 demonstrates. The broad form is that concentrating production in a jurisdiction subject to volatile trade policy is itself a risk worth reducing, and that form is robust and survives policy changes in either direction. A firm should say which form it is relying on, because a case built on the narrow form should be sized as a policy bet and a case built on the broad form should be evaluated as risk reduction.

The third is the lead-time argument, which is measurable and frequently the strongest. It should be quantified directly: the inventory reduction achievable from shorter replenishment, the value of a shorter design-to-delivery cycle for the specific product, and the reduction in forecast exposure. These are calculable from the firm's own data and they are frequently worth more than the unit-cost differential in categories with short life cycles or volatile demand. The fourth is the risk argument covering intellectual property, concentration, and geopolitical exposure, which resists quantification and should be stated as a judgment with the exposure named rather than converted into a number that implies a precision it does not have.

Separating them changes the governance of the decision in a specific and useful way. A committee presented with a single payback figure can only compare it against a threshold. A committee presented with four arguments, each with its own evidence and its own fragility, can ask which one is carrying the decision and interrogate that one properly. In most relocation cases this exercise reveals that the unit-cost argument was doing the presentational work while a lead-time or risk argument was doing the actual persuading, and once that is visible the firm can evaluate the argument it is actually acting on.

There is a fifth consideration that sits underneath the four arguments and is worth separating explicitly, which is customer requirement. In some sectors, particularly automotive, defence, and increasingly semiconductors and medical devices, a customer or a regulator may require production within a specified region as a condition of the contract or of market access. Where that is the case, the relocation is not an optimisation at all but a condition of continuing to hold the business, and evaluating it on payback is a category error.

Firms in that position should say so plainly rather than constructing a cost case that will not survive scrutiny. The relevant analysis is the value of the business being protected against the cost of the compliance, which is a straightforward comparison and usually a decisive one. Presenting a requirement as an optimisation obscures the actual decision and, more damagingly, prevents the firm from negotiating properly with the customer whose requirement is imposing the cost, because the firm has represented the move as something it wanted to do anyway.

13 A relocation protocol, and a scoring rubric

The principles above combine into a protocol for underwriting a relocation and a rubric a board or investment committee can apply to judge whether the case in front of it has been built properly. Figure 9 sets out the discipline.

Figure 9. Underwriting a relocation as four decisions, not one number: separate the arguments, model the tariff as a range, cost the transfer rather than the steady state, verify the supplier ecosystem, and stage the commitment.

The protocol runs as follows. State the four arguments separately, with the evidence and the fragility of each, and identify which one is carrying the decision. Model the tariff input as a range spanning plausible policy outcomes including the disappearance of the differential entirely, and state what the return becomes in that case. Cost the transfer itself, including ramp yield loss, dual running, engineering travel, working capital, and retraining, rather than modelling only the steady state. Verify the second and third-tier supplier ecosystem near the candidate site by visiting it rather than by assuming it. Build the labour input from the firm's own multi-site experience rather than from published rate comparisons, and include turnover. And prefer a staged commitment that preserves the ability to stop.

A scoring rubric

The dimensions below distinguish an underwritten relocation from a justified one.

Dimension	Underwritten	Justified only
Argument structure	Four arguments, separately evidenced	One payback figure
Tariff input	Range, including differential going to zero	The rate on the day the model was built
Labour input	Own multi-site data, productivity and turnover	Published hourly rate comparison
Transfer cost	Ramp, dual running, travel, working capital	Capital expenditure and moving costs
Supplier ecosystem	Walked and verified before site selection	Assumed to exist
Evidence base	Official statistics and filings	Advocacy totals and provider comparisons
Commitment shape	Staged, with a defined stop point	Single irreversible programme

A case scoring in the left column tells the committee which argument it is being asked to accept and what would have to happen for that argument to fail. A case scoring in the right column tells the committee a number assembled from published rates and a tariff schedule that may not survive the year. The rubric does not make relocation less attractive; applied properly it strengthens cases built on lead time and risk, which are the arguments most likely to be sound and least likely to appear in the model.

14 Staging, and the option to be wrong

A dimension that scarcely appears in relocation business cases, and that deserves weight alongside the return, is the shape of the commitment: whether the firm has preserved any ability to change course as information arrives.

Relocation decisions are frequently structured as a single programme with a defined end state, because that is how capital programmes are conventionally scoped and because vendors, site-selection advisers, and development agencies all prefer a committed customer. The alternative structure is a sequence: qualify a small volume at the new site while maintaining the existing line, establish the actual ramp curve and the actual supplier availability, transfer a second tranche once the first is stable, and retain the option to stop at each boundary. This costs more in aggregate if everything goes well, because dual running is expensive and the transfer takes longer, and it costs dramatically less if any of the five omitted-cost categories turns out worse than assumed.

The value of the staged structure is highest precisely where the evidence in this article suggests uncertainty is greatest. The ramp curve and the ecosystem depth are the two variables the model cannot establish in advance, and they are the two that a first tranche reveals directly. A firm that has transferred ten percent of volume and measured its actual yield, its actual expediting cost, and its actual reliance on distant suppliers knows something the model could not have told it, and it knows it before committing the remaining ninety percent. That is a substantial reduction in the risk of the largest capital commitment, purchased at the cost of a slower and less elegant programme.

The objection is that staging forfeits scale economics and prolongs the period of dual cost, and that objection is legitimate and should be evaluated rather than dismissed. Where the firm has high confidence in its volume outlook, has transferred comparable processes before, and has verified the supplier ecosystem, an integrated single-step transfer may well be correct and the staging premium is waste. The test is whether the firm can state, from its own prior experience, what the ramp curve looks like. A firm that can has earned the right to commit in one step. A firm that cannot is proposing to discover its ramp curve at full scale, and staging is the cheaper way to learn the same thing.

There is a further and less obvious benefit to preserving the existing site through the transition, which is that it maintains the option to reverse. The filings examined earlier record capacity decisions unwound at material cost, and in most of those the reversal was expensive precisely because the original capability had been dismantled. A firm that has kept a line warm can return to it; a firm that has sold the equipment and released the workforce cannot, and its only remaining options are to persist with a site that is not working or to rebuild from nothing. The cost of holding the option is visible and quantifiable, the cost of not holding it appears only in the scenario where it is needed, and business cases systematically overweight the first because it is the one that shows up in the budget.

A final observation about staging concerns how it interacts with the tariff argument specifically. Where a relocation is motivated wholly or partly by a tariff differential, and where that differential has been demonstrated to be capable of disappearing by judicial decision within a single week, the case for preserving optionality is at its strongest. A staged programme allows the firm to observe whether the

policy environment stabilises before committing the balance of the capital, and the cost of that patience is a delay measured in quarters against an exposure measured in the full programme value. Firms that committed fully to tariff-motivated relocations during 2025 did not have that option available to them by February 2026.

15 Conclusion: the arithmetic was never the argument

Relocating production is among the largest and least reversible decisions a manufacturing organization makes. It commits capital, it reshapes the supplier base, it moves employment, and it determines the cost structure for a decade. The analysis that supports it is typically a landed-cost comparison assembled from four published numbers, and the argument of this article is that this comparison is not the analysis the decision requires, that the variables it omits are the ones that determine the outcome, and that its most important input was demonstrated within the past year to be unstable in a way no model had contemplated.

The evidence assembled here is drawn from official statistics, court records, and securities filings rather than from case studies or advocacy totals. It shows a real reallocation of trade, with China's share of United States goods imports falling roughly eight points from its peak while several destinations each gained a couple of points. It shows that the convenient dismissal of that shift as transshipment is not supported, with independent estimates putting direct rerouting at a small fraction of the movement. It shows a record Mexican investment figure composed overwhelmingly of reinvested profits from firms already present, with new greenfield investment at its lowest proportion since 1993. It shows a construction boom in United States manufacturing that has already turned. And it shows a Supreme Court holding, in February 2026, that the tariff authority underpinning a year of relocation modelling did not permit tariffs at all.

What follows for an operations or supply chain leader is not that regional production is unwise. This article has set out the serious case that lead time compression is valuable and unmodelled, that reducing exposure to volatile trade policy is sound whichever direction policy moves, and that concentration and intellectual property risks are real exposures a board may properly decide to reduce at a cost. What follows is that these are the arguments, that they should be made explicitly and evidenced separately, and that the unit-cost comparison should be recognised as the weakest of the four rather than presented as the basis for the decision. A firm that argues its actual case can be challenged on it and can learn from the outcome. A firm that argues a cost case while acting on a risk judgment has produced a document that will teach it nothing.

The practical programme is short. Separate the four arguments and name the one carrying the decision. Model the tariff as a range that includes the differential disappearing, because it has now done so once. Cost the transfer rather than the steady state, and get the ramp curve from the firm's own history. Walk the second-tier supplier base before selecting the site. Build labour cost from internal multi-site evidence including productivity and turnover rather than from comparisons published by firms that sell relocation. And stage the commitment where the volume outlook is uncertain, holding the option to stop, because the two variables that most often break these programmes are precisely the two that only a first tranche can reveal. None of this is difficult and none of it is expensive relative to the capital at stake. It is simply a different analysis from the one that currently reaches the committee, and it is the one the decision actually requires.

16 Methodology, caveats, and sources

Methodology

- This article draws on official trade and investment statistics, court records, securities filings, peer-reviewed and central-bank research, and contemporaneous reporting, current to mid-2026. Supply Chain Research is independent and accepts no payment from the site-selection advisers, shelter-service providers, or advocacy organizations discussed.
- Cost and wage figures published by parties that sell relocation services are identified as interested sources throughout. Where no disinterested series exists, that absence is stated rather than filled with a promotional figure.

Caveats

- No disinterested published series comparing fully-loaded, productivity-adjusted manufacturing labour cost across the principal nearshoring destinations was located. The comparisons cited come from firms that provide relocation services and are flagged accordingly.
- Reshoring job totals come from an advocacy organization that compiles press announcements and sells related advisory services. They are announcement counts rather than employment measurements, and the organization's own material notes the gap between announced and filled positions.
- Transshipment estimates from different research groups measure different quantities using different methods and are not directly comparable. They are presented together to show a convergent direction, not a single figure.
- The tariff position described was current at the time of writing and remains subject to further litigation and legislative action. Readers should verify the applicable authority and rates before relying on any figure here.
- Figures 1, 6, and 9 are conceptual illustrations of structure rather than measured data. Figure 8 is an indexed presentation of an interested source's published figures.
- Company examples are drawn from public disclosures to illustrate that capacity decisions are reversed at material cost. They are not a sample and no rate of failure is claimed or implied.

Sources

[1] Supreme Court of the United States. Opinion in the consolidated tariff cases, No. 24-1287 (February 2026).

[2] Congressional Research Service. Legal sidebar on the tariff decision and successor authorities.

[3] Alfaro and Chor, CEPR. An update on the great reallocation in US supply chain trade.

[4] Board of Governors of the Federal Reserve System. Vietnam's export boom to the US and the role of Chinese firms, FEDS Notes.

[5] Freund, University of California San Diego. Tracking products to identify tariff evasion through transshipment.

[6] Mexico News Daily. Preliminary 2024 foreign direct investment figures and composition, citing Secretaria de Economia and Banxico.

[7] The Timken Company. Annual report on Form 10-K, fiscal year 2025 (plant closure disclosure).

[8] E2. Tracking of announced clean-economy project cancellations and scale-backs (interested source: environmental advocacy organization).

Additional context drawn from Harvard Business School research on trade rerouting, from United States Census Bureau manufacturing construction series, from published shelter-service provider wage comparisons which are identified as interested sources, and from reshoring advocacy publications which are identified as interested sources. This article is analysis, not investment, tax, customs, or legal advice, and its conclusions should be validated against your own circumstances and obligations before any decision.