

RESEARCH REPORT

Why American productivity gains differ across places

State productivity and county income accounts measure different regional outcomes. Official evidence and original accounts separate industrial composition, efficiency, capital, commuting, and resident purchasing power.

Information cutoff: 2026-10-27

<https://aesenate.org/research/why-american-productivity-gains-differ-across-places/>

Abstract

American regions differ in their industrial mix, the capital available to workers, the organization of firms, and the access businesses have to customers, suppliers, and specialized labor. Those differences can generate divergent output per hour without implying that workers in a slower-growing place exert less effort. They can also produce high workplace output alongside modest resident incomes when commuters, ownership, and local living costs are taken into account.

Official state productivity measures for 2025 provide a concrete starting point: output increased everywhere, yet productivity declined in several states because labor hours increased faster. County production and income accounts supply additional geographic detail but measure different populations and concepts. Research on manufacturing agglomeration further challenges the assumption that productive places are simply collections of firms that were already more efficient before locating there.

This report separates those mechanisms through observed state comparisons and original production, commuting, and income accounts. It assesses how infrastructure, housing, skills, competition, and business entry can affect productive capacity, and identifies the evidence required to distinguish gains created by a place from gains generated by changing its occupants. The analysis supports place-specific diagnosis rather than a league table that treats a single annual growth rate as a comprehensive measure of regional economic performance.

The output-per-hour question

Labor productivity is a relationship between real output and hours worked. It is not a direct measure of individual diligence, a wage rate, or a complete indicator of regional welfare. More capital, better processes, a change in sector mix, or the exit of low-output businesses can raise the ratio. A temporary decline in hours can also raise it during adjustment. Interpreting the movement requires evidence about both numerator and denominator and the activity that connects them.

Regional analysis adds a boundary problem. Output is generally assigned to the location of production, while people can live elsewhere. Profits may accrue to owners outside the region, and residents may receive retirement income or transfers unrelated to current local production. A region's GDP divided by population consequently combines a workplace numerator with a resident denominator. It can be useful for a specified comparison, but it should not be labeled average worker productivity or average household income.

The geographic unit changes the result. A county containing a business center can have substantial output and relatively few residents, while surrounding counties supply commuters. A state combines those areas but still includes industry and population differences. A metropolitan economy or commuting zone can better represent some labor-market relationships, yet its boundaries may not match available statistical series or policy authority. Research should choose the boundary for the question and then test its sensitivity.

A final distinction concerns growth and level. A place with low output per hour can grow rapidly without overtaking a high-productivity place. A mature economy can retain high productivity while recording little annual increase. Ranking growth alone changes the question from productive capacity to recent change. Both are relevant, but they should not be mixed in an argument about which location provides the strongest economic opportunities.

The state evidence rejects a single growth narrative

BLS's [May 28, 2026 state release](#) reports that private nonfarm output rose in all states in 2025, while labor productivity declined in eight. Alaska's output rose 3.8% and hours 6.2%, yielding a reported productivity decline of 2.3%. Arizona's output rose 2.5% while hours fell 1.8%, yielding productivity growth of 4.4%. These contrasting combinations show why output growth and efficiency growth cannot be used interchangeably.

Using the displayed rounded rates, Alaska's ratio is 1.038 divided by 1.062, less one, or approximately -2.26%. Arizona's is 1.025 divided by 0.982, less one, or approximately 4.38%. These checks reproduce the direction and approximate size of the published results; exact reconciliation requires unrounded indexes. They establish accounting consistency, not the cause of either state's change. Explaining that cause requires industry, capital, utilization, and workforce evidence.

The BLS technical note identifies a private nonfarm production boundary and hours worked by the covered workforce. It also explains that labor productivity reflects several joint influences and is not an estimate of labor's separate contribution. That qualification is particularly important when local growth involves the opening

of new facilities. Hiring and training can raise hours before a plant reaches mature output, temporarily weakening the ratio even when future capacity is improving.

An annual decline can thus occur during expansion, while an annual increase can occur during contraction in selected activities. A regional assessment should show output and hours alongside productivity, and place the observation in a multiyear context. It should investigate revisions and changes in coverage before treating a difference of a few tenths of a percentage point as economically decisive. The object is a production account that explains the result, not a ranking that supplies a verdict in advance.

The weight of a place also matters for national interpretation. A rapid gain in a small state can contribute less to the national outcome than a modest gain in a large one. BLS separately reports state contributions using output shares and identifies the limits of that decomposition. A contribution measures influence on an aggregate; a state growth rate measures its own change. Neither is a complete measure of the productivity level or the opportunities available to residents.

County GDP provides detail without supplying a productivity estimate

BEA's [February 5, 2026 county release for 2024](#) reports real GDP growth in 2,273 counties and declines in 809. County growth rates vary much more widely than state aggregates. The release also records changes in geographic presentation, including Connecticut planning regions and the discontinuation of several metropolitan-area series. Those details matter for continuity: an analyst should not silently join different boundaries or assume a discontinued aggregate remains a current published observation.

County GDP is valuable for locating production changes and identifying industries that warrant closer examination. It does not provide the corresponding hours needed for a labor-productivity calculation merely because its geography is detailed. Dividing it by employment from a differently defined source can introduce discrepancies in coverage, jobs versus people, self-employment, and commuting. A constructed local measure requires a documented bridge between those concepts.

Small denominators also amplify volatility. A new facility, an agricultural fluctuation, or the return of an idled operation can produce a large percentage change in a small county. That can be economically important locally without materially changing national output. Absolute changes, industry contributions, and several years of evidence help distinguish a substantial expansion from a dramatic percentage on a narrow base. Large counties can make a bigger aggregate contribution while growing more slowly.

The county release combines production and personal income for good reason, but their coexistence is not an identity. Personal income includes components with different geographic and economic relationships to current production. Comparing the two can reveal useful questions about commuting, ownership, retirement, or transfers. It does not justify subtracting one from the other and calling the residual a measure of economic leakage without reconciling their definitions.

A one-year ranking can reverse the longer-run interpretation

Washington illustrates the importance of retaining the period. In BLS's May 2026 release, its private nonfarm productivity grew 0.8% in 2025, below the national rate for that year. Over 2019-2025, however, its annualized growth was 3.5%, among the strongest state results. The two observations are compatible: a relatively weak final year can follow a strong cumulative expansion. They should not be combined into a claim that the state's productive capacity had become weak or that its longer-run advantage continued at an unchanged annual pace. [BLS, annual and multiyear state productivity results](#).

A 3.5% annualized increase over six years corresponds to a cumulative increase of approximately 22.9%, calculated as 1.035 raised to the sixth power, less one. That conversion uses the rounded published rate, so it approximates the change in the underlying indexes. It is not a comparison of Washington's productivity level with the national level. A rapidly improving place can still have a lower level than another, and a highly productive place can grow slowly during a particular year.

California illustrates a different distinction: influence on the aggregate. BLS reports that its 4.2% productivity growth in 2025 contributed nearly one third of the national 1.8% increase, reflecting its substantial output weight. A state with a larger growth rate but a much smaller economy need not make a larger national contribution. The release's contribution measure captures within-state changes and excludes the effects of shifting output and labor shares across states. It is therefore narrower than a complete regional reallocation account.

Those observations direct research toward industrial and firm detail. A state can gain from more efficient continuing businesses, a larger share of high-output activities, or the contraction of lower-output activities. The aggregate series does not identify which process dominated. Production growth in a capital-intensive industry can affect state output substantially with a limited direct hours increase. Expansion of a labor-intensive service can add more employment while producing a smaller rise in the aggregate ratio. Neither pattern is a sufficient verdict on household opportunity.

A descriptive industry decomposition would identify where the movement appears before attributing it to a local policy or a characteristic of workers. It should retain comparable prices, labor coverage, and sector definitions and recognize that chain-linked output components are not simply additive. A within-industry result can still contain changes in product mix, establishment entry and exit, or outsourcing. More detail narrows the mechanism; it does not turn an accounting residual into a causal estimate.

The practical implication is to compare annual movement, cumulative change, output contribution, and the distribution of employment together. Washington's slower final year, California's large national contribution, and the Alaska-Arizona output-hours contrast answer different questions. A mature regional assessment should explain those differences before selecting a policy response. It should also retain genuine uncertainty where industry or firm evidence is insufficient to identify the cause.

Capital intensity and the meaning of a productive place

Labor productivity can rise when workers use more or better capital. A region with energy installations, advanced manufacturing equipment, or digital infrastructure can generate high output per hour with relatively few workers. That does not mean all of the output is labor income. Returns to capital, depreciation, taxes, and payments to other factors help determine what residents receive. A regional development strategy should distinguish productive capacity from its distribution among people.

Capital quality matters as well as quantity. Installed machinery that is poorly maintained or lacks a reliable connection can provide fewer productive services than its acquisition cost suggests. An old asset can remain useful with appropriate maintenance, while a new asset may take time to reach expected performance. A useful account follows capital services and utilization rather than infer production capability directly from announced investment or the book value of assets.

Investment can also change the local demand for skills. Automation may reduce some routine tasks while raising requirements for maintenance, engineering, supervision, and complementary services. Aggregate hours can remain stable while occupations and earnings diverge. Training policy should therefore be connected to actual job tasks and employer demand. Counting course participants is an implementation measure; it does not establish a productivity gain or successful movement into higher-paid work.

The capital channel creates an identification challenge. Firms invest where they expect demand or local advantages to improve. A correlation between investment and later productivity is consistent with investment causing gains, firms anticipating gains, or both. Evaluating an incentive requires evidence on the unassisted decision and alternative locations. A productive facility in a supported place does not establish that the support created the facility or that its national location was additional.

Agglomeration and sorting are competing mechanisms

Proximity can improve productivity through specialized suppliers, a deeper labor pool, knowledge exchange, and shared infrastructure. Productive firms and workers may also select places offering better opportunities. Competition can force less efficient businesses to exit. These mechanisms can generate the same positive relationship between density and average productivity while implying different responses to policy. An urban productivity premium cannot, by itself, identify which mechanism dominates.

Andrea Stella's [2025 Federal Reserve working paper on manufacturing](#) estimates heterogeneous agglomeration benefits using plant data and a production-function framework. It finds larger benefits for lower underlying-productivity establishments, challenging a simple account in which the most efficient plants gain most from density. The result is specific to its manufacturing sample and model. It is not evidence that every cluster policy works or that worker sorting is unimportant in all activities.

The paper's identification depends on assumptions about input choices, productivity dynamics, and establishment location. That is substantively different from comparing output per resident between dense and sparse counties. Its value for policy is to make heterogeneity an empirical question. If benefits differ across

firms, attracting one highly visible facility may not maximize local learning or supplier gains. Conversely, supporting all firms in a cluster equally may ignore where the external benefit is strongest.

Agglomeration also has costs. Congestion, expensive land, commuting time, and pressure on public services can offset some production advantages. A firm can gain from proximity while workers face higher housing costs. A national strategy may benefit from enabling more people to access productive places, but that requires housing and transport capacity. Simply encouraging concentration without those complements can increase measured workplace productivity while limiting the spread of household gains.

Commuting separates production from resident prosperity

Consider a hypothetical central county with output of \$12 billion and 100,000 residents. It has 60,000 workplace jobs, of which 30,000 are held by commuters living elsewhere. Thirty thousand residents hold local jobs, another 10,000 work outside the county, and the remaining residents include children, retirees, unemployed people, and others outside employment. Output per resident is \$120,000. Output per workplace job is \$200,000. Neither equals income per resident or output per labor hour.

The example deliberately uses jobs without multiple-job holding to make the boundary transparent. Introducing multiple jobs would require another reconciliation. The large output-per-resident figure partly reflects production undertaken by nonresidents. Their earnings can support consumption and taxes in surrounding counties. Calling that entire difference an unexplained productivity advantage would combine a genuine business center with a geographic accounting effect.

Now suppose a transit improvement permits another 10,000 commuters to reach the center, with additional output of \$1 billion and no immediate change in resident population. Output per resident rises to \$130,000 while output per workplace job falls to roughly \$185,700. One ratio improves and another declines, despite an increase in production and access. Hours, wages, commuting costs, and the counterfactual work opportunities are needed to evaluate the benefit. The example shows why no single denominator can answer every regional question.

Remote work introduces related challenges. The location of an establishment, the residence of an employee, and the place where tasks are performed may differ. Statistical allocation follows defined methods rather than an intuitive map of every laptop. Changes in working arrangements can therefore affect the interpretation of regional series even when national production is little changed. Researchers should consult the relevant allocation rules before assigning a measured change to a physical relocation of productive capability.

Housing costs and the purchasing power of earnings

A nominal wage premium can compensate partly for a higher cost of living. BEA's [regional price parity release covering 2024](#) places California's overall price level at 110.7 relative to the national level of 100, compared with 86.9 in Arkansas. These are regional basket comparisons, not individual household budgets. They demonstrate why equal dollar incomes do not necessarily purchase equal quantities of local goods and services.

For illustration, a hypothetical \$80,000 income in a place with a price parity of 110 has purchasing power of about \$72,727 in national-price units. A \$70,000 income in a place with parity 90 corresponds to about \$77,778. The nominal ranking reverses after this simple adjustment. The incomes and price parities in that example are invented and are not claims about either named state. Household composition, housing tenure, taxes, and particular spending needs can produce different individual comparisons.

Regional price adjustment is not a universal deflator for business output. A household consumption basket differs from the prices of a factory's output, machinery, or intermediate inputs. Dividing GDP by a consumer price parity and calling the result technological efficiency would conflate production and consumption prices. The adjustment should match the object: household purchasing power, business costs, or real production. The relevant indexes are related but not interchangeable.

Housing also affects access to productive employment. A skilled worker may be willing to accept an offer in a high-productivity place but unable to find an affordable dwelling within reasonable travel time. Employers can face recruitment limits even when nominal wages are high. Increasing housing supply or transport access can then support both mobility and production. The policy's effect should be measured through feasible access and subsequent outcomes, rather than assumed from the wage premium alone.

Innovation diffusion can matter more than a new frontier

Regional gains do not always require inventing a new technology locally. Businesses can improve by adopting established practices in inventory control, maintenance, software, energy use, and customer management. The constraint may be managerial time, finance, data quality, or access to implementation expertise. A region with many smaller firms can therefore have substantial opportunities for diffusion even if it does not host the largest research centers.

The output effect depends on complementary changes. Purchasing software without redesigning the process can increase costs with little improvement in accepted output. Installing machinery without training or reliable inputs can leave capacity idle. A diffusion program should identify the operational bottleneck and follow actual use, quality, downtime, and unit cost. Participation and equipment purchases are necessary records but do not establish that productive performance changed.

There is also a selection problem. Businesses that volunteer for assistance may already be better managed or planning improvements. Comparing their later performance with every nonparticipant can overstate the program effect. Useful evidence can come from phased access, eligibility thresholds, or a carefully designed

comparison among similar applicants. Whatever the design, the outcome should be tied to the particular capability the assistance was intended to improve.

The regional spillover requires further evidence. An assisted firm can share trained workers, supply knowledge, or new demand with nearby businesses. It can also draw scarce staff away from them or capture their customers. Recipient gains are not automatically a net regional gain, and regional gains are not automatically a national gain if activity moved from elsewhere. Evaluation should retain these nested boundaries rather than stop at the most favorable one.

Entry, exit, and the distribution of firm performance

An average can rise because continuing firms become more productive, productive entrants expand, less productive firms exit, or workers shift between them. These processes differ in their implications for households and policy. Improving continuing businesses may preserve jobs, while exit can impose substantial transitional costs even when resources eventually move to more productive uses. A complete regional account follows the distribution and reallocation of activity rather than describe every increase in the mean as uniform progress.

A small hypothetical example makes this clear. Two firms initially employ 50 workers each and produce output per worker of 80 and 120, respectively. Average output per worker is 100 and total output is 10,000. If the lower-productivity firm closes and no worker is reemployed, the average among remaining workers rises to 120 while output falls to 6,000. The measured productivity gain is compatible with a severe loss of employment and income. This is an accounting example, not an estimate of typical closures.

If all displaced workers instead join the other firm at its existing productivity, output rises to 12,000 with average productivity still 120. That alternative requires demand, capital, skills, and organizational capacity to accommodate them. It cannot be assumed merely because the surviving firm was initially more productive. Reallocation is a process with frictions, and the speed and quality of that process determine the economic cost of transition.

Policy should therefore distinguish support for people from support for an unchanged business structure. Retraining, job matching, mobility assistance, and public services can improve adjustment without guaranteeing every incumbent firm's survival. Conversely, temporary support may preserve a productive network through a reversible shock. The correct instrument depends on whether the problem is a temporary disruption or a persistent mismatch, which requires firm and market evidence beyond a regional average.

Hours, employment, and access to productive work

A place can raise total production through additional employment, longer hours per worker, or greater output per hour. Those components should be kept separate because their welfare implications differ. Extending already long working hours is not equivalent to making a process more efficient. Enabling an unemployed resident to obtain regular work can improve welfare even if the new worker initially produces less per hour than incumbent workers and lowers the workplace average.

An exact identity helps organize the comparison. Let population be N , employed residents E , hours per employed resident h , and output per corresponding hour p . Where the production and residence boundaries are genuinely aligned, output per resident equals E divided by N , multiplied by h and p . A change in employment access can therefore raise output per resident without a change in hourly productivity. The alignment condition is essential: applying the identity with local workplace GDP and resident employment in a commuting center would be inconsistent.

In a constructed closed region, suppose 60 of every 100 residents work 1,500 hours annually and produce 100 real units per hour. Output per resident is 90,000 units. If employment rises to 63 with hours and productivity unchanged, it becomes 94,500, a 5% gain. If instead productivity rises to 105 with employment unchanged, the same result occurs. The two gains are numerically identical but have different sources. They also need not be equally distributed among households.

Now suppose additional workers enter at 80 units per hour while the original 60 retain productivity of 100. The region adds 360,000 units of output from three workers at 1,500 hours each. Output per resident rises to 93,600, but average hourly productivity falls to approximately 99.05. A declining aggregate productivity ratio can accompany improved inclusion and higher production. The example does not establish that lower productivity is desirable; it establishes that employment access and efficiency must be evaluated together.

This is relevant to policies addressing childcare, disability accommodation, transport, or scheduling barriers. Their purpose may be to expand feasible participation rather than raise the output of every incumbent hour immediately. A sound evaluation measures sustained employment, hours, earnings, and the cost of participation, alongside business performance. It should not dismiss an inclusion gain because a regional productivity rank declines, or declare success solely because participants are administratively recorded as employed.

Business costs do not equal household living costs

A high-cost place can offer productive advantages sufficient to justify higher wages, rent, or supplier charges. The relevant business comparison is the cost of delivering accepted output to customers. A lower hourly wage can be offset by additional training, downtime, logistics, or rework. A higher rent can be justified by access to customers or specialized staff. Comparing only one input price can therefore misstate the economics of a location.

For a hypothetical operation, assume labor compensation of \$40 per hour and accepted output of two units per hour. Direct labor cost is \$20 per accepted unit. At another location, compensation is \$32 but accepted output is 1.4 units, giving about \$22.86 per unit. The lower wage does not produce the lower direct labor cost. The comparison omits capital, materials, quality differences, and transport deliberately, so it cannot rank the total cost of the locations. Its purpose is to align the cost denominator with usable production.

Reliability can change that denominator further. A machine's nominal throughput is not accepted output if interruptions, defects, or missing components prevent delivery. A location incentive that lowers the acquisition cost of equipment can have little operational value if a power or logistics constraint keeps it idle. Conversely, a

shared infrastructure improvement can raise productive services from existing assets without a large new investment announcement. Research should examine the constraint that actually limits delivery.

The connection to household welfare then requires compensation and prices, rather than an assumption that business efficiency passes wholly into wages. Some gains accrue to customers through lower prices, some to owners, and some to workers. Local competition, bargaining, ownership, and mobility influence that distribution. A regional production improvement is valuable evidence, but identifying who benefits requires the additional income and consumption account developed above.

Evaluating a regional investment intervention

A new plant can raise a host area's output, wages, and land values while shifting activity from another location. A local authority may appropriately care about the host gain, but a national appraisal must consider the displaced alternative and the fiscal package. Announced jobs do not reveal either effect. The project should be followed through finance, construction, operation, hiring, and supplier development, with a defined baseline for each claimed outcome.

Location incentives are particularly vulnerable to selection. Firms choose among places, and governments offer support partly in response to expected investment. A recipient's subsequent success cannot be attributed to the incentive simply because the project used it. Internal decision records, competing sites, timing, and eligibility rules can help identify the relevant margin. Where confidential information limits disclosure, public analysis should state the resulting uncertainty rather than replace the counterfactual with an assumed cancellation.

The outcome should include local incumbents. A large facility can raise demand for suppliers while increasing wages, rents, and congestion. Some businesses gain customers; others lose staff or become unable to afford premises. Those effects can coexist. A narrow recipient account omits them, while an aggregate regional series may not identify their distribution. Linked business and household evidence is preferable where available and legally accessible.

An evaluation horizon should also reflect the mechanism. Recruitment during construction is different from sustained operating employment. Supplier learning can take longer than commissioning. A short-run tax-revenue increase may be offset by infrastructure costs or incentives paid later. A multiyear account should discount dated financial flows consistently and avoid counting the same wage or expenditure as an independent benefit several times. The appraisal must distinguish financial transfers, resource use, and welfare.

A practical diagnostic sequence

For a region with weak productivity growth, begin by reconciling output and hours over several years. Identify whether the result arises from slower output, faster labor input, or both. Then examine industries, major facilities, and revisions. This first stage can reveal a temporary ramp-up or commodity exposure that would otherwise be misread as a broad deterioration in business efficiency.

Next distinguish composition from within-industry performance and inspect the distribution of firms. Are successful businesses expanding? Are entrants improving the average? Are workers moving into higher-output activities, or has a low-output sector simply contracted without replacement? Those questions lead to different policy priorities. A uniform prescription for more research spending, lower taxes, or more training is unlikely to address every mechanism.

Then connect workplace outcomes to residents through employment access, commuting, compensation, ownership, transfers, and local prices. A productive business center can coexist with excluded neighborhoods or high housing costs. A less productive retirement destination can have substantial personal income unrelated to current production. The connection to household opportunity should be investigated directly, rather than inferred from GDP per resident or a nominal wage ranking.

Finally identify the constraint that a feasible public action can change. Infrastructure, housing permissions, training, information, and competition operate through different margins and timescales. Establish the baseline before implementation, define outcomes that can falsify the theory, and retain a comparison where possible. Evaluation should recognize improvements that do not change a ranking and reject apparent gains created solely by denominator or boundary changes.

Findings and analytical limits

American productivity differences reflect several overlapping processes: the mix of industries, capital services, firm efficiency, worker and firm selection, agglomeration, and the geography of work and residence. Official state observations demonstrate that strong output growth can coincide with declining productivity when hours grow faster. They do not identify the cause of every state's result. County and income accounts add evidence without eliminating the need for compatible labor measures.

All original regional accounts here are hypothetical. The published observations retain their release dates and statistical boundaries. The manufacturing research supplies evidence on a specific estimated mechanism, not a universal estimate of urban benefits. This report does not rank states by an unobserved underlying efficiency measure or estimate the return to a named incentive. Such claims require more granular data and an identified comparison.

The most useful policy objective is to improve the ability of firms and people to produce and receive valuable services at sustainable cost. Productivity is central to that objective, but its interpretation requires a complete account of production, access, and distribution. Keeping those relationships explicit makes regional research useful for investment and public decisions without confusing a place's statistical position with the opportunities available to its residents.

Related research

- [Explaining productivity differences across the Atlantic](#) examines cross-economy industrial composition and productivity.

- Housing supply constraints and the transmission of U.S. monetary policy connects access to employment with the physical housing stock.

Source addresses

The links below retain the source addresses cited in the analysis. References should be read within their stated periods and scope.

<https://www.bls.gov/news.release/prin4.htm>

<https://www.bea.gov/news/2026/gross-domestic-product-county-and-personal-income-county-2024>

<https://www.federalreserve.gov/econres/feds/files/2025031pap.pdf>

<https://www.bea.gov/data/prices-inflation/regional-price-parities-state-and-metro-area>