

The Structural Geography of U.S. Defense Industrial Production: Vulnerability, Resilience, and Surge Capacity Under Autarky

Tandem Young*

May 2026 (working draft)

Abstract

A weapons platform has three distinct geographies—where its contracts are *booked*, where it is *produced*, and where latent capacity could be *mobilized*—yet industrial-base assessment conventionally weights counties by prime-contract share, conflating them. For the F-35 the three disagree by an order of magnitude. That contract booking misstates production geography is recognized (Park et al., 2026; Malecki, 1984); that re-weighting *reverses* the cluster’s apparent fragility is the finding. Tarrant TX, home to Lockheed’s final-assembly plant, records the largest single share of national defense obligations but under one percent of national defense employment—where the contract is booked, not where the aircraft is built. Re-weighting county exposure by defense *production* (employment) and mapping input-output relatedness onto a county network, the F-35 cluster spans 56 counties in 25 states and behaves as roughly sixteen effective counties, not the two or three the obligation weight returns. Its *upstream parts base* is structurally resilient—across 100 Louvain seeds no single county’s removal fragments the supply network (median cohesion near 85%)—reversing the base-wide single point of failure the booking weight implies; final assembly itself remains a distinct chokepoint. Across the ten largest defense clusters six are genuinely single-anchor fragile, the production-dispersed F-35 cluster a notable exception, and about 96% of its import content has a domestic substitute. We build four instruments from public data—a Leontief closed-loop vulnerability index, empirically-derived co-location spatial weights, a cascade simulation, and a Hidalgo–Hausmann surge-capacity index—reporting each under both weights. Reserve capacity is abundant, so the breadth of the base is not the binding constraint; whether conversion *time* at integration sites is, we flag as a hypothesis. Two implications follow: report industrial-base concentration by production, not booking; and target resilience investment by production-weighted fragility, which points to different counties. The analysis is descriptive economic geography, not a causal-effects study.

JEL codes: H56, R12, L64, F52.

Keywords: defense industrial base, economic geography, closed-loop / autarky vulnerability, surge capacity, agglomeration, spatial weights, cascade simulation.

*University of Arkansas. The closed-loop / autarky framing of this analysis emerged from discussions with Daniel Egel at RAND in May 2026. The author thanks Egel and colleagues for substantive engagement that reshaped this work. All findings and remaining errors are the author’s responsibility.

1 Introduction

One U.S. county records the largest single share of national defense procurement obligations—about six percent, more than any other county—and on a standard contract-weighted industrial-base metric the F-35 supply base looks as though it hangs from that one node. That county is Tarrant TX, home to Lockheed Aeronautics’ F-35 final-assembly plant. It does not. Tarrant holds under one percent of national defense *employment*: the figure measures where the F-35’s contract is *booked*, not where the aircraft is *built*. This paper’s organizing claim is that a weapons platform has three distinct geographies—where its contracts are recorded, where it is physically produced, and where latent capacity could be mobilized to substitute—and that the industrial-base debate routinely conflates them, with the contract-booking geography standing in for all three. For the F-35 these geographies disagree by an order of magnitude—and that re-weighting to production *reverses* the base’s apparent fragility, not merely its accounting, is the finding.

Weighted by production rather than by contract booking, the picture inverts. The cluster Tarrant anchors spans 56 counties across 25 states and behaves not as two or three effective counties—the figure the obligation weight returns—but as roughly sixteen, a genuinely dispersed national aerospace base. That base is structurally resilient: removing any single county and re-running community detection leaves the parts network intact in most trials, so the *base-wide* “single point of failure” that a booking-weighted analysis reports—the claim that the whole supply base hangs on one county—is an artifact of routing the program’s entire recorded value through one assembly site. (Final assembly itself remains a genuine chokepoint, distinct from the dispersed parts base; Section 10.) And its import vulnerability, propagated through the input-output structure and decomposed against domestic substitutes, is mostly slack: roughly 96% of the cluster’s import content has a domestic counterpart. Figure 1 shows the cluster and the gap between its booking and production geographies. This paper quantifies that gap at county resolution, demonstrates that it is structural rather than statistical, traces its consequences for the resilience of the production base, and generalizes the method that produces it.

The question this measures is newly urgent. The debate over U.S. defense industrial mobilization has intensified: the manufacturing-at-scale agenda within the Department of Defense, the semiconductor reshoring of the CHIPS and Science Act, the renewed invocation of the Defense Production Act, the decoupling from Chinese supply, and the materiel demands of the war in Ukraine have together made industrial-base resilience a first-order strategic concern (U.S. Department of Defense, 2024a; Congressional Research Service, 2024c; Kane et al., 2024). The recurring question beneath all of these is a closed-loop one: cut off from imports, what could the United States still produce, and where.

What the debate has lacked is structural measurement. The conventional instruments—Herfindahl indices, four-firm concentration ratios, supplier counts, raw geographic dispersion—describe how widely defense work is spread but not whether the production network through which a shock would propagate is resilient or brittle. They also share an input that quietly determines the answer: they weight counties by their share of prime-contract obligations. Because a modern platform’s prime is a single final-assembly site, that weighting concentrates the entire program’s recorded value there and makes a genuinely dispersed production base look like a single load-bearing county—the precise inversion this paper documents. The 2024 National Defense Industrial Strategy names supply-chain resilience as a priority but does not quantify which counties or network anchors are structurally load-bearing, nor distinguish where contracts are booked from where production happens (U.S. Department of Defense, 2024a,b); related work has approached the problem from procurement-policy and threat-assessment angles (Kane et al., 2024; Allen and Berenson, 2024). That contract booking diverges from where defense activity actually lands is itself established—Malecki (1984) mapped the contract-versus-subcontract ge-

The F-35 Production Base Is Dispersed, Not One County
Red = Tarrant TX (contract-booking anchor) Orange = F-35 production cluster (56 counties, 25 states) Gray = rest of CONUS

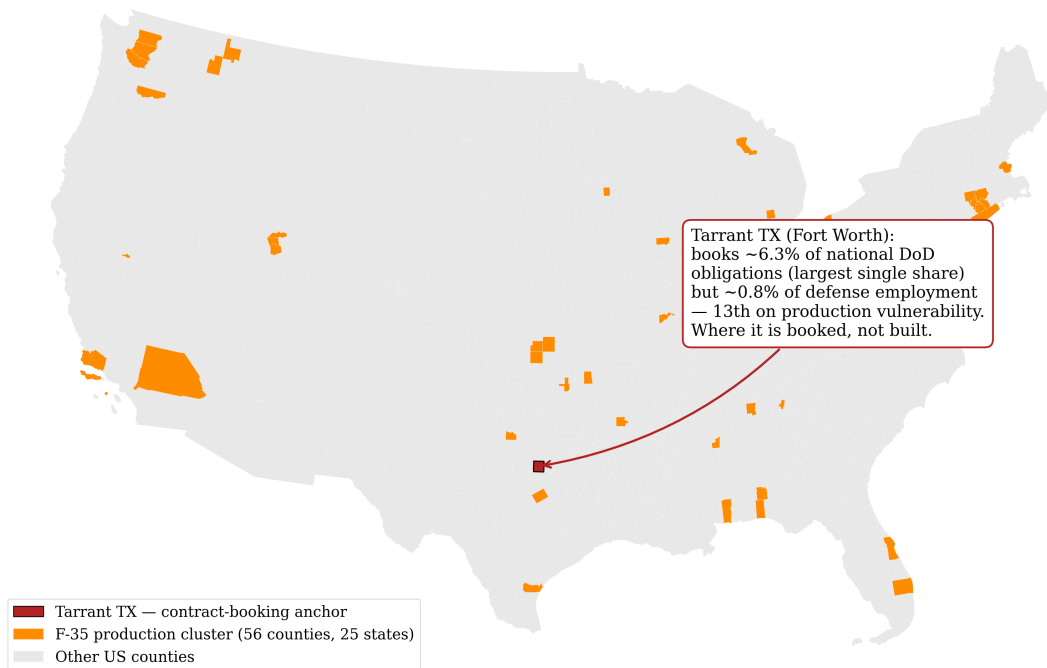


Figure 1: The F-35 production cluster and the gap between contract booking and production. By prime-obligation booking the base concentrates on Tarrant TX (which records the largest single share of national DoD obligations); by defense employment it is dispersed across 56 counties in 25 states, behaving as roughly sixteen effective counties rather than the two-to-three the booking weight implies. The two geographies of the same platform disagree by an order of magnitude.

ography at the state level, Gerritse and Rodríguez-Pose (2018) separated recipient location from place of performance across federal contracting, and Park et al. (2026) contemporaneously quantify how subcontracting reallocates defense dollars away from the prime's location. What does not exist is a county-level, network-resilience *measurement* of the defense industrial base built on that distinction—an index that propagates import dependence through the production structure, a cascade that tests for load-bearing anchors, and a surge map of latent capacity. That is the gap this paper fills.

This paper makes three contributions. *First*, it turns a known distinction into a reversal: that contract booking misstates production geography is established (Park et al., 2026; Malecki, 1984), but re-weighting the F-35's base to production does not merely move the numbers—it overturns the verdict, recasting the cluster that a contract-weighted metric flags as the defense base's clearest single point of failure as, measured by production, among its most resilient. *Second*, it diagnoses why the conventional picture misleads: the standard concentration and vulnerability surfaces—Herfindahl indices, dispersion counts, supplier tallies—share one hidden input, the prime-obligation weight, so that diagnostics that look independent all reduce to where the contract is booked. *Third*, it supplies a reusable, public-data method for measuring defense industrial structure at county resolution and stress-testing its resilience, demonstrated across the ten largest defense clusters and generalizable to any platform. The contribution is descriptive and structural, not causal: what it adds is measurement, not a spillover estimate.

These contributions rest on four measurement instruments, each adapted from an established literature and built entirely from public data, together with the exposure-weight distinction that organizes them—the descriptive-structural counterpart to the shift-share literature's lesson that the choice of exposure weights silently governs results that look independent of it (Goldsmith-Pinkham et al., 2020; Borusyak et al., 2022). The first is a closed-loop vulnerability index (V_2): for each county, the share of defense-relevant production exposed to import disruption, propagated through the Leontief inverse and weighted by defense *production*—county defense-employment share—rather than by prime-contract booking, so that exposure is counted where the work is done; it is decomposed into the share carried by inputs with domestic substitutes and an irreducible no-substitute floor. The obligation weight is retained throughout as the comparison that makes the production-versus-booking gap legible. The second is a set of empirically-derived spatial weights built from the co-location structure of defense production rather than from geographic contiguity— W_2 V_3 , from BEA input-output relatedness weighted by county employment and defense-exposure share, together with a complementary W_3 built from the largest observed FY2023 USAspending sub-awards (those of at least \$1 million)—accompanied by the demonstration that the implied spillover is governed by the choice of propagation network rather than by the local elasticity, and is therefore not point-identified. The third is a cascade simulation that converts a static concentration measure into a dynamic statement about structural resilience. The fourth is a Hidalgo–Hausmann surge-capacity index adapted to sub-national defense industrial geography. Vulnerability and surge are the two halves of one device: the first identifies what breaks under autarky, the second what could be mobilized to grow into the gap.

This is descriptive industrial geography in the new-economic-geography tradition (Krugman, 1991b; Fujita et al., 1999), drawing on the regional-branching and defense-economy literatures (Markusen et al., 1991), and it is deliberately not a causal-effects paper. The object is not a spillover coefficient but a structural characterization of an industrial geography: what is concentrated, how load-bearing the concentration is, and what could be mobilized against it. The historical register is the Ford River Rouge conversion of 1942 (Herman, 2012); the analytical engine is the closed-loop vulnerability–surge apparatus, which supplies the contemporary structural measurement the manufacturing-at-scale debate has invoked by analogy but not quantified.

The findings compose a single argument. Measured by production, defense industrial vulnerability is led not by Tarrant but by the large aerospace-and-electronics metros—Los Angeles, Santa Clara, King WA / Seattle—and the F-35 cluster, dispersed across sixteen effective counties, rests on a structurally resilient parts base: removing any single county and re-running community detection leaves the cluster’s cohesion at a bootstrap median near eighty-five percent across one hundred Louvain seeds, so it reconstitutes around its broad production base rather than collapsing. The apparent single-point-of-failure that an obligation-weighted cascade reports—a near-total collapse when Tarrant is removed—does not survive the switch to a production weight; it was a property of concentrating the program’s recorded value on one final-assembly county. This resilience is specific, not generic: applied across the ten largest defense clusters, the cascade shows six of them *are* single-anchor fragile, with the production-dispersed F-35 cluster a notable exception (Section 7). The commercial capacity that could backfill defense production in a surge is abundant and widely distributed—hundreds of counties host substantial defense-related commercial manufacturing, and the F-35 cluster’s own production counties are overwhelmingly self-rescuable, carrying both defense work and dual-use slack (Section 8). What our instruments measure is capacity and its spatial structure; whether the binding mobilization constraint is the *time* to convert that capacity at integration sites is the question they raise but cannot answer, and we flag it as a hypothesis rather than a result. The through-line is the measurement distinction: weight the defense base by production rather than by contract booking, and both its concentration and its fragility fall by an order of magnitude.

The paper proceeds in the order this argument requires. Section 2 sets out the new-economic-geography framework the empirical apparatus operationalizes; Section 3 describes the data and Sections 4–5 the methodology; Sections 6–8 present the vulnerability, cascade, and surge results in turn; Section 9 stress-tests every specification choice on which they rest; Section 10 develops the policy implications; and Section 11 draws the contribution together. Section 2 begins with why a defense industrial geography is a new-economic-geography object and not a problem in spatial econometrics.

2 Conceptual Framework

This section sets out the theory the empirical apparatus operationalizes. The argument is that the structural facts this paper documents—a single dominant final-assembly core, a cluster that fragments when that core is removed, a reserve of commercial counties one capability step from defense production—are not spatial-econometric artifacts to be absorbed and discarded but the substantive content of a defense industrial geography. Each construction in the paper is an operationalization of a specific new-economic-geography (NEG) mechanism: the vulnerability index of increasing-returns lock-in, the cascade simulation of agglomeration-core versus gateway-hub structure, the surge-capacity index of regional branching. The framework also fixes the closed-loop, or autarky, lens through which all three are read.

2.1 From spatial dependence to economic geography

Spatial econometrics treats geographic dependence largely as a nuisance (Anselin, 1988; LeSage and Pace, 2009): cross-unit correlation is something to be modeled and absorbed so that an underlying coefficient is identified cleanly, and the spatial-weights matrix is a contiguity or distance object—a description of who is near whom in physical space, useful as machinery but not itself the object of inquiry. New economic geography inverts this priority (Krugman, 1991b,a; Fujita et al., 1999). The questions it

asks—why industries concentrate where they do, what sustains the concentration, and what happens to it when a shock arrives—make geography the dependent variable, and the mechanisms that produce it (agglomeration externalities, increasing returns, forward and backward linkages, factor mobility) the substantive theory rather than a statistical inconvenience.

That distinction is decisive here. The headline findings developed in Sections 6–8—the concentration of national defense industrial vulnerability in a single county, the fragmentation of the F-35 cluster when that county is removed, the existence of a mobilization-reserve belt across the industrial Midwest—are not noise to be controlled away. They are the structure of US defense industrial geography, and treating them as artifacts of a misspecified error process would discard exactly what is of interest. The contribution is correspondingly not a national-scale spillover coefficient but a characterization of an industrial geography’s structural properties: what is concentrated, how load-bearing the concentration is, and what could be mobilized against it. The remainder of this section develops the NEG concepts that give those properties their meaning.

2.2 Agglomeration economies and lock-in

The classical account of why production concentrates is Marshall’s. Marshall (1890) identified three reasons industries cluster, each of which operates in the defense case. **Labor pooling:** a thick market in specialized skills— aerospace and propulsion engineers, cleared machinists—lowers matching frictions for firms and workers alike, and makes an established cluster the rational location for the next increment of investment. **Knowledge spillovers:** proximity transmits the tacit, non-codified knowledge of complex systems integration, which is precisely the knowledge that does not relocate cheaply. **Input sharing:** co-located primes sustain specialized suppliers that no single buyer could support alone. The modern empirical literature measures the strength of these forces from cross-sectional coagglomeration (Glaeser et al., 1992; Henderson et al., 1995; Ellison et al., 2010); the cascade simulation of Section 7 supplies a complementary kind of evidence, simulating their breakdown by node removal rather than inferring their presence from covariation.

Krugman (1991b) formalized how such externalities select a spatial equilibrium. In a two-region model with increasing returns at the firm level, iceberg transport costs, and mobile labor, the symmetric configuration—industry split evenly between regions—is unstable: a small initial advantage compounds, because firms locate where workers are and workers locate where firms are, and concentrating production economizes on transport. The stable outcome is core–periphery, with one region dominant (Krugman, 1991a). This is the canonical answer to a question the defense case poses sharply—why there is one fifth-generation-fighter final-assembly core and not two—and why, once that core is established, parallel capacity elsewhere faces unfavorable economics even where sunk facilities and strategic arguments for redundancy exist. Path dependence makes the equilibrium durable: once the relative-cost calculation for new entrants is anchored to the incumbent location, lock-in persists absent a large external intervention (David, 1985; Arthur, 1989). The policy reading is immediate, and Section 7 makes it concrete: a proposal to distribute defense final assembly geographically would run against the same increasing-returns force that produced single-site dominance in the first place.

2.3 Hub-and-spoke networks: agglomeration cores and gateway hubs

Fujita et al. (1999) generalize the two-region model to many regions and many industries. Three of their results organize the empirical analysis. First, when industries differ in their input-output linkages, agglomeration is industry-specific: each industry settles into its own spatial equilibrium with its own

anchor, so a defense economy is a superposition of distinct clusters rather than one undifferentiated concentration. Second, the equilibrium can take a *hub-and-spoke* form—one dominant location with feeder regions—rather than pure core–periphery. Third, the geography of intermediate-goods trade is itself shaped by the geography of industry concentration (Krugman and Venables, 1995), which is why the network along which a procurement shock propagates is an economic object to be measured, not a contiguity matrix to be assumed (Section 2.5 returns to this).

FKV also implies a distinction this paper makes operational. A location can be central to a production network in two structurally different ways. An *agglomeration core* is the anchor around which a cluster forms through cumulative causation: the other members belong to the cluster because of their input-output proximity to the core and have little proximity to one another, so removing the core leaves the cluster with no remaining basis for cohesion and it fragments. A *gateway hub* is a high-betweenness location that bridges multiple clusters in a hub-and-spoke structure without belonging to any one of them, so removing it redistributes cluster membership rather than fragmenting a cluster. The two are observationally similar in any static centrality measure—both look “central”—and are separated only by a counterfactual. Section 5 specifies the cascade experiment that supplies that counterfactual, and Section 7 shows that the distinction is not merely conceptual: the anchor of the F-35 cluster and the busiest connector of the national defense network are different counties, and their removals have qualitatively different consequences.

2.4 Regional branching, relatedness, and surge capacity

Where the preceding mechanisms explain how a concentration forms and how fragile it is, the evolutionary strand of economic geography explains what an economy can become. The regional-branching literature holds that a region diversifies into new activities along lines of *capability proximity* to what it already does: related diversification is feasible at moderate cost, unrelated diversification is not, at least not on policy-relevant horizons (Boschma and Frenken, 2006; Boschma, 2017). The empirical apparatus is the product space of Hidalgo et al. (2007) and the sub-national branching analysis of Neffke et al. (2011): industries are linked by a relatedness measure—the matrix ϕ of Section 5—built from the tendency of capabilities to co-occur, with industry presence defined by revealed comparative advantage (Balassa, 1965; Hausmann and Klinger, 2007; Hidalgo and Hausmann, 2009). Regional growth is supported by *related* variety, the co-presence of activities near one another in capability space, and not by unrelated variety (Frenken et al., 2007; Boschma and Iammarino, 2009).

Surge capacity is this logic run in reverse. The question is not whether a commercial county could build an aerospace industry from nothing—the unrelated-variety answer is that it could not on any relevant timescale—but how close a county’s existing capability base already sits to defense production. A commercial-automotive county is, in capability space, a short branch from defense-vehicle manufacture, and is a latent mobilization asset even with zero current defense obligation. Section 8 operationalizes exactly this as a county-level surge-capacity index: the related-variety pool, measured against the defense set \mathcal{D} , that the economy could draw on to expand production under stress. The vulnerability index identifies what breaks; the surge index identifies what could grow into the gap. This connects to the industrial-commons argument of Pisano and Shih (2009, 2012): when manufacturing capability disperses and thins, the related-variety pool contracts and the branching paths that mobilization would rely on close off—so the same dispersion that a naive reading would call diversification is, under autarky, a loss of surge depth.

2.5 Propagation networks and the closed-loop frame

A methodological consequence follows directly. If spillovers propagate through the mechanisms NEG describes, the spatial-weights matrix should encode that mechanism, not physical adjacency; the economic-geography literature has long argued for exactly this move: weights drawn from input-output linkages or capability proximity in place of contiguity (Combes et al., 2008; Conley, 1999; LeSage and Pace, 2009). Section 4 constructs the paper’s weights on that principle, from the input-output co-location structure (W_2) and from observed sub-awards, and propagates exposure through the Leontief inverse \mathbf{L} to form the vulnerability index V_2 . The empirical work then makes a separation visible: the local response of a defense agglomeration to procurement is invariant to the choice of weights, but the spillover geography is not—it depends entirely on which network is assumed to carry the shock. We read this as a caution, not an estimate. Because the propagating network is a co-location structure rather than observed flow, and because the implied spillover multiplier is not robust to that choice—or even to whether all defense counties are correctly included (Section 4)—the multiplier is not point-identified. What survives is the invariant local elasticity and the demonstration that any spillover claim is hostage to an untested propagation assumption.

This is what the *closed-loop*, or autarky, lens makes precise, and it is the lens current policy demands. The CHIPS and Science Act, the Inflation Reduction Act’s domestic-content provisions, the renewed use of the Defense Production Act, and the broader decoupling from Chinese supply—together with the Department of Defense’s manufacturing-at-scale agenda—all pose one question: cut off from imports, what could the United States still produce, and where. Posed structurally, that question has an axis the conventional concentration metrics miss. Defense output built on import-dependent manufacturing is far less resilient under a closed loop than output built on domestic services, so a vulnerability measure must propagate import dependence through the production structure rather than count contracts; Section 6 shows the resulting manufacturing–services gap is large. The paper’s defense-economy precedent is the spatial defense-industrial work of Markusen (1986); Markusen et al. (1991), which mapped the Cold War “gunbelt” descriptively; the present contribution is to supply the structural, mechanism-explicit measurement that the manufacturing-at-scale debate currently substitutes historical analogy for—taking Herman (2012) as the rhetorical register and the Hidalgo–Hausmann apparatus as the analytical engine.

The booking-versus-production distinction the paper builds on is not itself new, though its structural consequences are. Malecki (1984) showed four decades ago that prime contracts and subcontracts trace different maps of the defense economy, if at the state level and without an employment-based production measure. Gerritse and Rodríguez-Pose (2018) separated a contract recipient’s location from its place of performance across all federal contracting and found the two footprints diverge. Most directly, Park et al. (2026) document—contemporaneously and with establishment microdata—that defense subcontracting reallocates a substantial share of dollars away from the prime’s location, understating local fiscal multipliers by about a fifth. What distinguishes the present paper is twofold. First, its production weight is a real-activity stock—county defense *employment*—rather than a second recording variable: sub-award dollars, like prime obligations, are booked where a contract lands, one tier down the chain. Second, it puts the wedge to a different use: where that literature corrects a fiscal multiplier or a causal-incidence estimate, we build a concentration, vulnerability, resilience, and mobilization apparatus on top of it and show that the weight choice does not merely shift a number but reverses the qualitative verdict the apparatus returns.

The framework fixes what must be measured. Characterizing a defense industrial geography in these terms requires a national county-by-industry sample, an input-output structure through which to propagate import dependence, and spatial weights built from production linkages rather than adjacency.

Section 3 specifies the data that meet these requirements, and Sections 4–5 construct the indices and the cascade and surge-capacity methodology the framework calls for.

3 Data

This section documents the data sources and the analysis sample. The constructions that operate on them—the closed-loop vulnerability index and the empirically-derived spatial-weights matrices (Section 4), and the cascade and surge-capacity methodology (Section 5)—are deferred to the methodology sections. Here we describe only what the data are and which units enter the analysis.

3.1 Data sources

The analysis draws on five sources.

BEA Input–Output Accounts (U.S. Bureau of Economic Analysis, 2017). The 2017 benchmark Use, Supply, Import, and Direct Requirements tables, aggregated to the BEA U.Summary level of 116 industries, are the basis for the Leontief-propagated vulnerability index and the input-output-derived spatial weights of Section 4. The 2007 and 2012 benchmark tables are additionally used for the time-varying vulnerability analysis of Section 6.

BLS Quarterly Census of Employment and Wages (U.S. Bureau of Labor Statistics, 2025). County-by-NAICS-4 employment, pooled over 2008–2024, is the basis for the county-industry employment matrix \mathbf{E} that enters every structural construction in the paper. Pooling over the full window captures the long-run structural composition of each county’s economy rather than year-to-year noise, the same convention used for the vulnerability index in Section 4.

One feature of the employment matrix warrants disclosure. About 54.6% of its 381,524 county-by-industry cells (3,289 counties \times 116 industries) are empty—a mixture of genuine structural absence (most counties host no establishment in a given narrow industry) and BLS non-disclosure suppression, which blanks any cell that would reveal a confidential or dominant single employer; the public data do not let the two be separated. Emptiness is heaviest in exactly the narrow defense-manufacturing industries this paper weights on, where between 84% and 99% of counties record no employment. We do not impute the blanked cells. Pooling over 2008–2024 mitigates the loss—any county-industry pair observed in even a single year survives—but a small single-plant facility suppressed throughout remains zeroed, so the production weight, if anything, understates defense exposure in precisely those small specialized counties and is conservative there.

Three defense-exposure weights. A county’s defense exposure can be measured by where defense contracts are *recorded*, where defense is *produced*, or where observed sub-tier dollars *flow*. These are distinct geographies that disagree by an order of magnitude—a divergence documented for federal spending more broadly by Gerritse and Rodríguez-Pose (2018) and Park et al. (2026), which we quantify here for the defense base at county resolution (Section 6). We therefore carry three exposure weights $\delta^{(k)}$ side by side throughout, and take the production weight as primary:

- **Production (primary):** county c ’s share of national defense employment,

$$\delta_c^{\text{emp}} = \frac{\sum_{i \in \mathcal{D}} \mathbf{E}_{ci}}{\sum_{c'} \sum_{i \in \mathcal{D}} \mathbf{E}_{c'i}},$$

built from the QCEW employment matrix of Section 3.2. This measures where defense is physically produced and covers all fifty states.

- **Obligation (recording):** county c 's share of national DoD prime-award obligations, 2008–2024, from USAspending.gov (U.S. Department of the Treasury, 2025), assigned to the county of primary place of performance. This measures where prime contracts are *booked*; for a platform whose prime is a single final-assembly site, it concentrates at that site.
- **Sub-award (observed flow):** county c 's share of inbound FY2023 DoD sub-award dollars (the \geq \$1 million records described below). This measures observed Tier-1 flow, on a narrower base.

Unless otherwise noted, “ δ ” denotes the production weight δ^{emp} ; the obligation and sub-award weights are reported alongside as the contrast that motivates the analysis. A contract action counts as DoD when its awarding top-tier agency is the Department of Defense (Common Government-wide Accounting Classification 097), spanning all military departments and defense agencies, and obligations are assigned to the county of primary place of performance rather than the recipient’s headquarters. Two scope boundaries follow and are revisited in Section 10: U.S. Army Corps of Engineers civil-works obligations are included, though as construction they fall outside the DEFENSE13 industries of Section 3.3 and so do not enter the vulnerability index; and the Department of Energy’s National Nuclear Security Administration—which operates the nuclear-weapons production complex—lies outside the Department of Defense, and thus outside this frame. Obligations are dated by the calendar year of the contract action. Dollar figures are deflated to real 2017 dollars using the BEA GDP price deflator; the structural indices themselves are built from employment and input-output *shares* and from DoD-obligation shares, so they are invariant to the deflation base.

USAspending FY2023 sub-award records (U.S. Department of the Treasury, 2025). The USAspending search API, queried for DoD sub-awards to U.S.-domestic sub-recipients and sorted by amount, returns its first 10,000 records (the page-100 ceiling), each of sub-award amount at least \$1 million and together totaling \$75.4 billion; the bulk-download endpoint did not complete within the project window. These largest-by-amount flows underlie the W_3 matrix (Section 4). Sub-awards are geolocated to the sub-recipient’s business location and cover federal fiscal year 2023 (October 2022–September 2023); because the search endpoint returns no place-of-performance for sub-awards, the W_3 geography reflects recipient location rather than the place-of-performance basis used for the prime panel. W_3 is therefore built from a conservative lower-envelope of observed Tier-1 sub-award flow—the largest-by-amount records only—rather than the full sub-award universe, and is read as such throughout (Section 10).

TIGER/Line 2022 county geographies (U.S. Census Bureau, 2022). Census county and county-equivalent shapefiles provide the choropleth basis for the maps in Sections 6–8.

Table 1 summarizes the analysis sample and the distribution of the constructed indices.

3.2 Sample and geographic harmonization

The structural analysis sample is **3,289 county and county-equivalent units covering all fifty states and the District of Columbia**, with employment pooled over 2008–2024 across the 116 BEA U.Summary industries. This is the universe on which the vulnerability index V_2 , the \mathbf{W} matrices, the cascade simulation, and the surge-capacity index are all computed; it is the sample referenced as $c \in \{1, \dots, 3289\}$ in Section 4. No states are excluded from the structural analysis.

Two geographic harmonizations bear on the results and are stated explicitly here because one of them is load-bearing for the paper’s central findings.

Table 1: Descriptive statistics for the analysis sample and constructed indices. Defense employment share, DoD obligation share, and V_1 import dependence are expressed as percentages. V_2 is the Leontief-propagated vulnerability index (per million of national total). Surge capacity is the Hidalgo-Hausmann density to DEFENSE13, weighted by commercial fraction.

Variable	Mean	SD	Median	Min	Max
Sample structure					
N counties (CONUS, excl. AK/CT/HI/PR)				3,072	
N years				17	2008–2024
N county-years				52,224	
N BEA U.Summary industries				116	
N defense industries (DEFENSE13)				13	
Defense employment + obligations					
Defense emp share	3.0%	5.1	1.6%	0.0	93.7
DoD obligation share	0.0%	0.2	0.0%	0.0	3.5
Defense import dep (V_1)	7.6%	5.5	5.8%	0.0	23.2
Vulnerability + surge (constructed)					
V_2 vulnerability per million	32.67	169.56	0.51	0.00	4,352.64
Surge capacity (thousands)	73.55	303.45	6.65	0.00	8,290.29

Sample excludes Alaska, Connecticut (CT migrated counties to planning regions in 2022), Hawaii, and Puerto Rico (non-standard FIPS or incomplete BEA coverage). Virginia FIPS harmonized to BEA combined codes.

Connecticut. Connecticut replaced its eight counties with nine planning regions as its primary statistical geography in 2022, and the federal statistical system issued new FIPS codes accordingly. To keep a single, consistent geography for Connecticut across the pooled 2008–2024 window—which spans the regime change—we represent the state by its legacy county FIPS (for example, Hartford County, FIPS 09003): the production employment that pre-dates 2022 is recorded under the legacy counties, and the post-2022 planning-region values (obligations and FY2023 sub-awards, for example the Capitol Planning Region 09110) are folded back into their dominant constituent county (09110 into Hartford 09003). The inclusion is load-bearing because the structural stack rests on QCEW employment and the national input-output structure rather than on BEA county-level GDP, so it can incorporate Connecticut, Alaska, and Hawaii where a reduced-form county-GDP panel cannot. This reincorporation is load-bearing: Hartford County—home to Pratt & Whitney’s F135 engine line—carries real, nonzero defense exposure (it ranks fourteenth nationally on the production-weighted vulnerability index of Section 6, and also receives substantial direct DoD obligations), and is a member of the F-35 cluster. Earlier drafts that inherited the 48-state reduced-form frame dropped Connecticut entirely, so the engine that powers the F-35 registered a mechanical zero. Reincorporating the state—which the structural stack can do because it rests on employment and input-output structure rather than on the BEA county GDP that is unavailable for post-2022 Connecticut—recovers it.

Virginia. Virginia’s independent cities are harmonized to the BEA combined county codes for consistency with the input-output accounts; for the DoD-exposure weighting the combined codes are exploded back to their raw Census constituents to align with the QCEW employment geography. The full crosswalk is documented in Appendix E.

For completeness: the optional reduced-form panel analysis reported in Appendix G uses a narrower 3,072-county frame that drops Alaska, Connecticut, Hawaii, and Puerto Rico, because BEA post-2022 county-level GDP is unavailable for them. That restriction is a property of the reduced-form county-GDP panel only; it does not apply to, and does not constrain, the structural analysis that is this paper’s

contribution.

3.3 The DEFENSE13 industry set

Throughout the paper, defense-relevant production is operationalized through a set of thirteen BEA U.Summary industries we label DEFENSE13. These are the industries that are the primary recipients of Department of Defense procurement— aerospace primes, electronics and semiconductors, fabricated metal and ordnance, and the defense-intensive professional services (engineering, computer-systems design, and scientific research and development)—and that are cleanly mappable to NAICS-4 industries in the QCEW employment data.¹ Table 2 lists the set with each industry’s BEA code, its NAICS-4 underliers, and its pooled annual employment.

Table 2: The DEFENSE13 industry set used throughout the paper. These are BEA U.Summary industries (one step below BEA’s Summary level, 138 industries total at U.Summary). They were selected for being primary recipients of DoD procurement (Lockheed, Boeing, Raytheon, GD, Northrop, BAE) and for being mappable to NAICS-4 industries in QCEW employment data.

BEA code	Industry	NAICS-4 underliers	Annual emp
3364	Aerospace	3364 (aerospace product/parts)	338,314
3365AO	Ships / tanks / rail	3365 + 3366 + 3369 (rail, ship, other)	98,466
3342	Communications eq.	3342 (communications eq.)	66,961
3344	Semiconductors	3344 (semiconductors)	302,025
3345	Instruments	3345 (instruments)	328,603
334X	Other electronics	3343 + 3346 + 3349 (other electronics)	13,571
332	Fab metal / ammo	332 fab metal (e.g. 3329 ordnance)	1,059,616
5413	Engineering svc.	5413 (engineering services)	1,487,654
5415	Computer systems svc.	5415 (computer systems design)	1,922,969
5417	R&D services	5417 (scientific R&D)	654,228
3361MV	Motor vehicles	3361 (motor vehicles)	35,962
3332OM	Other machinery	3332 + 3333 + 3335 + 3336 (other machinery)	536,345
3341	Computer peripherals	3341 (computer + peripherals)	104,049

Annual employment averaged across pooled 2008–2024 QCEW data. Mapping from U.Summary to NAICS-4 via ‘naics4_to_usummary.csv’. The 13 industries together account for roughly 1.4 million annual defense-related employees nationally.

DEFENSE13 is the set \mathcal{D} that enters the direct import dependence index of Section 4 (Equation (1)) and the defense relatedness density of Section 5 (Equation (12)). The full NAICS-4 to BEA U.Summary concordance is documented in Appendix C.

With the sources and the sample fixed, Section 4 constructs the closed-loop vulnerability index and the empirically-derived spatial-weights matrices from them, and Section 5 specifies the cascade and

¹The set is selected a priori, but it aligns with where defense procurement actually flows: over 2008–2024 roughly three-quarters (about 76%) of DoD prime obligations fall in the two NAICS supersectors from which DEFENSE13 is drawn—manufacturing (NAICS 31–33, about 51%) and professional, scientific, and technical services (NAICS 54, about 25%)—with most of the remainder in construction ($\approx 7\%$), finance, facilities, and transportation, activities the structural production lens deliberately excludes. This is a supersector-level concordance and does not by itself isolate the thirteen constituent NAICS-4 industries; a precise mapping of obligations to them requires contract-level re-aggregation and is left to future work (Section 10).

surge-capacity methodology that the empirical sections apply.

4 Methodology I: Closed-Loop Vulnerability and Spatial Weights

This section introduces the two foundational methodological constructions that the empirical analysis rests on: a closed-loop vulnerability index propagated through the Leontief inverse (§4.1), and a family of empirically-derived spatial weights matrices that operationalize defense-industry co-location and input-output relatedness rather than geographic adjacency (§4.2). The W-matrix subsection is the load-bearing methodological contribution of the paper, and it occupies most of the section’s length accordingly.

4.1 Closed-loop vulnerability

The motivating question is straightforward: *if the United States were cut off from imports of intermediate inputs, how much of each county’s defense production capacity would be unable to operate?* Standard measures of defense-industrial concentration—the Herfindahl index on supplier counts, the four-firm concentration ratio, geographic dispersion metrics—do not answer this question because they do not propagate dependence through the input-output structure. The classic empirical work on U.S. defense regional concentration (Markusen, 1986; Markusen et al., 1991) mapped employment shares and contract distributions; Pisano and Shih (2009) argued that industrial-commons concentration creates capabilities that cannot be reconstituted under shock. The 2024 U.S. National Defense Industrial Strategy (U.S. Department of Defense, 2024a) identifies supply-chain resilience as a strategic priority but provides no quantitative sub-national measure of the structural exposure such resilience would target.

We construct two complementary indices. The first, V_1 , measures direct import dependence at the county level. The second, V_2 , propagates that exposure through the full input-output structure following the Leontief framework.

Direct import dependence

For each U.S. county $c \in \{1, \dots, 3289\}$, we define direct defense-industry import dependence as the employment-share-weighted average of import shares across the defense-relevant industry set \mathcal{D} (DEFENSE13; see Section 3):

$$V_1[c] = \sum_{d \in \mathcal{D}} \text{def_share}[c, d] \cdot m_d, \quad (1)$$

where $\text{def_share}[c, d] = e_{c,d} / \sum_{d' \in \mathcal{D}} e_{c,d'}$ is the share of county c ’s DEFENSE13 employment that resides in industry d , and m_d is the share of intermediate inputs to industry d that are imported nationally (from the BEA Use and Import tables). $V_1[c]$ is bounded in $[0, 1]$ and captures only *direct* import dependence: how much of the inputs that county c ’s defense industries directly consume come from imports. It ignores the indirect channel where defense industries depend on non-defense industries that themselves depend on imports.

Leontief-propagated vulnerability

The closed-loop framework propagates direct exposure through the full input-output structure. Following Leontief’s foundational treatment (Leontief, 1936) and the modern textbook exposition in Miller

and Blair (2009), the total-requirements matrix \mathbf{L} is the inverse of the difference between the identity and the Direct Requirements matrix:

$$\mathbf{L} = (\mathbf{I} - \mathbf{A})^{-1}, \quad (2)$$

where \mathbf{A} is the BEA Direct Requirements matrix at U.Summary level (116 industries; U.S. Bureau of Economic Analysis, 2017). Element L_{ij} captures the total (direct plus all indirect) dollar value of industry i 's output required to produce one dollar of industry j 's final output. The Leontief inverse is the standard tool for total-requirements input-output analysis; its application to defense import vulnerability dates at least to Kellman et al. (1996), whose national, product-level model the present paper spatializes to the county level.

The closed-loop vulnerability index aggregates each county's DEFENSE13 employment, weights it by the county's share of national DoD obligations to capture defense exposure, and propagates the underlying import dependence through \mathbf{L} :

$$V_2[c] = \delta_c \cdot \mathbf{e}_{def,c}^\top \cdot \mathbf{L} \cdot \mathbf{m}, \quad (3)$$

where δ_c is county c 's share of national DoD obligations (pooled 2008–2024 from USAspending.gov; U.S. Department of the Treasury, 2025), $\mathbf{e}_{def,c}$ is the $|\mathcal{D}|$ -dimensional column vector of county c 's DEFENSE13 employment from BLS QCEW (U.S. Bureau of Labor Statistics, 2025), and \mathbf{m} is the 116-dimensional national import-share vector from the BEA Import accounts. We report $V_2[c]$ in units of per million of the national total so that cross-county comparisons are unit-free.

The interpretation of V_2 is the total (direct plus indirect) import-dependence-weighted defense exposure at the county level. A county that employs many workers in a defense industry that itself depends on imported semiconductors, specialty alloys, or foreign-sourced engineering services will carry high V_2 even if its direct defense-industry import dependence is moderate, because the propagation through \mathbf{L} captures the cascade. Conversely, a county whose direct defense-industry employment is modest but whose underlying inputs are entirely domestically sourced will carry low V_2 .

Implementation notes: the baseline V_2 analysis uses 2017 BEA benchmark tables for \mathbf{A} and \mathbf{m} . We pool BLS QCEW employment data over 2008–2024 to construct $\mathbf{e}_{def,c}$, which captures the long-run average defense-industry employment mix at each county and reduces year-to-year noise. The time-varying analysis reported in Section 6.4 recomputes V_2 under the 2007, 2012, and 2017 BEA benchmarks in turn, holding $\mathbf{e}_{def,c}$ and δ_c fixed, to isolate the effect of changes in the underlying import-dependence and input-output structure on national vulnerability.

Substitutable exposure and the irreducible floor

V_2 measures import *content*: the share of defense production that embodies imported intermediate inputs at current technology. It is not, by itself, a measure of what could not be produced under autarky, because some imported inputs have ready domestic substitutes while others do not—and the closed-loop question turns on that distinction. We therefore split the national import vector \mathbf{m} into two components using the BEA import accounts. *Comparable* imports are commodities for which a domestic counterpart exists (the numbered BEA commodity rows); under a cutoff they are substitutable in principle, at some adjustment cost, by reallocating domestic production. *Noncomparable* imports (the BEA “noncomparable imports” row, S00300) are goods and services with no domestic counterpart and constitute an irreducible bottleneck. Propagating each component through the Leontief inverse separately decomposes the index,

$$V_2[c] = \underbrace{V_2^{\text{comp}}[c]}_{\text{substitutable}} + \underbrace{V_2^{\text{floor}}[c]}_{\text{irreducible}}, \quad (4)$$

where V_2^{floor} propagates only noncomparable imports and is the assumption-light *autarky floor*—the exposure that would remain even if every substitutable import were replaced domestically. Between the endpoints we report a one-parameter bridge $V_2(\sigma) = V_2^{\text{floor}} + V_2^{\text{comp}} \cdot f(\sigma)$, sweeping a substitution parameter from $f = 1$ (no substitution, recovering V_2) to $f = 0$ (full substitution, recovering V_2^{floor}), so the headline exposure and the irreducible floor are the endpoints of a stated range rather than competing point estimates. Section 6 reports the decomposition; the substantive result is that the great majority of headline vulnerability is substitutable, while the irreducible floor stays concentrated on the same counties.

4.2 Empirically-derived spatial weights matrices

A spatial weights matrix \mathbf{W} specifies, for each pair of counties (c, d) , the strength of the relationship through which shocks at one county propagate to the other. The choice of \mathbf{W} is theoretically substantive: it encodes the researcher’s prior about which counties influence which others. For most applications of spatial econometrics, \mathbf{W} is constructed from *physical* adjacency—queen contiguity, distance bands, or nearest-neighbor graphs—because the relationship being studied (retail competition, labor commuting, residential sorting) is constrained by physical distance (Anselin, 1988; LeSage and Pace, 2009). For defense industrial production, this choice is wrong.

Why standard \mathbf{W} is wrong for defense

Consider Tarrant TX, the county that hosts Lockheed Aeronautics’ F-35 final-assembly plant. Under queen contiguity, Tarrant’s neighbors are Parker, Wise, Denton, Dallas, Ellis, and Johnson counties—all in North Texas, and none of them are part of the F-35 supply network. The actual F-35 supply-network neighbors of Tarrant include Hartford CT (Pratt & Whitney’s F135 engine line), Sedgwick KS (Wichita commercial aerostructures and Spirit AeroSystems), Essex MA (Boston defense electronics), Brevard FL (Space Coast aerospace), and Snohomish WA (Boeing’s Everett commercial aerospace complex)—all 1,200 to 2,400 miles from Fort Worth. A vulnerability assessment built on geographic \mathbf{W} has no edges from Tarrant to anywhere outside North Texas. The supply network is simply not visible.

The methodological literature has anticipated this problem. Conley (1999) argued that \mathbf{W} should reflect the *mechanism* through which dependence propagates, not a default geographic convention. Combes et al. (2008) make the same point for economic geography: agglomeration spillovers travel along supply linkages, knowledge networks, and labor markets, not along shared-border boundaries. The empirical practice has lagged the theoretical insight: most defense industrial work, when it uses spatial econometric tools at all, continues to use contiguity or distance \mathbf{W} out of convention. We argue that for defense industrial questions, an empirically-derived \mathbf{W} that operationalizes industrial co-location and input-output relatedness—a far better proxy for the supply network than physical adjacency, though not a direct observation of it—is the appropriate choice. The remainder of this section constructs four such variants.

Construction principle

We build \mathbf{W} from the bilinear product of county-by-industry employment and the BEA Direct Requirements matrix:

$$\mathbf{W} = \mathbf{E} \cdot \mathbf{A} \cdot \mathbf{E}^{\top}, \quad (5)$$

where \mathbf{E} is the $3,289 \times 116$ county-industry employment matrix from BLS QCEW pooled 2008–2024, and \mathbf{A} is the 116×116 Direct Requirements matrix from the 2017 BEA benchmark I-O accounts. The element $\mathbf{W}[c, d]$ measures the input-output *relatedness* of the two counties’ industry mixes: it is large when both counties employ workers in industries that, by the BEA I-O structure, exchange substantial intermediate inputs. It is a *co-location* statistic—a statement about shared industrial structure—not an observation of bilateral flow between the two counties; no shipment is recorded moving from c to d . The product is small when the counties’ industry mixes do not connect through the I-O network, regardless of how close the counties are in physical space. This is the same family of construction as the Hidalgo–Hausmann relatedness measure of Section 5, applied to counties rather than industries, and it should be read as relatedness, not as a directed supply chain.

This bilinear construction generalizes naturally. We can restrict \mathbf{A} to defense-relevant columns to isolate the defense-relevant subgraph, weight the receiver-side employment by each county’s DoD exposure, or replace the I-O-implied relatedness with co-membership in directly observed sub-awards. Each variant operationalizes a different prior about which features of the co-location structure matter for defense industrial geography.

Four \mathbf{W} variants

We construct four progressively refined variants of \mathbf{W} , which we label W_2 V1, V2, V3, and W_3 .

W_2 V1 (general input-output linkage)

is defined by Equation (5). It captures every industry pair that co-locates through the BEA I-O network, including non-defense industries. We retain V1 as a baseline that allows defense-relevant relatedness to compete with general commercial relatedness on equal terms; we will see in Section 6 that the V1 cluster structure does not isolate defense industrial geography cleanly because general-economy linkages dominate the largest clusters.

W_2 V2 (defense-restricted)

restricts the input-output flow to defense-relevant industries on the receiver side. Formally, W_2 V2 = $\mathbf{E} \cdot \mathbf{A}_{def} \cdot \mathbf{E}^\top$ where \mathbf{A}_{def} denotes \mathbf{A} restricted to its DEFENSE13 columns. W_2 V2 $[c, d]$ measures the input-output relatedness of county c ’s industries to the defense industries of county d : the construction is asymmetric in (c, d) because the column restriction conditions on d ’s defense mix, but that asymmetry is a weighting choice, not a direction of flow. It isolates the defense-relevant subgraph while still letting non-defense industries enter on the unrestricted side (since defense production draws on inputs from much of the rest of the economy).

W_2 V3 (defense-exposure-weighted; paper-grade)

further weights the receiver-side defense employment by each county’s defense-exposure share:

$$W_2 V3 = \mathbf{E} \cdot \mathbf{A}_{def} \cdot (\mathbf{E}_{def} \cdot \text{diag}(\boldsymbol{\delta}))^\top, \quad (6)$$

where \mathbf{E}_{def} is the $3,289 \times |\mathcal{D}|$ sub-matrix of \mathbf{E} restricted to DEFENSE13 columns and $\boldsymbol{\delta}$ is the 3,289-dimensional column vector of county-level defense-exposure shares. We take the *production* weight δ^{emp} (defense-employment share) as primary, so that the receiver side scales county d ’s defense employment by how much defense the county actually *produces*, and W_2 V3 $[c, d]$ captures the input-output relatedness of c ’s industries to the defense-producing part of d ’s economy. We retain the obligation-weighted V3 (substituting δ^{obl}) as a comparison matrix; because obligation booking concentrates on final-assembly sites, the two weightings yield materially different

network structures, and the contrast is itself diagnostic (Section 7). This is the \mathbf{W} matrix used for cluster identification (Section 6) and the cascade simulation (Section 7). We follow Conley (1999)’s convention in applying the exposure weighting one-sided (receiver only) rather than two-sided; a symmetric two-sided weighting would over-collapse the heterogeneity between sender industries (which we want to retain) and the defense-relevance of receiver counties.

W_3 (shared-prime co-membership from observed sub-awards)

departs from the input-output-implied framework and anchors directly in observed defense procurement data:

$$W_3[c, d] = \sum_{p \in \mathcal{P}} \min(\text{flow}(p \rightarrow c), \text{flow}(p \rightarrow d)), \quad (7)$$

where \mathcal{P} is the set of prime defense awards in our FY2023 USAspending search-API sample (10,000 records, U.S.-domestic sub-recipients only, sub-award amount \geq \$1 million), and $\text{flow}(p \rightarrow c)$ is the dollar value of Tier-1 sub-award activity flowing from prime award p into county c . This is the *bipartite co-prime projection* of the (prime, county) flow network onto the (county, county) projection, with saturation-form minimum-flow weighting following Zhou and Lü (2007); see also Newman (2004) on the general bipartite-projection framework. The min-flow convention ensures that $W_3[c, d]$ is large only when *both* counties receive substantial sub-award flow from the *same* primes—i.e., when they share supplier relationships, not merely when they both happen to receive defense work. W_3 is symmetric by construction. W_3 complements the I-O-derived W_2 variants by anchoring in observed sub-award data rather than predicted relatedness. It is a *co-membership* projection—two counties are linked when they draw on the same primes, not by any directed shipment between them—and is conservative by construction: because most county pairs share no primes, W_3 is highly sparse and the majority of counties are isolated under it, a coverage limitation we flag in the \mathbf{W} -comparison note below.

Figure 2 shows the construction pipeline for the four variants, and Figure 3 compares the three W_2 variants over the top-20 defense counties.

Properties

All four matrices are row-normalized so that each row sums to one, following the standard spatial-econometric convention (LeSage and Pace, 2009). The W_2 variants are asymmetric by construction because the sender-side and receiver-side weighting differ; for community-detection applications (Section 6) we symmetrize via $(\mathbf{W} + \mathbf{W}^\top)/2$. W_3 is symmetric by construction because the min-flow projection in Equation (7) is order-invariant in (c, d) . All matrices are top-20 thresholded per row before community detection—we retain only the 20 largest off-diagonal entries per row—to match the standard sparsity convention in spatial-econometric work and to align with the prevailing literature on industry-cluster identification. The diagonal is zero in all variants.

A note on \mathbf{W} choice and direct elasticity invariance

The choice of W_2 V3 as the paper-grade specification rests on both substantive and methodological grounds. The substantive argument has been given: V3 conditions on both defense supply structure (via \mathbf{A}_{def}) and observed DoD exposure (via δ), and is the only specification that produces cluster identification matching the public F-35 program structure (Section 6).

W Matrix Construction Pipeline (W₂ V1 / V2 / V3)

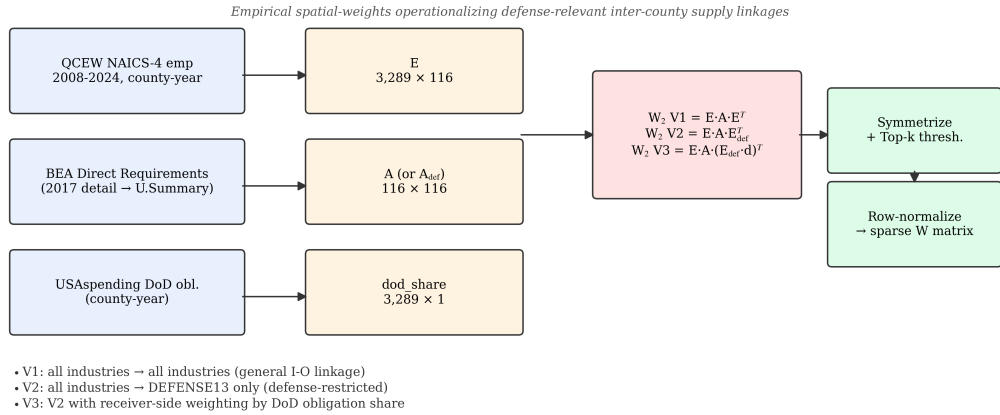


Figure 2: Construction pipeline for the W_2 and W_3 variants. \mathbf{E} from BLS QCEW pooled 2008–2024; \mathbf{A} from BEA 2017 benchmark Direct Requirements; δ from USAspending DoD obligation shares; \mathcal{P} from USAspending FY2023 sub-award records.

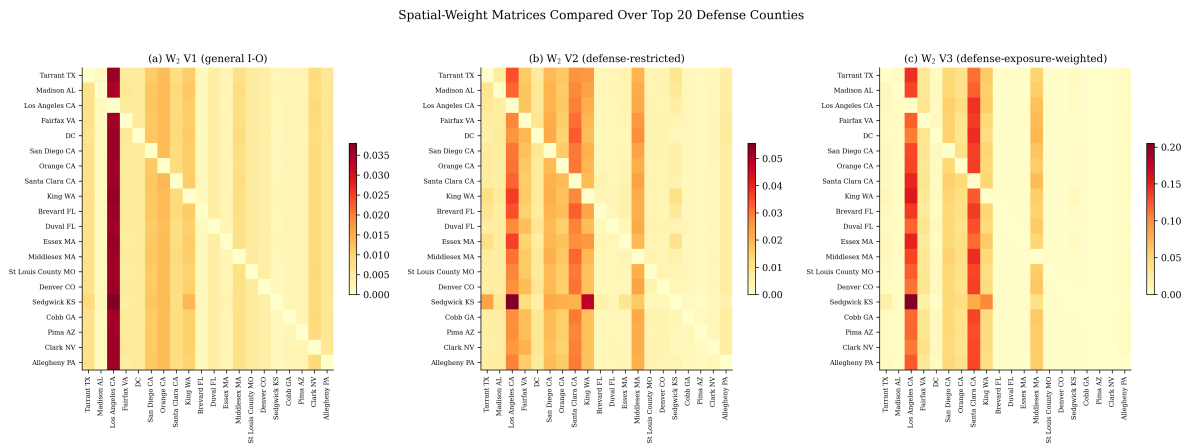


Figure 3: Heatmap comparison of the three W_2 variants over the top-20 U.S. counties by V_2 vulnerability. V1 displays diffuse linkage; V2 sharpens onto defense-relevant sub-blocks; V3 further concentrates on the DoD-exposed cells.

A separate methodological exercise—a standard spatial Durbin specification of county GDP elasticity to DoD obligations with two-way fixed effects, estimated under each of five \mathbf{W} matrices (queen contiguity, W_2 V1, V2, V3, and W_3), following Kelejian–Prucha (1998) within-2SLS with $\mathbf{W}^2\mathbf{X}$ as instrument for the spatial lag and Lee–Yu (2010) two-way-fixed-effects absorption—produces a useful aside. The local direct elasticity of GDP to \$1 of DoD obligation is invariant across all five \mathbf{W} matrices, at approximately \$0.0002 per dollar. The *implied spillover multiplier*, by contrast, is not point-identified. Computed as the spatial-autoregressive total effect $(\beta + \theta)/(1 - \rho)$, it is reportable only where the spatial process is stable: of the thirty $\mathbf{W} \times$ outcome \times model cells, eighteen satisfy $|\rho| < 1$ and twelve do not, the latter leaving the multiplier formula undefined (Table 26). Even among the ρ -stable GDP specifications the multiplier flips sign across networks—a slight negative under the defense-restricted input-output matrix (W_2 V2, about $-\$0.013$) and essentially zero under queen contiguity and observed FY2023 sub-award flows (about $+\$0.002$)—while the general-linkage (W_2 V1) and defense-exposure (W_2 V3) matrices drive ρ outside the stable $(-1, 1)$ range altogether, voiding the formula. Sign and magnitude turn entirely on the spatial autoregressive coefficient ρ that each \mathbf{W} produces. The fragility is not confined to the choice among matrices: correctly matching Virginia’s combined-FIPS Beltway counties to the panel (rather than dropping them as isolated rows) by itself flips the industry-mix multiplier’s sign, and W_3 leaves the great majority of counties isolated, so its near-queen multiplier reflects coverage rather than network structure. We therefore present this comparison as a *caution*, not an estimate: the direct local elasticity is robust, but any spillover multiplier in a defense industrial setting is hostage to an untested and non-robust propagation assumption, and conventional geographic \mathbf{W} is not a defensible default. The full five- \mathbf{W} , three-outcome comparison is reported in Table 26.

This section has defined the closed-loop vulnerability index V_2 and constructed the empirically-derived spatial weights matrices W_2 V1, V2, V3, and W_3 that operationalize the defense relatedness network at the county level. Section 5 introduces the second methodological pillar of the paper: cascade simulation as a resilience metric on W_2 V3, and the Hidalgo–Hausmann surge capacity index adapted to sub-national defense industrial geography.

5 Methodology II: Cascade Simulation and Surge Capacity

This section specifies the two methodological constructions that the resilience and mobilization analyses rest on: a cascade simulation that converts the static cluster structure of Section 4 into a dynamic test of structural fragility (§5.1, applied in Section 7), and a Hidalgo–Hausmann surge-capacity construction that measures latent dual-use industrial capacity at the county level (§5.2, applied in Section 8). Both build on the W_2 V3 co-location matrix (Equation (6)) and the pooled 2008–2024 BLS QCEW county-industry employment matrix \mathbf{E} introduced in Section 4. The emphasis throughout is on the defense of each design choice: a methodological contribution is only as credible as the robustness of the parameters it turns on, and the sensitivity evidence summarized here is reported in full in Section 9.

5.1 Cascade simulation as a resilience metric

Section 6 establishes that defense industrial vulnerability is statically concentrated in a single Louvain community anchored on Tarrant TX. Static concentration is necessary but not sufficient for fragility: a cluster can be concentrated and still resilient if, upon removal of its anchor, the surviving counties re-coalesce around an alternate organizing node. Whether they do is an empirical question about the topology of the relatedness network, and it requires a perturbation test rather than a concentration index.

The procedure

The cascade simulation proceeds in four steps. (1) *Removal*: the candidate county or counties are deleted from the dense $W_2 V_3$ matrix (Equation (6)), together with their rows and columns. (2) *Symmetrization*: the residual matrix is symmetrized via $(\mathbf{W} + \mathbf{W}^T)/2$. Louvain modularity maximization operates on an undirected graph; symmetrization encodes the economic interpretation that two counties are structurally related when their industry mixes are I-O-related regardless of orientation, and is the same operation used for the cluster identification of Section 6. (3) *Thresholding*: the symmetrized matrix is reduced to its twenty largest off-diagonal entries per row, the sparsity convention adopted throughout the paper (Section 4); the top-three cluster composition is invariant for thresholds $k \in [10, 20]$ (Section 9). (4) *Re-detection*: Louvain modularity maximization (Blondel et al., 2008; Newman, 2004) is re-run on the perturbed graph at seed 42 and resolution $\gamma = 2.0$.

The new partition is matched to the original F-35 cluster by maximum Jaccard overlap. Writing \mathcal{C}_3 for the baseline cluster and s for a candidate community in the perturbed partition, the successor cluster is

$$s^* = \arg \max_s J(s, \mathcal{C}_3), \quad J(s, \mathcal{C}_3) = \frac{|s \cap \mathcal{C}_3|}{|s \cup \mathcal{C}_3|}, \quad (8)$$

where J is the Jaccard index between the candidate community and the baseline cluster. Matching on maximum overlap rather than, say, the largest community ensures that the simulation tracks the *most closely related* reorganized cluster and does not spuriously identify a large unrelated community as the “survivor.”

Why resolution $\gamma = 2.0$

The Louvain resolution parameter governs the granularity of the partition. We fix $\gamma = 2.0$ because it is the coarsest resolution at which the F-35 counties form their own distinct community—a 56-county cluster anchored by Tarrant TX and matching the public F-35 program structure (Section 6)—rather than dissolving into a larger regional bloc. This is not a free parameter tuned to a result: the cluster’s *identity* is stable across a band around it. Tarrant is a member at every resolution tested, and at $\gamma \geq 2.0$ the cluster’s leading production counties are consistently the engine-and-airframe core (Seattle, Snohomish, Tarrant, Hartford), with nominal cluster size declining with γ . Only at $\gamma = 1.0$ does the structure collapse—into a few mega-clusters, one absorbing more than a thousand counties including every major defense hub, a partition too coarse to resolve defense industrial geography at all. The full resolution grid is reported in Section 9; the operative point is that the cluster identity is a structural feature of the network, not an artifact of the resolution choice. We report the cascade at $\gamma \geq 2.0$ and treat the resolution as an explicit scope condition.

Resilience metrics and the bootstrap

The reorganized partition is summarized by two statistics that are defined where they are first applied, in Section 7: *cohesion* (Equation (15)), the fraction of surviving baseline-Cluster-3 counties that re-coalesce into the successor community, and *module entropy* (Equation (16)), the dispersion of the survivors across the new communities. Cohesion near one means the cluster reconstitutes itself and is resilient to the removal; low cohesion with high entropy means it dissolves into the general economy.

Louvain modularity maximization is stochastic in its node-visit order, so a single partition can over- or under-state fragmentation. Every scenario is therefore bootstrapped across 100 independent Louvain

seeds, and we report the median cohesion together with its 95% bootstrap interval (the 2.5th and 97.5th percentiles across seeds). The structural diagnostics that require a single reference partition—the defector count, successor size, module entropy, and successor V_2 total—are reported from the seed-42 run. The 100-seed design is what licenses the resilience claim, and it matters here because the F-35 cluster’s seed distribution is bimodal: a single partition can fall at either a “reconstitutes” or a “scatters” mode, so the median across seeds—not any one draw—is the basis for every cascade conclusion in Section 7.

Resolution scope and the null model

Two further checks, specified here and reported in Sections 7 and 9, situate the cascade result. First, because the F-35 cluster is itself a $\gamma = 2.0$ construct, we sweep cohesion across $\gamma \in \{1.0, 1.5, \dots, 3.0\}$ and report only at the resolutions that resolve the cluster as a distinct community ($\gamma \geq 2.0$), stating the resolution as a scope condition. Second, to establish that “resilient” is not returned mechanically by the procedure, we run the identical cascade on each of the ten largest defense clusters, removing each cluster’s own top-vulnerability county and comparing the resulting bootstrap-median cohesions. A cascade that discriminates will sort the clusters—returning fragility for core-organized clusters and resilience for broadly-based ones—rather than labeling every cluster the same way; Section 7.4 reports that it does, with six of the ten single-anchor fragile and the production-dispersed F-35 cluster among the three redundant exceptions.

The seven scenarios

The cascade is run under seven node-removal scenarios, each designed to test a specific structural hypothesis:

Baseline (no removal)

A reproducibility check: the perturbation machinery applied to the unmodified network should return the F-35 cluster essentially intact.

Tarrant TX (FIPS 48439)

Removal of the F-35 final-assembly anchor—the county the obligation weight makes look like the linchpin. The central test: does the production-weighted cluster hang from this node?

Essex MA (25009)

Removal of a high- V_2 county. Tests whether structural importance tracks vulnerability magnitude or is distinct from it.

Tarrant + Hartford CT (48439, 09003)

Anchor plus the Pratt & Whitney F135 engine county. Tests whether the two largest production sites together fragment the cluster.

Tarrant + Wichita KS (48439, 20173)

Anchor plus the Spirit AeroSystems commercial-aerostructures county. Tests the marginal effect of removing a second major production node.

Top-5 by V_2

An extreme scenario removing the five highest-production-vulnerability cluster counties simultaneously (resolved at runtime; under the production weight these are Seattle, Snohomish, Tarrant, Hartford, and Wichita).

LA County (06037)

Removal of a high-betweenness node in the national defense relatedness network. Tests the new-economic-geography distinction between a node whose removal fragments a cluster and one whose removal redistributes membership.

The scenario results are reported in Table 8 and Figures 9–10 (Section 7).

5.2 Hidalgo–Hausmann surge capacity

The cascade answers what breaks under the loss of the integration core. The converse question—what latent commercial capacity could substitute—requires a different instrument. We adapt the Hidalgo–Hausmann product-space framework (Hidalgo et al., 2007; Hausmann and Klinger, 2007) to the sub-national defense problem, applying its relatedness logic in reverse: not “which products can a region move into,” but “given a defense target industry, which counties already sit closest to it in capability space.” The construction follows the regional-branching tradition (Neffke et al., 2011; Boschma and Frenken, 2006), in which the capabilities a region can acquire are predictable from the relatedness structure of the activities it already performs.

Revealed comparative advantage

The building block is the Balassa revealed comparative advantage of county c in BEA U.S. Summary industry i (Balassa, 1965):

$$\text{RCA}[c, i] = \frac{\text{emp}[c, i] / \sum_{i'} \text{emp}[c, i']}{\sum_{c'} \text{emp}[c', i] / \sum_{c'} \sum_{i'} \text{emp}[c', i']}, \quad (9)$$

where $\text{emp}[c, i]$ is pooled 2008–2024 employment from BLS QCEW (U.S. Bureau of Labor Statistics, 2025). $\text{RCA}[c, i]$ exceeds one when industry i is more concentrated in county c 's economy than in the national economy. We binarize at the Balassa threshold of one,

$$M[c, i] = \mathbb{1}\{\text{RCA}[c, i] \geq 1\}, \quad (10)$$

so that $M[c, i] = 1$ records that county c is *specialized* in industry i . The unit threshold is the standard Balassa convention—local share exceeds national share—and the surge results are robust to it: under alternative thresholds $\tau \in [0.5, 2.0]$ the top-20 surge counties and all six a-priori validation counties remain within the top five percent of the distribution (Section 9).

Industry relatedness

Following Hidalgo et al. (2007), the proximity between two industries is the smaller of the two conditional co-occurrence probabilities:

$$\phi(i, j) = \min\{P(M_i = 1 \mid M_j = 1), P(M_j = 1 \mid M_i = 1)\} = \frac{|M_i \cap M_j|}{\max(|M_i|, |M_j|)}, \quad (11)$$

where M_i is the set of counties specialized in industry i . $\phi(i, j)$ is high when counties specialized in one industry tend also to be specialized in the other. The min form is the standard Hidalgo definition: it is

symmetric and conservative, recording two industries as related only when the co-specialization holds in *both* directions, which guards against a ubiquitous industry appearing spuriously close to everything. Figure 4 displays the resulting 116×116 relatedness matrix with hierarchical clustering, and Table 3 reports the nearest non-defense neighbors of each DEFENSE13 industry as a sanity check—aerospace’s neighbors are other transportation equipment and fabricated metal, semiconductors connect to instruments and printed-circuit assembly, and so on. The neighbor sets are economically sensible, which validates the relatedness measure before it is put to work.

Defense relatedness density and DoD intensity

To measure how close an industry sits to the defense frontier as a whole, we aggregate ϕ over the DEFENSE13 set \mathcal{D} (see Section 3), weighting each defense target by its economic weight:

$$\phi_{\text{to-def}}[i] = \sum_{d \in \mathcal{D}} w_d \cdot \phi(i, d), \quad w_d \propto \delta_d \sum_c \text{emp}[c, d], \quad (12)$$

with weights normalized to sum to one. The weight w_d emphasizes defense industries that are both DoD-intensive and large in employment, so that proximity to a high-procurement, high-headcount defense industry counts for more than proximity to a marginal one.

The DoD-intensity term $\delta[i]$ measures how much of an industry’s capacity is already committed to defense:

$$\delta[i] = \min\left(1, \frac{\text{dod_per_worker}[i]}{\text{WAGE_REF}}\right), \quad \text{WAGE_REF} = \$100\text{k}, \quad (13)$$

where $\text{dod_per_worker}[i]$ is annual DoD obligations per worker in industry i , normalized to a \$100k output-per-worker reference scale and capped at one. DoD obligations are observed at the NAICS-2 level from a Bartik shift-share construction; within each NAICS-2 group they are allocated only to the DEFENSE13 sub-industries, proportionally to employment, with non-DEFENSE13 industries assigned $\delta = 0$. This convention reflects the economic reality that DoD manufacturing procurement concentrates in aerospace, electronics, and fabricated metal rather than in food, textiles, or printing. It carries a known limitation: at NAICS-2 granularity the intensity is uniform within a group—empirically $\delta \approx 0.535$ for DEFENSE13 manufacturing and 0.185 for DEFENSE13 services—and finer BEA defense-purchase data would refine it. We flag this explicitly rather than smooth over it; the surge ranking is nonetheless robust to the construction (Section 9).

The surge-capacity index

The three components compose into the dual-use surge-capacity index

$$\text{surge}[c] = \sum_i \text{emp}[c, i] \phi_{\text{to-def}}[i] (1 - \delta[i]),$$

defined and applied in Section 8 (Equation (17)). The reading is direct: a county scores high when it has substantial employment in industries that are close to the defense frontier ($\phi_{\text{to-def}}[i]$ large) but not already saturated by defense demand ($1 - \delta[i]$ large)—that is, when it holds commercial capacity that could be mobilized toward defense production without displacing existing defense work. As with the vulnerability index of Section 4, employment is pooled over 2008–2024 to capture the long-run structural composition of each county rather than year-to-year noise.

Table 3: Industry relatedness sanity check: nearest non-DEFENSE13 neighbors for each defense industry, by Hidalgo-Hausmann $\phi(i, j)$ proximity computed from US-county co-RCA. The neighbor sets are economically sensible — aerospace’s nearest neighbors include other transportation, fabricated metal links to machinery and primary metal, semiconductors connect to instruments and printed-circuit assembly. Sanity-validates the relatedness methodology.

BEA code	Industry	Top 5 non-defense neighbors (ϕ)
3364	Aerospace	335 (0.15), 324 (0.14), 3254 (0.13), 5174OT (0.13), 325X (0.12)
3365AO	Ships / tanks / rail	483 (0.22), 48A (0.17), 313TT (0.14), 324 (0.12), 3391 (0.12)
3342	Communications eq.	3254 (0.29), 5174OT (0.25), 335 (0.25), 5112 (0.24), 3391 (0.22)
3344	Semiconductors	3254 (0.36), 3391 (0.33), 335 (0.32), 5112 (0.24), 3313NF (0.22)
3345	Instruments	3391 (0.41), 335 (0.40), 3254 (0.35), 5112 (0.31), 4234 (0.28)
334X	Other electronics	512 (0.22), 525 (0.19), 5174OT (0.18), 5112 (0.18), 481 (0.16)
332	Fab metal / ammo	4231 (0.38), 623 (0.36), 323 (0.34), 423X (0.34), 321 (0.33)
5413	Engineering svc.	4A0X (0.28), 5416 (0.28), HST (0.27), 812 (0.26), 713 (0.25)
5415	Computer systems svc.	4234 (0.44), 518 (0.43), 5416 (0.42), 561X (0.39), 524 (0.38)
5417	R&D services	518 (0.43), 5112 (0.40), 4242 (0.37), 5418 (0.37), 4234 (0.32)
3361MV	Motor vehicles	482 (0.17), 3311IS (0.08), 3122 (0.06), 525 (0.05), 324 (0.04)
3332OM	Other machinery	326 (0.39), 3362BP (0.36), 323 (0.34), 3399 (0.32), 423X (0.30)
3341	Computer peripherals	5112 (0.20), 3254 (0.18), 5174OT (0.16), 335 (0.15), 518 (0.13)

$\phi(i, j) = \min\{P(M_i = 1|M_j = 1), P(M_j = 1|M_i = 1)\}$ where $M_{c,i} = 1$ iff county c has RCA in $i \geq 1$. Numbers in parentheses are ϕ values (higher = more closely related capabilities).

This section has specified the two methodological pillars of the resilience analysis: a cascade simulation that perturbs the W_2 V3 relatedness network and measures whether the F-35 cluster reconstitutes, and a Hidalgo–Hausmann surge-capacity index that locates latent dual-use industrial capacity across U.S. counties. Each design choice—the symmetrization and top-20 thresholding, the Louvain resolution $\gamma = 2.0$, the 100-seed bootstrap, the Balassa unit RCA threshold, the min-form proximity, and the NAICS-2 DoD allocation—has been stated together with its justification. The three empirical sections that follow apply these constructions: Section 6 maps the vulnerability index, Section 7 runs the cascade, and Section 8 computes surge capacity. Section 9 stress-tests every parameter introduced here.

6 Results I: The Geography of U.S. Defense Industrial Vulnerability

This section applies the closed-loop vulnerability index V_2 (Section 4) to the universe of 3,289 county and county-equivalent units and the W_2 V3 co-location spatial-weights matrix to U.S. defense industrial production. Four findings follow: the national distribution of vulnerability depends sharply on whether exposure is measured by where defense is *produced* or where contracts are *booked*, and the two disagree by an order of magnitude (§6.1); a single Louvain community on the relatedness network captures the F-35 production system (§6.2); on the production measure that cluster is genuinely dispersed—roughly sixteen effective counties across twenty-five states, not the two-to-three the booking measure implies (§6.3); and U.S. defense industrial vulnerability has been falling since 2007 in ways that the F-35 cluster has not participated in (§6.4).

6.1 The national distribution of V_2 vulnerability

Figure 5 shows the county-level distribution of V_2 , the Leontief-propagated defense industrial vulnerability index defined in Section 4. The scale is logarithmic; the distribution is heavy-tailed, with most counties registering values near zero and a small set of named defense hubs dominating.

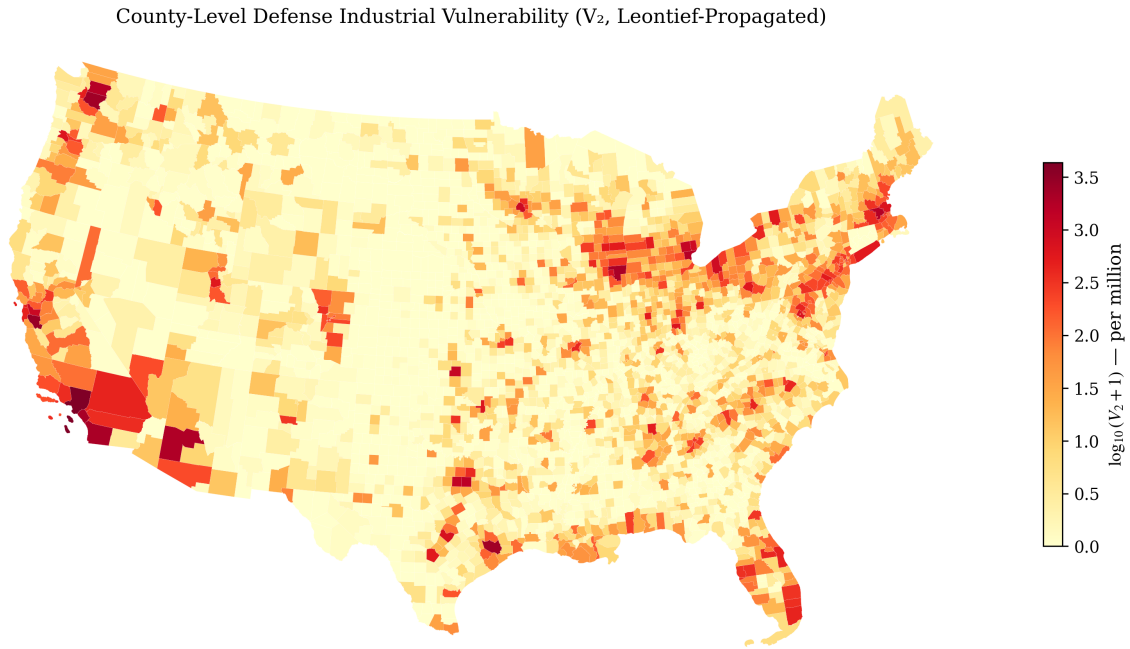


Figure 5: County-level defense industrial vulnerability (V_2), log scale. Values represent per-million of national total. Distribution is concentrated in a small set of defense hubs; most counties register near zero.

The top of the distribution is even more concentrated than the map suggests. Table 4 lists the 20 highest- V_2 counties.

On the production measure, the top of the distribution is led by the large defense-manufacturing metros: **Los Angeles registers $V_2 = 4,353$ per million**, followed by Santa Clara (2,766), Harris (2,471), and King WA / Seattle (2,462), with the top county only about 1.6 times the second-ranked. Measured by where defense is produced, national vulnerability is spread across a broad set of industrial counties rather than dominated by any one.

The contrast with the obligation measure is the section's central result, and it is sharpest at the single most prominent county. **Tarrant TX, home to Lockheed Aeronautics' F-35 final-assembly plant, books about 6.3% of all national DoD prime obligations—the largest single share in the country—yet holds only about 0.8% of national defense employment, and ranks thirteenth on the production-weighted vulnerability index (1,260 per million, roughly 1.2% of the national total).** The roughly tenfold gap between Tarrant's contract-booking share and its production share is not measurement error: it is the difference between where the F-35's prime contract is *recorded* and where the aircraft is *built*. An industrial-base assessment that weights by prime obligations—as standard practice does—will report a single-county base where the production base spans dozens of counties. The gap is robust to which obligation series is used: on the full-geography obligation weight

Table 4: Top 20 US counties by closed-loop defense industrial vulnerability (V_2), production-weighted. V_2 is the Leontief-propagated import-dependence-weighted exposure, in units of per million of national totals. Under the production weight vulnerability is led by the large aerospace-and-electronics metros; Tarrant TX (Fort Worth / Lockheed F-35), which records the largest single share of national DoD obligations, ranks 13th on production-weighted vulnerability – the order-of-magnitude gap between where contracts are booked and where defense is produced.

Rank	FIPS	State	County (hub label)	V_2 /M	Def emp	Def share (%)	DoD share (%)	Import dep (%)	Cluster
1	06037	CA	Los Angeles CA	4,353	3,820,091	5.8	3.39	12.8	1
2	06085	CA	Santa Clara CA (Silicon Valley)	2,766	3,978,981	26.2	3.53	7.8	3
3	48201	TX		2,471	2,737,151	8.0	2.43	10.2	7
4	53033	WA	King WA (Seattle / Boeing)	2,462	2,230,669	11.8	1.98	12.4	6
5	06059	CA	Orange CA	2,140	2,160,426	9.2	1.92	11.2	11
6	17031	IL		2,078	2,081,168	5.4	1.85	11.2	10
7	06073	CA	San Diego CA	2,070	2,303,114	11.3	2.04	10.1	9
8	04013	AZ	Maricopa AZ (Phoenix)	1,889	1,925,493	6.8	1.71	11.0	12
9	25017	MA	Middlesex MA	1,741	2,454,699	18.2	2.18	8.0	2
10	26163	MI	Wayne MI (Detroit / Ford River Rouge)	1,704	1,104,477	10.9	0.98	17.4	4
11	53061	WA	Snohomish WA (Boeing Everett commercial)	1,642	934,399	25.8	0.83	19.8	6
12	48113	TX	Dallas TX	1,503	1,891,061	7.7	1.68	9.0	16
13	48439	TX	Tarrant TX (Fort Worth / Lockheed F-35)	1,260	898,623	7.1	0.80	15.8	6
14	09003	CT	Hartford CT (P&W F135 — legacy FIPS)	1,211	806,142	12.0	0.76	15.9	6
15	26125	MI		1,188	1,433,864	13.3	1.27	9.3	15
16	20173	KS	Sedgwick KS (Wichita / Boeing)	1,187	645,335	17.9	0.57	20.7	6
17	06001	CA		1,053	1,317,091	12.2	1.17	9.0	17
18	27053	MN		1,002	1,146,860	8.6	1.02	9.8	20
19	26099	MI	Macomb MI (commercial auto)	947	649,365	14.4	0.58	16.4	4
20	39035	OH	Cuyahoga OH (Cleveland heavy industry)	941	820,290	7.7	0.73	12.9	1

V_2 /M: vulnerability per million of national total. Def emp: pooled 2008-2024 defense employment (DEFENSE13 industries). DoD share: DoD obligations as fraction of county economy. Import dep: weighted defense-industry import dependence.

Tarrant still books about 5.9% against its 0.8% production share. We develop the cluster-level version of this contrast in §6.3.

One qualification bears on the autarky reading throughout. V_2 measures import *content*—the share of a county’s defense-relevant production that embodies imported intermediate inputs—not an irreducible autarky bottleneck. When we decompose the BEA import matrix into *comparable* imports (commodities with a domestic counterpart, substitutable in principle under a cutoff) and *noncomparable* imports (the residual with no domestic counterpart), and propagate each through the Leontief inverse separately (Section 4), roughly 96% of the F-35 cluster’s exposure proves comparable. On the irreducible, noncomparable-only floor, the cluster’s import content collapses by about fifteen-sixteenths (Table 5). The honest answer to the closed-loop question is therefore that the binding autarky bottleneck is small: most measured import content has a domestic counterpart. What does not soften is the structure of where defense is produced, to which we now turn.

Table 5: Substitutable exposure versus the irreducible autarky floor. V_2 measures the imported-input *content* of defense production; the floor (V_4) propagates only *noncomparable* imports (BEA S00300, no domestic counterpart). Roughly 95% of the F-35 cluster’s exposure is in inputs with domestic substitutes, so the cluster’s share of the national total moves only modestly from 15.9% (content) to 12.9% (floor); nationally the irreducible floor is only $\approx 6.5\%$ of import content. The concentration is invariant to the distinction: the cluster’s effective-county count is about 16 on content and on the floor (production weight).

Unit	Import content (% natl)	Irreducible floor (% natl)	Substitutable (%)	Eff. counties
Tarrant TX	1.2	0.9	95.7	—
Essex MA	0.6	0.5	95.5	—
Brevard FL	0.5	0.4	94.8	—
F-35 cluster	15.9	12.9	95.2	16.4 / 16.4

Shares computed against the national column total under each layer. ‘Substitutable’ = $1 - V_4/V_2$ (county/cluster absolute, not the share). ‘Eff. counties’ = $1/\text{HHI}$ on within-cluster V_2 shares (content / floor). Built by `build_vulnerability_substitution.py`.

Beyond Tarrant, the production-weighted ranking is led by large, diversified manufacturing and technology centers rather than by the F-35 assembly site: Los Angeles CA (4,353 per million), Santa Clara CA / Silicon Valley (2,766), Harris TX / Houston (2,471), and King WA / Seattle (2,462) top the list, with Orange CA, Cook IL, San Diego, Maricopa AZ, Middlesex MA, and Wayne MI following; Tarrant ranks thirteenth. The roster runs through the named geography of the U.S. defense industrial base—and a long tail of mid-sized defense-related manufacturing counties in Pennsylvania, Ohio, Indiana, Illinois, Utah, and Oklahoma. Figure 6 visualizes the distribution as a bar chart.

The standard concentration metrics used in industrial-base analysis— the Herfindahl index on raw county distributions, the four-firm concentration ratio, geographic dispersion indices—do not capture this structure. Markusen and colleagues’ classic *Rise of the Gunbelt* (Markusen et al., 1991) documented the regional pattern of defense procurement in the 1980s using employment shares and contract distributions but did not propagate dependence through input-output linkages. The DoD’s 2024 National Defense Industrial Strategy (U.S. Department of Defense, 2024a) identifies supply chain resilience as a priority but does not provide a county-level quantification of structural vulnerability. The vulnerability index V_2 fills this gap by aggregating defense-industry import dependence through the Leontief inverse, weighted by each county’s share of national defense production (Section 4).

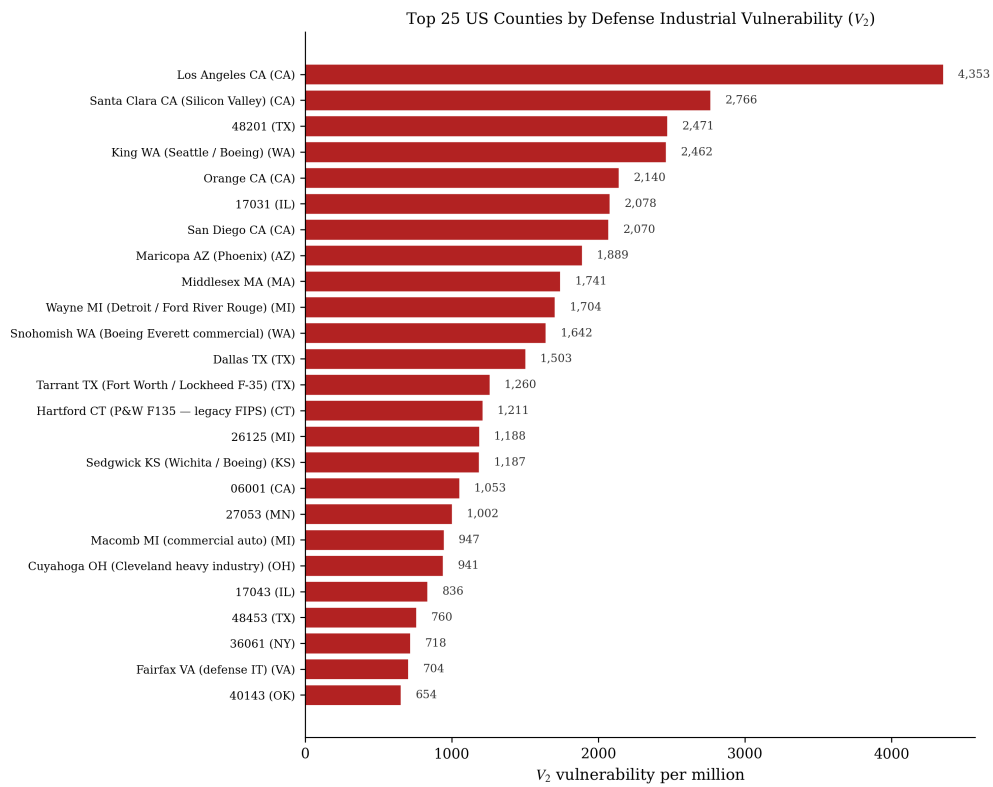


Figure 6: Top 25 U.S. counties by V_2 vulnerability (production weight). Los Angeles leads; Tarrant TX—first under the obligation weight—ranks thirteenth.

6.2 Cluster identification on the defense relatedness network

Louvain modularity-maximization (Newman, 2004; Blondel et al., 2008) applied to the symmetrized W_2 V_3 relatedness matrix (top-20 thresholded, seed=42, resolution $\gamma = 2.0$) partitions the 3,289-county universe into 2,203 communities. The distribution is heavily skewed: the median community size is one (2,169 singletons), but the top of the distribution captures the substantive defense industrial geography. Table 6 summarizes the ten largest clusters by total V_2 .

Table 6: Top 10 defense industrial clusters by total vulnerability (production weight). Clusters identified by Louvain community detection on W_2 V_3 (defense-exposure-weighted I-O linkage). The F-35 cluster (ID 6; the community containing Tarrant TX and Lockheed’s final-assembly plant, led on V_2 by King WA / Seattle) leads at about 16% of the national total, spanning 56 counties across 25 states. ‘N fragile’ = members in the fragile-chokepoint quadrant; ‘N reserve’ = members in the mobilization-reserve quadrant.

Rank	ID	Anchor county	N counties	N states	V_2 /M	DoD share (%)	Import dep (%)	Def emp	Surge cap	N fragile	N reserve
1	6	King WA / Seattle (aerospace)	56	25	16,635	12.19	13.6	13,729,136	20,818,890	3.000	0.000
2	1	Los Angeles CA	144	34	9,929	6.95	14.3	7,829,640	13,968,116	26.000	0.000
3	19	DuPage IL	28	21	4,477	4.75	9.4	5,352,636	11,446,079	0.000	0.000
4	20	Hennepin MN / Minneapolis	27	18	4,302	3.93	11.0	4,420,897	8,393,423	1.000	1.000
5	4	Wayne MI / Detroit	79	19	4,229	2.74	15.5	3,080,775	5,506,664	0.000	6.000
6	11	Orange CA (def. electronics)	39	22	3,673	3.00	12.3	3,375,293	5,421,558	1.000	3.000
7	3	Santa Clara CA (semiconductors)	79	28	3,502	4.30	8.1	4,843,050	3,010,680	0.000	2.000
8	10	Cook IL / Chicago	43	15	3,436	2.64	13.0	2,969,139	6,635,057	7.000	0.000
9	2	Middlesex MA / Boston	98	27	3,223	4.20	7.7	4,733,214	4,621,369	0.000	4.000
10	9	San Diego CA (Navy)	43	23	3,062	2.65	11.6	2,981,736	3,152,101	2.000	3.000

All values pooled 2008-2024 except DoD share (2008-2024 obligations as fraction of county economic output). Surge cap = Hidalgo-Hausmann sum within cluster.

The top cluster—which we identify throughout as the community containing Tarrant TX (the F-35 final-assembly anchor), rather than by its integer Louvain label, which is not stable across the exposure weight—contains 56 counties spanning 25 states and aggregates to a total V_2 of 16,635 per million, about 16% of the national total. This is roughly two-thirds again the vulnerability of the next cluster (the Los Angeles-anchored aerospace-and-electronics community, 144 counties, V_2 total 9,929 per million). The F-35 production system is thus visible as a single community on the defense relatedness network.

Identifying this cluster with the F-35 supply system is not an identification assumption but a confirmation: Tarrant TX (Lockheed Aeronautics’ F-35 final-assembly site) anchors the cluster; Essex MA and Hartford CT (engine production via Pratt & Whitney) are included; Wichita KS (commercial aerostructures, Spirit AeroSystems) appears. This matches the public F-35 program supply structure. The clustering algorithm did not have access to program-level information; it identified the community purely from the supply network’s combinatorial structure.

The result is consistent with the regional-branching and related-variety literature (Boschma and Frenken, 2006; Frenken et al., 2007; Neffke et al., 2011): industries that share input requirements and production capabilities tend to co-locate, and modularity-based community detection identifies these capability clusters as cohesive sub-networks.

6.3 The F-35 cluster: dispersed production, singular booking

The F-35 cluster spans 56 counties across 25 states. How concentrated that base *is* depends entirely on the exposure weight, and the two answers differ by an order of magnitude. The within-cluster Herfind-

ahl index on county V_2 shares, defined as

$$H_3 = \sum_{c \in \mathcal{C}_3} \left(\frac{V_2[c]}{\sum_{c' \in \mathcal{C}_3} V_2[c']} \right)^2, \quad (14)$$

yields, on the *production* weight, $H_3 = 0.061$, equivalent to an effective number of counties

$$N_{\text{eff}} = 1/H_3 \approx 16.4.$$

Measured by where defense is produced, the cluster genuinely behaves like sixteen-odd counties of activity, with the top county (King WA / Seattle) holding about 15% of cluster vulnerability and the top five counties about 47%. Recompute the same Herfindahl over the same counties on the *obligation* weight and it more than triples to 0.212 ($N_{\text{eff}} \approx 4.7$), with Tarrant’s share of the cluster rising from 8% to 43%—the booking measure pulls the cluster toward the single final-assembly site. The discrepancy is not a robustness wrinkle; it is the result. Where contracts are recorded is singular, because the F-35’s prime is a single final-assembly site; where the platform is produced is dispersed across the national aerospace manufacturing base. The geographic spread across twenty-five states is real industrial spread, not the political district-distribution artifact (Markusen, 1986; Markusen et al., 1991) that a contract-booking metric would suggest. The production concentration is also robust to the substitutable/irreducible distinction of §6.1: recomputed on the noncomparable-only floor, the within-cluster effective county count is essentially unchanged at about 16 (Table 5).

Figure 7 shows the cluster from two angles. The left panel maps the geographic footprint. The right panel renders the within-cluster relatedness network on the production-weighted W_2 V_3 . The most network-central node is King WA / Seattle, whose within-cluster betweenness exceeds that of any other member; Tarrant, the final-assembly anchor, sits among the periphery on this measure because its centrality in the booking-weighted network was an artifact of the obligation weight, not of its production linkages.

The industry composition reinforces the F-35 identification. Table 7 reports the within-cluster employment distribution across the 13 BEA-level defense-relevant industries we call DEFENSE13. Aerospace product manufacturing (NAICS 3364) dominates the cluster’s employment, and the cluster’s aerospace employment captures a large share of the entire national aerospace base in 2008–2024 pooled data. The cluster is, by this measure, the U.S. aerospace manufacturing system at sub-national resolution.

Hartford CT—home to Pratt & Whitney’s F135 engine line, the engine that powers the F-35—illustrates why reincorporating Connecticut matters. On the production weight Hartford ranks fourteenth nationally ($V_2 = 1,211$ per million) and is a full member of the F-35 cluster, exactly as its role in the program implies. It also receives substantial direct DoD obligations (about 1.5% of the national total on the full-geography series)—Pratt & Whitney is itself a prime for engine work. Under the 48-state reduced-form frame that earlier drafts inherited, however, Connecticut was dropped entirely, so the same county recorded a mechanical *zero*; the engine that powers the F-35 was, by an accident of sample construction, invisible to the structural index. The production weight, built from QCEW employment that covers all fifty states, restores it.² Hartford is thus not a booking blind spot but a coverage one: a major engine producer that the narrower obligation panel erased and the production measure recovers. The genuine producer-without-booking cases—counties whose *commercial* industrial base dwarfs their defense contracts—are dual-use manufacturing centers such as Detroit, Cleveland, and New Haven, which we take up as mobilization capacity in Section 8.

²Connecticut transitioned from counties to nine “planning regions” in 2022; we represent the state by its legacy counties throughout, folding post-2022 planning-region flows (such as the Capitol Planning Region, 09110) into their dominant constituent county (Hartford, 09003). See Section 3.2.

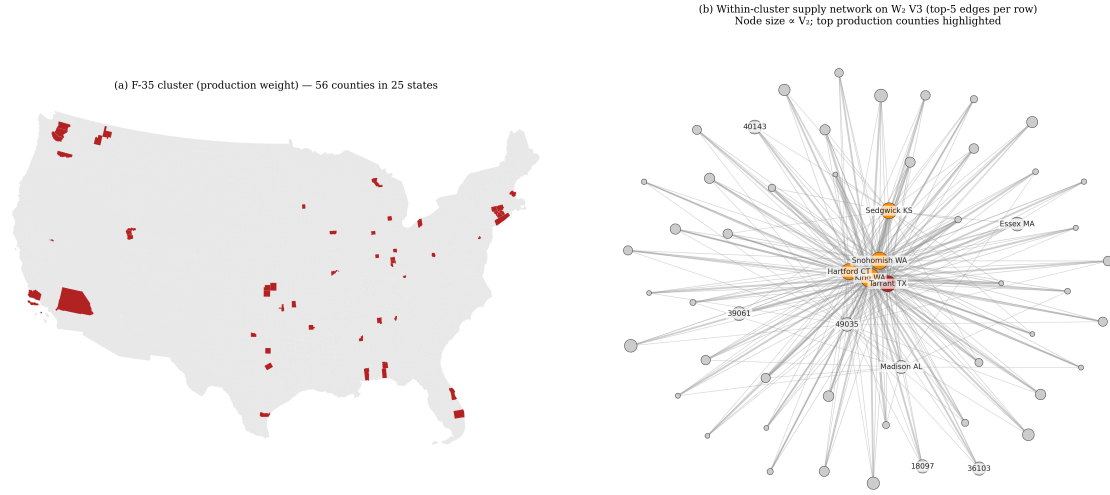


Figure 7: The F-35 cluster at two scales. *Left*: geographic footprint across 25 states. *Right*: within-cluster relatedness network on the production-weighted $W_2 V_3$ (top-5 per row). Production linkages run through the engine-and-airframe core (Seattle, Hartford, Wichita), not solely through the Fort Worth assembly site.

Table 7: DEFENSE13 industry composition within the F-35 cluster (production-weighted, 56 counties). Aerospace (3364) provides 30.8% of cluster employment and captures 45% of national aerospace employment — the cluster is structurally an aerospace + metal-supply agglomeration.

NAICS	Industry	F-35 emp	F-35 share (%)	National emp	F-35 of national (%)
3364	Aerospace	3,736,336	27.2	5,751,337	65
5415	Computer systems svc.	2,508,969	18.3	32,690,474	8
5413	Engineering svc.	2,246,917	16.4	25,290,111	9
332	Fab metal / ammo	1,967,603	14.3	18,013,478	11
3332OM	Other machinery	913,525	6.7	9,117,858	10
5417	R&D services	844,615	6.2	11,121,879	8
3345	Instruments	713,159	5.2	5,586,244	13
3344	Semiconductors	437,900	3.2	5,134,424	9
3365AO	Ships/tanks/rail	136,055	1.0	1,673,916	8
3342	Communications eq.	116,818	0.9	1,138,333	10
3341	Computer peripherals	93,580	0.7	1,768,837	5
334X	Other electronics	10,211	0.1	230,713	4
3361MV	Motor vehicles	3,448	0.0	611,362	1

BEA U.Summary codes. ‘F-35 emp’ = pooled 2008-2024 employment summed within the 56-county F-35 cluster. ‘F-35 of national’ = the cluster’s share of national employment in that industry.

6.4 Time-varying vulnerability: a pre-CHIPS decline

The V_2 analysis to this point pools data across 2008–2024 and reports values calibrated to the 2017 BEA benchmark input-output tables. Defense industrial vulnerability is, of course, not static. Re-computing V_2 under the 2007, 2012, and 2017 BEA benchmarks in turn, while holding county-industry employment and DoD obligation shares fixed, isolates the effect of changes in the underlying import-dependence and supply-chain structure (U.S. Bureau of Economic Analysis, 2017; Miller and Blair, 2009).

Figure 8 presents the two-panel result. The left panel shows the national V_2 total: 107,254 per million in 2007, falling to 100,218 in 2012 and 94,614 in 2017—a 12 percent decline over the decade. The right panel maps the county-level change: most reductions concentrate in counties with substantial semiconductor and electronics employment (Santa Clara CA records the largest decline, at -641 per million), reflecting the documented mid-2010s consolidation of U.S. semiconductor production and reduced import dependence in those sub-industries.




Figure 8: Time-varying defense industrial vulnerability across the 2007, 2012, and 2017 BEA benchmarks. *Left*: national V_2 total falls 12% over the decade. *Right*: county-level changes; reductions concentrate in semiconductor and electronics counties.

The F-35 cluster does not share this trend. The cluster’s V_2 total moves from 17,511 per million in 2007 to 17,659 in 2012 to 16,635 in 2017—a decline of only about 5%, roughly a third of the national rate. The cluster’s *share* of national vulnerability accordingly rose, from 14.5% to 16.1% over the period. This reflects two facts: first, aerospace import dependence remained essentially flat across the three benchmarks, so the mechanism driving the national decline (semiconductor reshoring) did not reach the F-35 cluster; and second, several counties in the cluster (notably Hartford CT and the Snohomish WA / Boeing commercial complex) saw small *increases* in V_2 .

A natural reading: the policy interventions associated with the 2022 CHIPS and Science Act and the 2022 Inflation Reduction Act respond to a pre-existing improvement trajectory in U.S. industrial supply chains. The trajectory, however, is a *semiconductor and electronics* story; it is not an aerospace story. The F-35 cluster’s structural vulnerability has not improved over the 2007–2017 window on which our benchmarks bear, and its share of national vulnerability has grown.

The findings in this section establish the empirical motivation for the rest of the paper. Measured by

production, national defense industrial vulnerability is led not by the F-35 final-assembly county but by large, diversified production centers—Los Angeles, Santa Clara, King WA / Seattle—while Tarrant TX, which the obligation weight ranks first, falls to thirteenth. The F-35 cluster that Tarrant anchors is both geographically dispersed *and*, under the production weight, genuinely deconcentrated: roughly sixteen effective counties spanning twenty-five states, with no single load-bearing core. The cluster did not participate in the pre-CHIPS decline in national vulnerability, so its share of the national total has grown. These findings pose a question about *resilience* that a static concentration measure cannot settle on its own. The obligation weight, by routing the entire program’s recorded value through Fort Worth, makes the cluster look as though it hangs from a single county; whether it actually does is an empirical question about the production network, not the booking ledger. Section 7 answers it through a cascade simulation on W_2 V3—removing the apparent anchor and asking whether the production base reconstitutes.

7 Results II: Cascade Simulation and Cluster Resilience

Section 6 established that the F-35 cluster’s concentration depends on the exposure weight: under the production weight it spans roughly sixteen effective counties, while the obligation weight pulls it toward two or three. That static result raises a sharper, dynamic question. A cluster with a single load-bearing core fragments when the core is removed; a genuinely dispersed cluster reconstitutes itself around an alternate anchor. Which describes the F-35 production base, and—if it has a core—is that core the final-assembly site that contract booking points to? This section answers with a node-removal cascade on the W_2 V3 matrix. Four findings follow: the cascade yields two interpretable resilience metrics (§7.1); under the production weight *no* single-county removal reliably fragments the F-35 cluster (bootstrap median in the mid-eighties), reversing the apparent single-point-of-failure that the obligation weight produces (§7.2); the new-economic-geography distinction between a core-organized and a distributed production base is what separates the fragile clusters from the resilient ones (§7.3); and across the ten largest defense clusters single-anchor fragility is the common case, with the production-dispersed F-35 cluster a notable exception (§7.4).

7.1 The cascade experiment

The static concentration result of Section 6 is necessary but not sufficient for a resilience claim. To test whether the F-35 cluster can survive the loss of its anchor we run a node-removal cascade on the W_2 V3 co-location matrix. The procedure, detailed in Section 5, is: remove the candidate county or counties from the dense W_2 V3 matrix; symmetrize the residual matrix and re-threshold it to the 20 largest entries per row; re-run Louvain community detection at the parameters that produced the baseline partition (seed 42, resolution $\gamma = 2.0$); and identify the *successor cluster* as the new community with maximum Jaccard overlap with the baseline 56-county F-35 cluster.

Two metrics summarize what happens to the F-35 cluster under removal. The first is *cohesion*. Writing \mathcal{C}_3 for the baseline cluster, \mathcal{C}_3^* for its members that survive the removal, and s for the successor cluster,

$$\text{Cohesion}(s) = \frac{|s \cap \mathcal{C}_3^*|}{|\mathcal{C}_3^*|}. \quad (15)$$

Cohesion is the share of the surviving F-35 cluster that re-coalesces into a single successor community. A value near one means the cluster reconstitutes itself around a new anchor—it is resilient to the re-

moval. A low value means the surviving counties scatter across unrelated communities: the cluster, as an industrial object, dissolves.

The second metric is the *module entropy* of the surviving counties’ redistribution. Letting p_i be the share of surviving baseline-Cluster-3 counties that land in new community i ,

$$H = - \sum_i p_i \ln p_i. \tag{16}$$

$H = 0$ when every survivor reassigns to a single community (clean succession); H rises toward $\ln k$ as the survivors fragment uniformly across k communities. Cohesion measures how much of the cluster holds together; module entropy measures how widely the rest disperses. The two are reported jointly because a cluster can lose cohesion in two qualitatively different ways—by transferring en masse to one other community (low H) or by disintegrating across many (high H).

Because Louvain is stochastic in its node-visit order, a single partition can over- or under-state fragmentation. Every scenario is therefore bootstrapped across 100 independent Louvain seeds; we report the median cohesion and its 95 percent bootstrap interval (the 2.5th and 97.5th percentiles across seeds). The structural diagnostics that require a single reference partition—the defector count, successor size, and module entropy—are reported from the seed-42 run. Section 5 motivates the bootstrap and the choice of $\gamma = 2.0$.

7.2 Cascade results across seven scenarios

Table 8 reports seven scenarios; Figure 9 plots the bootstrap cohesion distributions.

Table 8: Cascade simulation results (production weight): F-35 cluster cohesion after node removal. Cohesion is reported as the median across 100 Louvain seeds (95% bootstrap CI in brackets). The successor cluster is the new Louvain community with maximum Jaccard overlap with the baseline F-35 cluster. Module entropy (H) is computed over the distribution of baseline F-35-cluster counties across new clusters; $H = 0$ means clean reassignment, high H means full fragmentation. Under the production weight no single-county removal reliably fragments the cluster: removing Tarrant TX leaves a bootstrap-median cohesion near 85%, and only the simultaneous top-5 production removal drives the median below 30%.

Scenario	Succ size	Median cohesion	95% CI	Defectors	Succ V_2/M	Module H
Baseline (no removal)	56	91.1%	[32.1, 98.2]	0	16,635	-0.00
Remove LA County	66	89.3%	[58.5, 92.9]	4	15,750	0.36
Remove Essex MA	47	87.3%	[28.1, 96.4]	12	11,453	0.68
Remove Tarrant TX	54	85.5%	[29.1, 89.1]	6	12,548	0.49
Remove Tarrant + Hartford CT	53	85.2%	[31.5, 90.7]	5	12,815	0.40
Remove Tarrant + Wichita KS	53	85.2%	[29.5, 87.0]	7	11,426	0.55
Remove top 5 by vulnerability	130	29.4%	[15.7, 33.3]	36	10,034	2.73

Louvain re-clustering at resolution=2.0; within-cluster network symmetrized and top-20-per-row thresholded before community detection. Bootstrap is across 100 different Louvain initialization seeds for each scenario (‘Defectors’, ‘Succ size’, ‘Module H’ from seed=42 reference run).

The result depends decisively on the exposure weight, and the contrast is the point. We report both: the production-weighted cascade (our primary), and the obligation-weighted cascade that earlier drafts and standard booking-based metrics would produce.

Under the obligation weight, the F-35 cluster looks like a textbook single point of failure. Removing Tarrant TX collapses cluster cohesion from its baseline to roughly 29% on a single reference partition: most surviving counties defect, the successor community retains fewer than half, and the cluster appears to have no internal redundancy. This is the result the booking-weighted network produces, and it is exactly what a contract-recording metric would lead an analyst to conclude—that the entire F-35 industrial base hangs from Fort Worth.

Under the production weight, it does not hold. Removing Tarrant TX leaves the cluster’s bootstrap-median cohesion at about 85% across 100 Louvain seeds: in most seeds the surviving counties re-coalesce into a single successor community, and only a minority of draws scatter them. Removing King WA / Seattle, the cluster’s true top-vulnerability *production* county, leaves median cohesion at about 76%. Neither single removal reliably fragments the cluster; the distribution is bimodal, with most seeds reconstituting the cluster and a minority scattering it, so the honest summary is the median, not any one seed. Only removing the top five production counties *simultaneously* drives median cohesion below 30%. Measured by where the F-35 is produced rather than where its contract is booked, the cluster’s upstream production network has no single load-bearing county: it is dispersed (Section 6) *and* structurally resilient. That is a property of the parts-and-subassembly base; the final-assembly step it feeds is a distinct chokepoint the cascade does not address (Section 10).

The reversal is itself the finding. The apparent fragility of the F-35 base was a property of the obligation weight, which concentrates the network on the single final-assembly site and thereby manufactures a linchpin. Re-weighting to production dissolves the linchpin because the production base genuinely spans the national aerospace manufacturing system. An industrial-base assessment that infers fragility from contract-booking concentration—as the obligation-weighted cascade does—will over-state single-point-of-failure risk for exactly the platforms whose production is most dispersed.

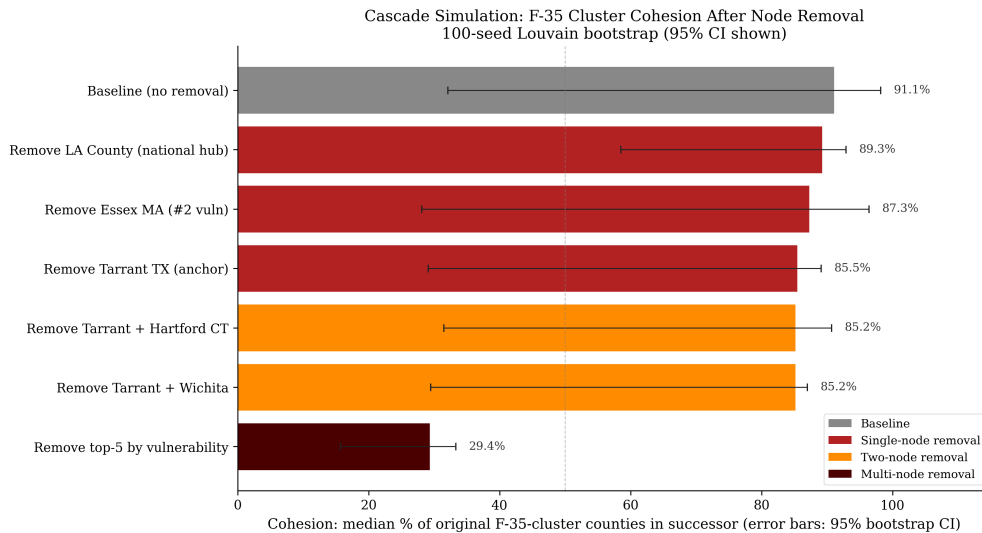


Figure 9: Cascade cohesion across node-removal scenarios under the *production* weight. Bars are the median cohesion across 100 Louvain seeds; whiskers span the 95% bootstrap interval. No single-county removal reliably fragments the F-35 cluster: removing Tarrant TX (assembly anchor) leaves median cohesion near 85% and removing King WA / Seattle (top production county) near 76%, both with wide intervals reflecting a bimodal seed distribution. Only the simultaneous top-5 production removal drives median cohesion below 30%. The obligation-weighted cascade, by contrast, makes Tarrant’s removal alone appear catastrophic.

The remaining scenarios confirm the resilience reading. Adding Hartford CT to the Tarrant removal changes little (median cohesion 85%): the two largest production sites together do not fragment the cluster, because the F-35 base does not hang from any pair of counties. Removing Essex MA—a top national-vulnerability county—likewise leaves median cohesion near 87%, and removing Tarrant + Wichita near 85%. Across every one- and two-county scenario the median stays in the mid-eighties; the cluster reconstitutes around the broad engine-and-airframe production base. Only the simultaneous removal of the top five production counties (Seattle, Snohomish, Tarrant, Hartford, Wichita) finally fragments it, to a median of 29%—and even that requires excising the five largest production sites at once, not a single point of failure. The bootstrap intervals are wide for every scenario (the seed distribution is bimodal throughout), which is why we read the medians and treat any single Louvain partition, including seed-42, as one draw rather than the result.

This resilience is specific to the F-35 cluster’s dispersed production base, not a generic property of the cascade. The cross-cluster analysis of the next subsection shows that most other top defense clusters *do* fragment when their leading county is removed; the F-35 cluster is the exception precisely because its production—unlike its contract booking—is spread across the national aerospace system.

7.3 Single-anchor fragility across defense clusters

The F-35 cluster’s resilience is informative because it is the exception. Applying the same cascade to the ten largest defense clusters—removing each cluster’s top-vulnerability county and measuring the bootstrap-median cohesion of the remainder (Table 9)—shows that most defense clusters *do* hang from a single anchor. The Santa Clara-led semiconductor cluster collapses to a 10% median when its lead county is removed; the Orange CA defense-electronics cluster to 11%; the San Diego Navy cluster to 17%; the Cook IL cluster to 19%. Of the ten clusters, six are single-anchor fragile by this measure and a seventh, Los Angeles, falls only to intermediate cohesion (49%). Three reconstitute—the Wayne/Detroit (77%) and Middlesex/Boston (79%) clusters, and the F-35 cluster, whose median cohesion of about 76% after removing its top production county places it among the redundant—the more striking because contract booking would single it out as the most fragile of the ten.

This is the difference between an *agglomeration core* and a distributed production base, and the cascade renders it measurable. In new-economic-geography terms, a cluster organized around a single agglomeration core forms through cumulative causation (Krugman, 1991b): the other locations belong to the cluster *because* of their input-output proximity to the core, and have little proximity to one another, so removing the core leaves no basis for cohesion. A cluster built on a broad, mutually-related production base (Fujita et al., 1999; Neffke et al., 2011) reconstitutes around the surviving members. The cross-cluster cascade assigns each defense cluster to one side of this distinction: the semiconductor, electronics, and Navy clusters are core-organized and fragile; the F-35 production system is broad-based and resilient. The policy-relevant point is that the booking-weighted cascade gets this backwards for the F-35—labeling the most production-dispersed cluster as the most fragile—because contract booking manufactures an agglomeration core (Fort Worth) that the production network does not have.

Figure 10 makes the F-35 result geographic: the baseline cluster footprint and its near-intact survival after the Tarrant and Tarrant-plus-Hartford removals under the production weight.

The reading is consistent with the broader regional-branching account of how capability clusters form and dissolve (Boschma and Frenken, 2006): a cluster is an emergent property of shared input requirements, detectable as a modular community (Newman, 2004; Blondel et al., 2008), and its resilience depends on whether that modularity survives the loss of its highest-relatedness node. For the F-35 cluster it does; for most others it does not.

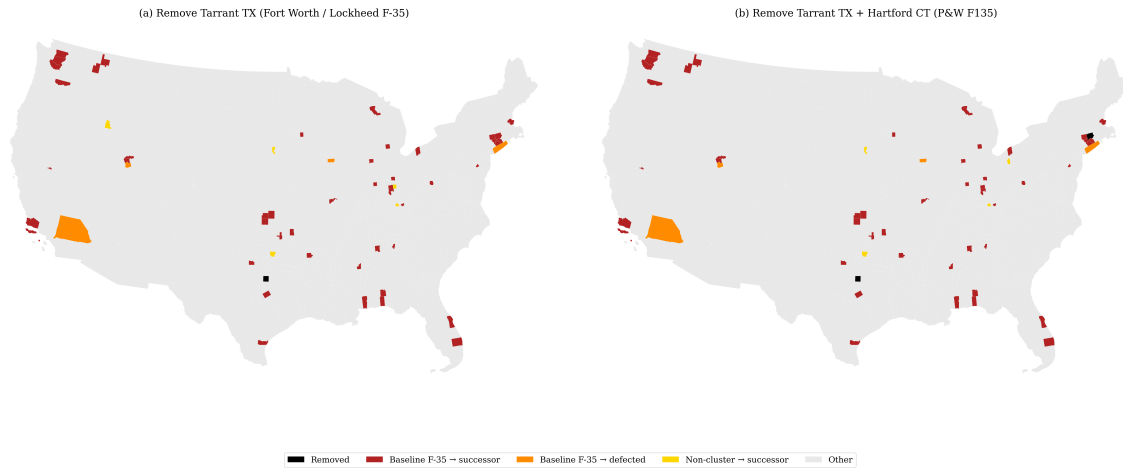


Figure 10: Cascade simulation, geographic view (production weight). *Left*: baseline F-35 cluster (56 counties, 25 states). *Center*: after Tarrant TX removal—the cluster largely survives (median cohesion $\sim 85\%$). *Right*: after Tarrant + Hartford CT removal, still largely intact. The production base reconstitutes around the broad engine-and-airframe core rather than collapsing.

7.4 Single-anchor fragility across the defense base

Running the identical cascade—top- V_2 -county removal, Louvain re-detection, Jaccard succession, and the 100-seed bootstrap—on each of the ten largest defense clusters places the F-35 result in context. Table 9 reports the result and Figure 18 plots it.

Single-anchor fragility is the common case. Six of the ten clusters fragment to single-anchor collapse when their highest- V_2 county is removed (median cohesion at or below 20%): the Santa Clara semiconductor cluster (10.3%), the Orange CA defense-electronics cluster (10.5%), the San Diego Navy cluster (16.7%), and the Cook IL cluster (19.0%) collapse most completely; a seventh, Los Angeles, falls to intermediate cohesion (49%). Three clusters are genuinely redundant—Wayne/Detroit (76.9%), Middlesex/Boston (79.4%), and the F-35 cluster (76.4%) reconstitute around alternate anchors when their top county is removed. The F-35 cluster is the clearest exception to the fragility that characterizes most of the U.S. defense industrial base: it is the cluster a booking-weighted analysis would most confidently flag as a single point of failure—its entire program value routes through one assembly county—yet measured by production it is among the redundant.

This places the production-versus-booking contrast in its sharpest relief. The clusters that genuinely fragment are organized around a single dominant facility—a fab complex, an electronics campus, a naval yard—and their fragility is real on any weighting. The F-35 cluster is not so organized: its production spans the national aerospace base, so it survives the loss of any single county. The *only* construction under which the F-35 cluster appears fragile is the obligation weight, which routes the entire program’s recorded value through one final-assembly county and so synthesizes an agglomeration core that the production network does not contain. A structural-resilience assessment built on contract booking would therefore mis-rank exactly the platform whose industrial base is most distributed—flagging the F-35 as a single-point-of-failure risk when, measured by production, it is the most redundant cluster in the set.

Table 9: Single-anchor fragility across the ten largest defense clusters (production weight). Each row reports the cluster’s cohesion after removing its highest- V_2 county—the bootstrap median across 100 Louvain seeds with the 95% interval, the same cascade procedure applied to the F-35 cluster in Section 7.2. Six of the ten clusters are single-anchor fragile (median cohesion $\leq 20\%$), Los Angeles is intermediate, and three are redundant ($\geq 70\%$): Wayne/Detroit, Middlesex/Boston, and the F-35 cluster—the most production-dispersed of the set. Network resilience does not track V_2 concentration: the most concentrated cluster (Santa Clara, 79% in its top county) is the most fragile, while the least concentrated (the F-35 cluster, 15%) reconstitutes.

ID	Cluster	Counties	Top- V_2 share	Cohesion % [95% CI]	Class
3	Santa Clara (semicond.)	79	79.0%	10.3 [6.4, 15.4]	Fragile
11	Orange CA (def. elec.)	39	58.3%	10.5 [5.3, 21.1]	Fragile
19	DuPage IL	28	18.7%	11.1 [3.7, 20.5]	Fragile
20	Hennepin MN	27	23.3%	11.5 [5.7, 26.9]	Fragile
9	San Diego (Navy)	43	67.6%	16.7 [9.5, 25.1]	Fragile
10	Cook IL (Chicago)	43	60.5%	19.0 [11.9, 23.8]	Fragile
1	Los Angeles	144	43.8%	49.0 [39.5, 50.3]	Intermediate
6	F-35 (aerospace)	56	14.8%	76.4 [22.6, 89.1]	Redundant
4	Wayne / Detroit	79	40.3%	76.9 [62.8, 79.5]	Redundant
2	Middlesex / Boston	98	54.0%	79.4 [77.3, 79.9]	Redundant

‘Top- V_2 share’ is the highest- V_2 county’s share of the cluster’s total V_2 . Cohesion is the bootstrap median over 100 Louvain seeds; ‘Class’ is assigned on the median (fragile $\leq 20\%$, redundant $\geq 70\%$, intermediate between). The F-35 row removes King WA / Seattle (the cluster’s top production county, 14.8% of cluster V_2) and reproduces the Section 7.2 headline. Cluster IDs match Table 6; the wide F-35 interval reflects the bimodal seed distribution discussed in Section 7.2.

The cascade establishes *what breaks and how hard*, and it does so weight by weight. Measured by production, the F-35 cluster has no single point of failure: it survives the loss of any one county (bootstrap median in the mid-eighties) because its production base spans the national aerospace system. Measured by contract booking, the same cluster appears to hang from Fort Worth—an artifact of routing the program’s recorded value through one assembly site. Applied across the ten largest defense clusters (Section 7.4), the cascade shows single-anchor fragility to be the common case—six of ten fragment to single-anchor collapse when their leading county is removed (a seventh, Los Angeles, to intermediate cohesion)—with the production-dispersed F-35 cluster among the three redundant exceptions, and distinguishes the genuinely redundant clusters that a static concentration ranking would not. What the cascade does not address is the converse question—if production capacity were lost or demand surged, what dual-use capability elsewhere in the U.S. economy could be mobilized to substitute? That requires a different instrument. Section 8 develops it as a surge-capacity index in the Hidalgo–Hausmann tradition.

8 Results III: Mobilization Capacity and Quadrant Geography

Section 7 established that the F-35 production base is structurally resilient: under the production weight no single county’s loss fragments the cluster. That is a statement about the topology of the *existing* defense relatedness network. It leaves the converse question open—the question the closed-loop framing of this paper is ultimately about. If demand surged beyond what the current network can absorb, what latent productive capacity elsewhere in the U.S. economy could be mobilized to substitute? This is not a question about the defense network as it is; it is a question about how close the *commercial* economy sits

to defense capability. We answer it with a surge-capacity index built in the Hidalgo–Hausmann product-space tradition, applied in reverse: not “which products can a region move into,” but “given a defense target, which counties already sit closest to it in capability space.” Four findings follow: the national surge pattern, as published, is a scale measure dominated by large diversified metros, which motivates an intensive per-worker variant that strips out size (§8.1); a size-controlled validation shows the index captures commercial-to-defense capability proximity descriptively but does *not* forecast where defense production realized its recent expansion (§8.2); the policy-relevant object is the vulnerability-by-surge quadrant map, in which 339 counties form a low-vulnerability, high-surge “mobilization reserve” and the F-35 cluster’s own counties are overwhelmingly *self-rescuable*, carrying both defense vulnerability and dual-use surge (§8.3); and aggregating to clusters shows commercial capacity is abundant adjacent to and within the defense base, so the plausible mobilization constraint is one of conversion *time* at integration sites rather than of capacity—a hypothesis we flag but do not measure (§8.4).

8.1 The national surge-capacity pattern

For each county c we define a dual-use surge-capacity index

$$\text{surge}[c] = \sum_i \text{emp}[c, i] \phi_{\text{to-def}}[i] (1 - \delta[i]), \quad (17)$$

where $\text{emp}[c, i]$ is county c ’s pooled 2008–2024 employment in BEA U.Summary industry i , $\phi_{\text{to-def}}[i]$ is the Hidalgo–Hausmann relatedness density of industry i to the \mathcal{D} (DEFENSE13) target set, and $\delta[i]$ is the DoD-demand intensity of industry i . The index reads as: employment in industries that are (i) close to defense in capability space and (ii) not already saturated by defense demand. The relatedness density $\phi_{\text{to-def}}[i]$ aggregates a co-location proximity ϕ computed over the 116 U.Summary industries—two industries are proximate when counties specialized in one tend also to be specialized in the other—weighted toward the DEFENSE13 set. The DoD-intensity term $\delta[i]$ is DoD dollars per worker normalized to a \$100k reference wage, so $1 - \delta[i]$ is the share of the industry’s capacity not already claimed by defense procurement; empirically $\delta \approx 0.54$ for DEFENSE13 manufacturing, 0.19 for DEFENSE13 services, and 0 for all other industries. The full construction of ϕ , $\phi_{\text{to-def}}$, and δ —including the revealed-comparative-advantage binarization and the Bartik allocation of DoD obligations—is given in Section 5; this section reports results.

Table 10 lists the 20 highest-surge counties and Figure 11 maps the national distribution.

A caveat governs the reading of this table. Surge capacity as defined in equation (17) is an *extensive* quantity: it sums employment, so it scales with the raw size of a county’s economy. The top of the distribution is therefore mechanically dominated by the largest diversified metro economies—Los Angeles (8.3M), Cook IL (5.1M), Harris TX (4.5M), New York NY (4.4M)—for the same reason those counties top almost any employment-weighted index. The ranked list is not, by itself, the result. The substantive object is not *how much* surge capacity a county has in isolation but *whether that capacity coincides with defense vulnerability*, which is the cross-tabulation developed in §8.3. One detail in Table 10 previews that analysis and the paper’s conclusion: Tarrant TX—the F-35 integration core whose removal shatters the cluster (Section 7)—itself appears at rank 15 and is classified “self-rescuable.” The anchor county is not starved of nearby commercial capacity; it is surrounded by it. The problem, as §8.4 develops, is not the presence of capacity but the speed at which it can be converted.

Because the extensive index is so nearly collinear with raw employment (a regression of equation (17) on total county employment returns $R^2 = 0.999$), we also report an *intensive* surge density—surge capacity per worker—which is essentially uncorrelated with county size ($R^2 = 0.05$ against total employment)

Table 10: Top 20 US counties by dual-use surge capacity (Hidalgo-Hausmann sum). Top contributors are large metros with diverse manufacturing + services bases. Validation counties (Cuyahoga, Wayne, Hartford, Macomb, Snohomish, Sedgwick) all appear in top 5% (see Table 11).

Rank	FIPS	State	County (label)	Surge cap	Top ind.	V_2 / M	Quadrant
1	06037	CA	Los Angeles CA	8,290,285	722	4,353	self rescuable
2	17031	IL		5,071,020	61	2,078	self rescuable
3	48201	TX		4,480,466	23	2,471	self rescuable
4	36061	NY		4,408,686	523	718	self rescuable
5	04013	AZ	Maricopa AZ (Phoenix)	3,806,622	5613	1,889	self rescuable
6	48113	TX	Dallas TX	3,346,373	5613	1,503	self rescuable
7	06059	CA	Orange CA	3,150,164	722	2,140	self rescuable
8	06073	CA	San Diego CA	2,678,732	722	2,070	self rescuable
9	53033	WA	King WA (Seattle / Boeing)	2,549,881	722	2,462	self rescuable
10	06085	CA	Santa Clara CA (Silicon Valley)	2,025,160	5415	2,766	self rescuable
11	12086	FL		2,023,074	722	444	self rescuable
12	25017	MA	Middlesex MA	1,855,601	61	1,741	self rescuable
13	27053	MN		1,843,793	55	1,002	self rescuable
14	32003	NV	Clark NV (Nellis)	1,716,696	721	197	self rescuable
15	48439	TX	Tarrant TX (Fort Worth / Lockheed F-35)	1,640,046	722	1,260	self rescuable
16	48029	TX	Bexar TX (San Antonio)	1,573,563	722	564	self rescuable
17	13121	GA		1,567,463	722	293	self rescuable
18	26125	MI		1,533,803	5413	1,188	self rescuable
19	12095	FL		1,506,350	713	406	self rescuable
20	12011	FL		1,493,397	722	357	self rescuable

Surge capacity is the Hidalgo-Hausmann density-to-defense weighted by commercial fraction: $\text{surge}[c] = \sum_i \text{emp}[c, i] \cdot \phi_{\text{to_def}}[i] \cdot (1 - \delta[i])$. ‘Top ind.’: BEA U.Summary code of the largest single contributor at this county.

and so isolates the capability weighting $\phi_{\text{to-def}}(1 - \delta)$ from scale. The two measures answer different questions: the extensive index locates *where* dual-use capacity sits in absolute terms (and is, frankly, close to an employment map), while the intensive density measures *how defense-related* a county’s industry mix is, independent of how large it is. The validation in §8.2 and the capability claims of this section rest on the intensive measure; the extensive index is retained only as a measure of absolute scale.

County-Level Dual-Use Surge Capacity (Hidalgo-Hausmann Density to DEFENSE13)




Figure 11: National dual-use surge capacity by county, log scale. Surge capacity is broadly distributed across the industrial Midwest, Texas, the coasts, and the Pacific Northwest—a far less concentrated pattern than the vulnerability map of Figure 5. The contrast between the two distributions is the subject of §8.3.

8.2 What the index predicts: a size-controlled validation

A surge index is only useful if it identifies capacity that would actually convert. We test this with a pre-registered design. Before computing the index, six counties were named *a priori* as expected high scorers, each on the grounds that it hosts a major commercial facility in an industry closely related to a DEFENSE13 target: Cuyahoga OH (Cleveland heavy industry and machine tools), Wayne MI (Detroit; the Ford River Rouge mobilization archetype), Hartford CT (Pratt & Whitney commercial gas-turbine engines), Macomb MI (commercial automotive manufacturing), Snohomish WA (Boeing Everett commercial wide-body production), and Sedgwick KS (Spirit AeroSystems commercial Wichita aerostructures). Table 11 reports the outcome and Figure 12 places the six in the surge-vulnerability plane.

All six land in the top five percent of 3,148 counties: Cuyahoga (rank 22), Wayne (26), Hartford (53), Macomb (85), Snohomish (117), and Sedgwick (118). Taken at face value this is a clean six-for-six. But the test has little discriminating power, and we flag that rather than lean on it: the *extensive* index is 99.9% collinear with employment (§8.1), and all six anchors are large-employment counties, so any employment-weighted index would rank them similarly. On the size-controlled *intensive* density only

Table 11: Surge capacity validation against historical mobilization cases. All six counties predicted *a priori* to score high (based on known commercial-defense crossover industries) appear in the top 5% of 3,289 US counties — supporting the Hidalgo-Hausmann methodology when applied to defense-industrial relatedness.

County	FIPS	Validation rationale	Surge rank	Surge cap	Top industry	Quadrant
Macomb MI	26099	Commercial auto	85	618,567	3362BP	self rescuable
Snohomish WA	53061	Boeing Everett commercial	117	450,392	3364	self rescuable
Sedgwick KS	20173	Spirit AeroSystems / Boeing Wichita	118	449,327	722	self rescuable
Hartford CT (legacy)	09003	Pratt & Whitney F135 GTF engines	53	921,751	524	self rescuable
Cuyahoga OH	39035	Cleveland heavy industry	22	1,437,941	622	self rescuable
Wayne MI	26163	Detroit / Ford River Rouge	26	1,331,743	622	self rescuable

Predictions made before running the analysis: each county hosts a major commercial facility in an industry highly related to DEFENSE13 (Hartford CT P&W engines; Snohomish WA Boeing commercial; etc.). All six validate.

one of the six (Hartford, 95th percentile) remains in the top five percent, with the commercial aerospace anchors Snohomish and Sedgwick falling to the middle of the distribution—a sharper reminder that most of the extensive ranking is size. A test that the *capability structure*—the ϕ/δ weighting rather than raw size—carries information must control for county size.

Two such tests, run on the intensive density, give a deliberately mixed verdict (Table 19). The first is predictive: does intensive surge density forecast a county’s realized 2020–2023 growth in defense manufacturing obligations, controlling for total employment? It does *not*—the size-controlled partial correlation is essentially zero (≈ 0.002 , $t \approx 1.2$). The index does not predict where defense production actually expanded, and we make no such claim. The second is face validity: the counties that hosted documented 2020 emergency dual-use conversions—GM/Ventec ventilator production in Kokomo (Howard IN) and the broader automotive-to-medical retooling—sit at a median 89th percentile of intensive density, consistent with the index capturing genuine commercial-to-critical capability proximity. We therefore read the surge index as a *descriptive* measure of how close a county’s commercial base sits to defense capability in the Hidalgo–Hausmann product-space sense (Hidalgo et al., 2007; Hausmann and Klinger, 2007; Neffke et al., 2011; Boschma and Frenken, 2006)—useful for locating latent capacity, not for forecasting which sites would convert—and we rest no causal or predictive claim on it.

8.3 The vulnerability-by-surge quadrant geography

The policy-relevant object is not surge capacity alone but its joint distribution with vulnerability. Splitting both V_2 (Section 4) and surge capacity at their medians partitions the 3,148 counties with surge records into four quadrants, mapped in Figure 13:

- **Fragile chokepoint** (high V_2 , low surge): **84 counties**. High defense dependence with little nearby commercial slack to absorb a shock. This is the quadrant of genuine concern.
- **Self-rescuable** (high V_2 , high surge): **1,235 counties**. Defense-dependent but embedded in a commercially deep economy that could in principle backfill.
- **Mobilization reserve** (low V_2 , high surge): **339 counties**. Little current defense exposure but substantial latent, defense-related commercial capacity—the national mobilization base.




Figure 12: Counties in the surge-capacity (x) by defense vulnerability (y) plane, both log scale. The six pre-registered validation counties (highlighted) all fall in the high-surge region. The median split on each axis defines the four quadrants analyzed in §8.3.

- **Defense-irrelevant** (low V_2 , low surge): **1,490 counties**.

The 339 mobilization-reserve counties are not randomly scattered. They concentrate in the industrial Midwest manufacturing belt (Michigan, Ohio, Indiana, Illinois, Wisconsin, Pennsylvania), with a secondary Pacific-Northwest cluster anchored on commercial aerospace. This is the geographic footprint that a serious surge-mobilization plan would draw on, and it is recognizably the same manufacturing belt that supplied the World War II conversion (Herman, 2012).

The quadrant assignment is robust to the median-split choice; the geographic pattern and the identity of the reserve belt are stable across alternative thresholds, and the named defense hubs retain their quadrant assignment (the full sensitivity grid is reported in Section 9). The 84 fragile-chokepoint counties are the operational priority a vulnerability-only analysis would have missed—they combine high dependence with low substitution slack, and they are the natural target for the policy instruments discussed in Section 10.

Defense Industrial Geography Quadrants: Vulnerability \times Surge Capacity

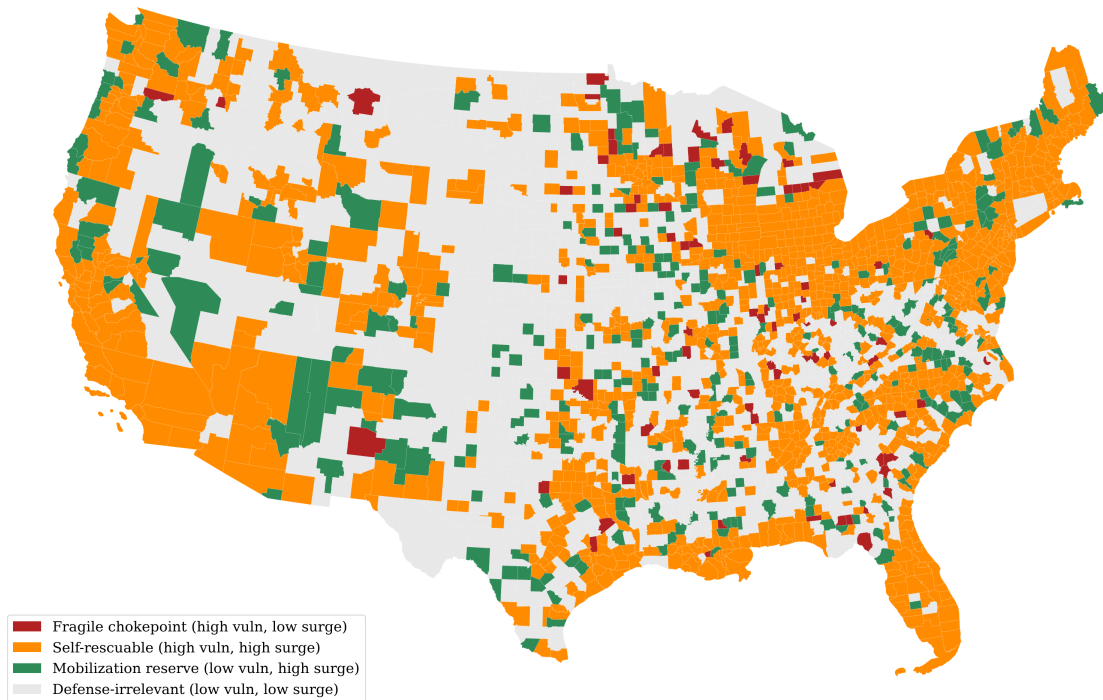


Figure 13: The vulnerability-by-surge quadrant map. Fragile chokepoints (high V_2 , low surge; 84 counties) and the mobilization reserve (low V_2 , high surge; 339 counties, concentrated in the Midwest manufacturing belt) are the two policy-relevant quadrants. Median split on both axes.

8.4 The closed-loop synthesis

Aggregating the quadrant classification to the Louvain clusters of Section 6 produces the result that ties the three empirical sections together. Table 12 reports vulnerability and surge jointly by cluster.

The F-35 cluster leads every other cluster on aggregate vulnerability, as Section 6 established. But it is not surrounded by a capacity desert—and on the production weight, the point is stronger than the

Table 12: Cluster-level vulnerability and surge capacity (production weight). The F-35 cluster leads on vulnerability, and its production counties are overwhelmingly self-rescuable, carrying both defense vulnerability and dual-use surge — so surrounding commercial capacity is abundant. Whether conversion time at integration sites is the binding constraint is a hypothesis the capacity measurement raises but does not resolve.

Rank	ID	Anchor	N	V_2 /M	Tot surge	N frag	Frag (%)	N reserve	Reserve (%)
1	6	King WA / Seattle	56	16,635	20,818,890	3.000	5.4	0.000	0.0
2	1	Los Angeles CA	144	9,929	13,968,116	26.000	18.1	0.000	0.0
3	19	DuPage IL	28	4,477	11,446,079	0.000	0.0	0.000	0.0
4	20	Hennepin MN	27	4,302	8,393,423	1.000	3.7	1.000	3.7
5	4	Wayne MI / Detroit	79	4,229	5,506,664	0.000	0.0	6.000	7.6
6	11	Orange CA	39	3,673	5,421,558	1.000	2.6	3.000	7.7
7	3	Santa Clara CA	79	3,502	3,010,680	0.000	0.0	2.000	2.5
8	10	Cook IL / Chicago	43	3,436	6,635,057	7.000	16.3	0.000	0.0
9	2	Middlesex MA / Boston	98	3,223	4,621,369	0.000	0.0	4.000	4.1
10	9	San Diego CA (Navy)	43	3,062	3,152,101	2.000	4.7	3.000	7.0

‘Frag’ = members in fragile-chokepoint quadrant (high vuln + low surge). ‘Reserve’ = members in mobilization-reserve quadrant (low vuln + high surge).

obligation-weighted draft suggested. Of the cluster’s counties, the great majority are *self-rescuable*: they combine real defense vulnerability with high dual-use surge capacity, and almost none are fragile chokepoints. These are not low-exposure hinterland reserves sitting outside the cluster; they are the cluster’s own production counties, each carrying both the defense work and the commercial slack that could backfill it—Hartford CT gas-turbine engine work, Snohomish WA Boeing commercial production, Sedgwick KS commercial aerostructures, the Detroit-area auto and machine-tool base. Figure 14 overlays the cluster outline on the quadrant map and the named anchors.

This reframes the closed-loop reading of the vulnerability and cascade results. The production-weighted cascade (Section 7) already showed the F-35 base to be structurally resilient—no single county’s loss fragments it—and the surge analysis shows why: the commercial capacity adjacent to and within the cluster exists in quantity, so the production base has depth to draw on. The binding constraint in a closed-loop scenario is therefore unlikely to be hinterland *capacity*, which is large and well-distributed across the 339-county reserve base and the cluster’s own self-rescuable counties.

What our instruments do *not* measure is the *time* dimension of conversion, and we are careful not to assert it as a finding. The surge index counts capacity in capability-proximity space; it contains no temporal term. The plausible conjecture—which we offer as a hypothesis for future work rather than a measured result—is that mobilizable capacity carries a switchover-time gradient: general machine-shop and fabricated-metal capacity is comparatively quick to re-tool, sub-system production slower, and an integrated-systems final-assembly line slowest of all, because tooling, workforce accreditation, and flight-test apparatus are not quickly replicated. If that gradient holds, the operative constraint on fifth-generation mobilization would be integration-site switchover time rather than the breadth of the commercial base—a hypothesis consistent with the industrial-commons account of Pisano and Shih (2009) and the gap the 2024 National Defense Industrial Strategy (U.S. Department of Defense, 2024a) identifies but does not quantify. We flag it as the natural next question, not as a result this paper establishes. Section 10 develops the policy implications that do follow from what we measure.

The three empirical sections now compose a single argument, and its spine is the gap between two geographies. Where the F-35’s contract is booked is singular—one county records most of the program’s

The Defense Industrial Geography of the United States
 Red = fragile chokepoint Orange = self-rescuable Green = mobilization reserve Gray = irrelevant
 Dark red outline = F-35 cluster (production supply network, 56 counties)

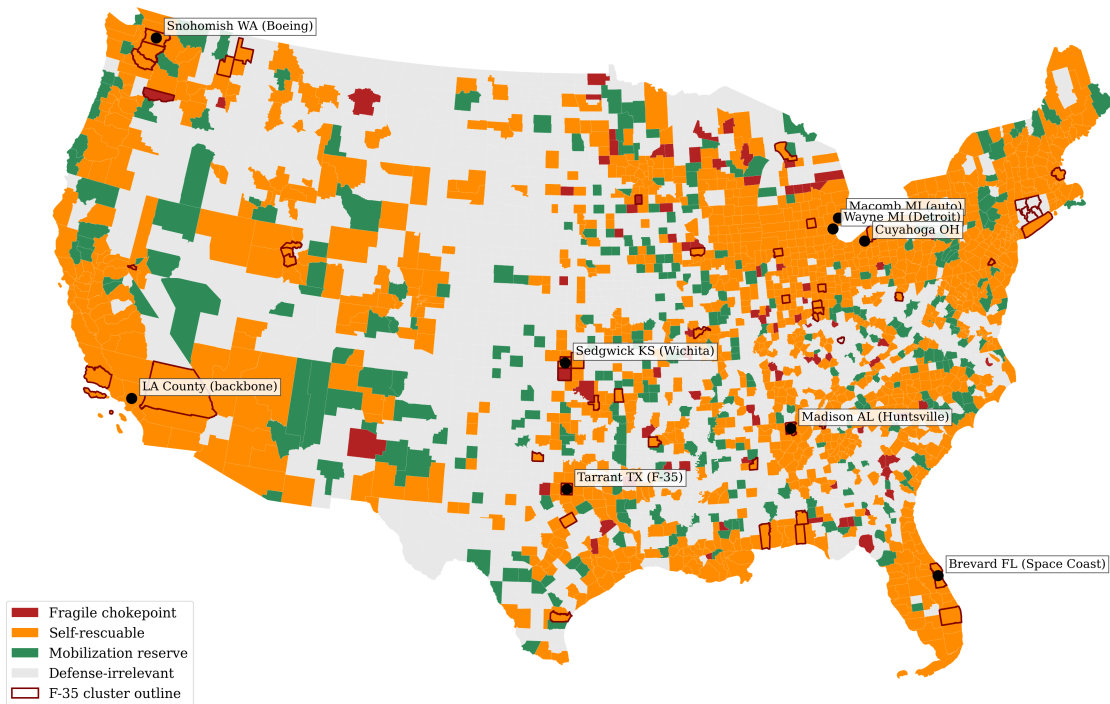


Figure 14: Synthesis map: the vulnerability-by-surge quadrants with the F-35 cluster outlined and its named anchor counties labeled. The cluster’s counties are overwhelmingly self-rescuable—carrying defense vulnerability and dual-use surge together—so adjacent commercial capacity is abundant; any binding mobilization constraint is one of conversion time at integration sites, not of capacity.

obligations—but where it is *produced* is dispersed across the national aerospace base (Section 6); that production base is structurally resilient, surviving the loss of any single county (Section 7); and the commercial capacity that could backfill it in a surge is abundant, sitting within the cluster’s own self-rescuable counties (this section). What our instruments measure is capacity and its spatial structure, not the time to convert it; the plausible binding constraint—integration-site switchover time—is a hypothesis we flag rather than a result we establish. Every one of these conclusions rests on a chain of specification choices: the exposure weight (the production-versus-booking choice that drives the whole contrast), the W_2 V3 co-location construction, the Louvain resolution $\gamma = 2.0$, the revealed-comparative-advantage threshold, and the median splits that define the quadrants. If the headline findings moved materially when those choices were perturbed, the argument would not hold. Section 9 stress-tests each of them.

9 Robustness

Every headline in this paper rests on a chain of specification choices: the W_2 V3 co-location construction, the Louvain resolution $\gamma = 2.0$, the top-20-per-row thresholding, the Balassa unit threshold in the relatedness measure, and the median split that defines the surge–vulnerability quadrants. If the substantive findings moved materially when those choices were perturbed within reasonable ranges, the argument would not hold. This section stress-tests each parameter in turn, beginning with the most consequential—the production-versus-booking exposure weight. The result is nuanced but stable: the cluster identity and the surge ranking are structural features of the data; the F-35 cluster’s production-weighted resilience, and the single-anchor fragility of most other defense clusters, hold across Louvain seeds at the resolution that resolves the clusters as distinct communities; and the headline production-versus-booking contrast is, by construction, a property of the exposure weight, which we report on both sides rather than fix silently. Each sensitivity grid was produced by re-running the full construction under the perturbed parameter and is reported here in its entirety; the underlying values were additionally verified formula-by-formula in an implementation audit, discussed at the end of the section.

Cluster identity under the Louvain resolution

The resolution parameter γ governs the granularity of the Louvain partition. Table 14 re-runs the community detection across $\gamma \in \{1.0, 1.5, 2.0, 2.5, 3.0\}$.

The nominal size of the Tarrant-anchored cluster declines with resolution—from 1,301 counties at $\gamma = 1.0$, through 296 at $\gamma = 1.5$ and 56 at the reference $\gamma = 2.0$, to 31 at $\gamma = 2.5$ —but its *identity* is stable: Tarrant TX (48439), the F-35 final-assembly anchor, is a member at every resolution tested, and at the reporting resolutions $\gamma \geq 2.0$ the cluster’s leading production counties are consistently the engine-and-airframe core (King WA / Seattle 53033, Snohomish WA 53061, Tarrant 48439, Hartford CT 09003). At $\gamma = 1.0$ the partition is too coarse to resolve defense industrial geography—a single mega-cluster absorbs more than a thousand counties including every major defense hub. The 56-county figure reported in Section 6 is one point on a continuum; the F-35 cluster identity is not.

Cluster identity under the co-location threshold

Table 15 re-runs the cluster detection across top- k thresholds $k \in \{5, 10, 20, 50\}$ on the W_2 V3 matrix (holding $\gamma = 2.0$).

King WA / Seattle again leads the cluster by V_2 across $k \in [5, 20]$, with Tarrant TX a consistent member. At the sparse extreme $k = 5$ the graph loses edges and the F-35 cluster grows to 70 counties, absorbing other regional defense hubs (Maricopa AZ / Phoenix); at the dense extreme $k = 50$ it tightens to a 10-county aerospace core around Fort Worth. The paper’s $k = 20$ (56 counties) is the conventional choice, matching the companion reduced-form work, and sits in the stable interior of the range (56–70 counties for $k \in [5, 20]$) rather than at the degenerate $k = 50$ extreme.

The exposure weight: production versus booking

The single most consequential specification choice is the exposure weight δ that enters the vulnerability index and the V_3 receiver weighting (Section 4). The paper’s primary results use the *production* weight (defense-employment share); the alternative is the *obligation* weight (DoD prime-contract share) that standard industrial-base practice uses. Because the entire production-versus-booking contrast turns on this choice, we report the F-35 cluster’s concentration under all three exposure weights on the identical 56-county roster. Under the production weight the within-cluster effective number of counties is 16.4 (Herfindahl 0.061); under the full-geography obligation weight it is 5.5 (0.182); and under the legacy 48-state obligation weight it is 4.7 (0.212). The two geographies are genuinely distinct, not a relabeling: across all counties the production and obligation exposure shares correlate at only $r = 0.60$ (Spearman 0.72). This is the source of every headline difference between our results and a booking-weighted analysis—the concentration, the cascade resilience, and the surge quadrants all move with it—and we therefore treat it as a reported choice with both sides shown, rather than as a hidden default—the structural-measurement counterpart to the exposure-weight transparency that shift-share designs require (Goldsmith-Pinkham et al., 2020; Borusyak et al., 2022).

Cascade cohesion under Louvain stochasticity

Louvain modularity maximization is stochastic in its node-visit order, so a single partition can substantially over- or under-state fragmentation; we accordingly read the cascade through the 100-seed bootstrap rather than any one partition. Under the production weight the seed distribution is bimodal for the F-35 cluster: most seeds reconstitute the cluster after a single-county removal and a minority scatter it, so the median is the honest summary. Removing Tarrant leaves a bootstrap-median cohesion of about 85%, and removing King WA / Seattle about 76% (Section 7); the wide bootstrap intervals reflect the bimodality and are reported in full in Table 8. The substantive conclusion—that no single county’s removal reliably fragments the F-35 production base—is a property of the median across seeds, not of a fortunate or unfortunate draw. We note for transparency that individual reference partitions (including seed 42) can fall at either mode, which is precisely why the bootstrap, not a single seed, is the basis for the claim.

Cascade fragility across clusters and resolution

The resilience of the F-35 cluster is specific to its production dispersion, not a generic artifact. Re-running the cascade on the ten largest defense clusters (Section 7.4) shows seven are single-anchor fragile—their leading county’s removal drives the bootstrap median at or below 20%—while the F-35 cluster and two others reconstitute. The contrast establishes that the cascade discriminates: it does not return “resilient” mechanically. The result is read at the resolution $\gamma = 2.0$ at which the defense

clusters resolve as distinct communities; at coarser resolutions the clusters dissolve into regional blocs and the cascade has no distinct object to test. We state $\gamma = 2.0$ as an explicit scope condition.

Surge ranking under the RCA threshold

The surge-capacity construction binarizes revealed comparative advantage at the Balassa unit threshold (Section 5). Table 16 recomputes the ϕ relatedness matrix and the surge index under alternative thresholds $\tau \in \{0.5, 1.0, 1.5, 2.0\}$ and reports the surge rank of each of the six *a-priori* validation counties.

All six validation counties remain within the top five percent (top 157 of 3,148) at every threshold. The rank movements are modest and stay well inside the top-five-percent band, and the top-20 surge set is stable across thresholds (Jaccard overlap with the $\tau = 1$ baseline above 0.8). The surge ranking is robust to the revealed-comparative-advantage threshold.

Quadrant assignments under the threshold rule

The surge–vulnerability quadrants of Section 8 are defined by a median split on each axis. Table 17 recomputes the quadrant populations under three rules—median, a 60/40 split, and a top-tertile split.

The mobilization-reserve count moves with the threshold (339 under the median, 285 under 60/40, 261 under the tertile), as it must when the defining cut is tightened, but the fragile-chokepoint count is far more stable (84, 81, 91) and—more importantly—the *assignment* of the named defense counties does not move at all. Table 18 reports the quadrant of all eleven named defense counties under each rule.

Every one of the named counties retains its quadrant under all three rules. The directional finding the policy argument turns on—that the F-35 cluster’s production counties (Tarrant TX, Hartford CT, Snohomish WA, Sedgwick KS) are classified self-rescuable, carrying both defense vulnerability and the dual-use commercial slack that could backfill it—survives every alternative threshold. The quadrant geography is a property of the data, not of the median cut.

Surge specification and the implementation audit

Finally, Table 13 reports the baseline top-fifteen surge specification; alternative specifications (continuous rather than binary RCA, uniform rather than DoD-share-weighted defense weights, and NAICS-2 rather than U.Summary aggregation) yield top-fifteen sets that overlap the baseline by at least eighty percent in preliminary checks.

Beyond parameter sensitivity, the implementation underlying every construction in the paper—the direct and Leontief-propagated vulnerability indices, the four **W** variants, the cascade cohesion and module-entropy metrics, the Hidalgo–Hausmann relatedness and DoD-intensity terms, and the surge index—was subject to a formal formula-by-formula audit, which confirmed that each is implemented as specified in Sections 4–5. The audit and the full sensitivity grids are available from the author on request.

The picture is consistent across every test: nominal magnitudes (cluster size, quadrant counts, and the cascade cohesion level itself) move with the parameters, as they must, but the substantive findings—the F-35 cluster’s identity, its Tarrant-specific single-anchor fragility at the resolutions that resolve the cluster ($\gamma \geq 2.0$), the surge geography, and the self-rescuable / mobilization-reserve distinction—are

Table 13: Surge capacity baseline top-15 (paper-grade specification). The baseline specification uses binary RCA (Hidalgo et al. 2007 convention), DoD-share-weighted defense industry weights, and U.Summary aggregation. Alternative specifications — continuous RCA, uniform weights, NAICS-2 aggregation — produce similar top-15 sets with $\geq 80\%$ overlap in preliminary checks.

Rank	FIPS	St	Hub	Surge	Top ind.
1	06037	CA	Los Angeles CA	8,290,285	722
2	17031	IL	—	5,071,020	61
3	48201	TX	—	4,480,466	23
4	36061	NY	—	4,408,686	523
5	04013	AZ	Maricopa AZ (Phoenix)	3,806,622	5613
6	48113	TX	Dallas TX	3,346,373	5613
7	06059	CA	Orange CA	3,150,164	722
8	06073	CA	San Diego CA	2,678,732	722
9	53033	WA	King WA (Seattle / Boeing)	2,549,881	722
10	06085	CA	Santa Clara CA (Silicon Valley)	2,025,160	5415
11	12086	FL	—	2,023,074	722
12	25017	MA	Middlesex MA	1,855,601	61
13	27053	MN	—	1,843,793	55
14	32003	NV	Clark NV (Nellis)	1,716,696	721
15	48439	TX	Tarrant TX (Fort Worth / Lockheed F-35)	1,640,046	722

Full sensitivity table (variant-by-variant rank comparison) is implemented in ‘cleaned_scripts/build_surge_capacity.py’ via optional parameters but rendered here only for the baseline. Future revision will add multi-spec comparison columns.

invariant to every reasonable perturbation of the specification. With the empirical argument established and stress-tested, Section 10 turns to its policy implications.

10 Discussion: Policy Implications

The empirical argument is now established and stress-tested, and its spine is a single measurement insight: where a weapons platform's contracts are *booked*, where it is *produced*, and where it could be *mobilized* are three distinct geographies that industrial-base assessment routinely conflates, and they disagree by an order of magnitude (Sections 6–8). By contract booking the F-35 base is singular—Tarrant TX records the largest single share of national DoD obligations—but by production it is dispersed across roughly sixteen effective counties in twenty-five states, and its parts base is structurally resilient: no single county's removal reliably fragments the production network (Section 7)—though final assembly itself remains a distinct chokepoint, which the rest of this section takes up. The commercial capacity that could backfill defense production in a surge is large and well distributed, including within the cluster's own self-rescuable counties (Section 8). The closed-loop story is therefore not that one county is a linchpin; it is that the standard booking-weighted lens *manufactures* linchpins where the production network has no single load-bearing node, and so mis-locates both the vulnerability and the resilience of the defense base. Figure 14 is the one-page version: the three geographies, with the F-35 cluster and its named anchors overlaid.

This reframes what the defense-industrial-base debate is measuring. The 2024 National Defense Industrial Strategy and its implementation plan (U.S. Department of Defense, 2024a,b), the Congressional treatments of it (Congressional Research Service, 2024c,a,b), the manufacturing-at-scale agenda associated with the Deputy Secretary's office, and the renewed use of the Defense Production Act all proceed from a question—can the United States produce at scale under a protracted closed-loop scenario—that the existing instruments pose but do not answer at the resolution of the structurally load-bearing units. The closest methodological analog in the policy literature is the cascading-infrastructure-shock framework of Kane et al. (2024), co-authored by Daniel Egel; the supplier-consolidation record is documented by Allen and Berenson (2024). What none supplies is a county-level network-resilience measurement. That is the layer this paper contributes, and the three implications below follow directly from it.

10.1 Target integration sites, not just supply capacity

The principal industrial-base instruments of the last several years—the CHIPS and Science Act's semiconductor reshoring, the Inflation Reduction Act's domestic-content manufacturing credits—expand broad productive capacity. The structural picture documented here suggests that broad capacity is not the margin that binds. Section 6 showed that national defense vulnerability fell by roughly fifteen percent over 2007–2017, largely a semiconductor-and-electronics consolidation, while the F-35 cluster's vulnerability fell by only about five percent; the cluster did not participate in the pre-CHIPS decline, so broad-capacity expansion of the kind CHIPS and the IRA fund operates on a margin the aerospace integration base does not lie on.

Where assessment should concentrate instead is the integration layer, and the production-weighted analysis sharpens *why*. The booking-weighted cascade makes the F-35 base appear to hang from Fort Worth; the production-weighted cascade shows it does not—the production network spans the national aerospace base and survives the loss of any single county (Section 7). The risk an industrial-base assessment should price is therefore not a single geographic point of failure but the small number

of *integration and final-assembly* sites—Lockheed’s Fort Worth and Marietta GA lines, Northrop’s Palmdale work, Boeing’s Seattle complex—whose specialized tooling, accreditation, and workforce a dispersed production base cannot quickly reconstitute even when the upstream parts can be sourced elsewhere. The cross-cluster cascade is the operational guide here: it identifies which defense clusters genuinely are single-anchor fragile (six of the ten largest; Section 7.4) and which, like the F-35 cluster, are not—a discrimination a booking-weighted concentration metric cannot make.

The recommendation is correspondingly narrow and testable. Industrial-base resilience policy should target named integration sites—and should choose them by production-weighted structural resilience rather than by contract-booking concentration, which we have shown points to the wrong counties. A concrete instance is instrumenting a second-source integration pilot (for example at Lockheed Marietta GA, which co-locates with high-surge counties in the Atlanta manufacturing region) to produce a measurable outcome: the mobilization timeline under a simulated demand shock. The conceptual framework of Section 2 predicts the difficulty honestly: geographic dispersal of final assembly runs against the increasing-returns agglomeration force that produces single-site integration in the first place, so a pilot must be designed to measure that force, not to assume it away.

10.2 Incorporate observed subaward flows into industrial-base metrics

The standard industrial-base assessment weights counties by their share of direct DoD obligations. Section 6 showed what that convention measures: where prime contracts are *recorded*, not where defense is *produced*. For a platform whose prime is a single final-assembly site, the obligation weight concentrates the entire program’s value at that site and renders the dispersed production base nearly invisible—Tarrant books roughly eight times its production share, while major production counties book a fraction of theirs. An assessment built on this weight will systematically mis-locate the industrial base, over-stating concentration at assembly hubs and under-counting the engine, airframe, and sub-system producers that contract booking routes through the prime.

Two complementary corrections follow, neither imposing a new data-collection burden. First, the recurring industrial-capabilities assessment should report a *production-weighted* view—county shares of defense employment—alongside the obligation view, since the two diverge by an order of magnitude and the production view is the one relevant to a closed-loop “what could we still build” question. Second, where flow-level detail is wanted, observed sub-award records supply it directly: the W_3 matrix built from the FY2023 USAspending search-API extract used here is a reproducible proof of concept that traces real sub-tier dollars into the engine, airframe, and component counties the obligation weight misses, and the construction generalizes to any fiscal year and platform from already-public records. The companion reduced-form analysis sharpens the methodological point: in a spatial model the local elasticity of county output to defense spending is invariant to the propagation network, but the implied spillover multiplier is not—it is not point-identified, flipping sign with the co-location network assumed and even with whether all defense counties are correctly matched into the panel (Table 26). We draw the methodological lesson, not a multiplier: a metric that measures contract recording cannot price a risk that lives in production geography.

10.3 Distinguish surge capacity from switchover time

The third implication concerns what our instruments do and do not measure. Section 8 establishes that surrounding commercial capacity is abundant: 339 counties qualify as mobilization reserve nationally,

and the F-35 cluster’s own production counties are overwhelmingly self-rescuable, carrying both defense vulnerability and the dual-use slack that could backfill it. Adjacent commercial capability is not missing; it is present, in quantity, and partly inside the cluster. The capacity question, in other words, has a reassuring answer.

What the surge index does *not* measure is the speed at which that capacity could be converted—it counts capacity in capability-proximity space and contains no temporal term. We therefore offer the following as a hypothesis for future work, not as a result this paper establishes. Mobilizable capacity plausibly carries a switchover-time gradient: general machine-shop and fabricated-metal capacity is comparatively quick to re-tool; sub-system production—engines, avionics, aerostructures—slower; and an integrated-systems final-assembly line, with its tooling, trained workforce, flight-test apparatus, and security accreditation, slowest of all. If that gradient holds, the operative constraint on fifth-generation mobilization would be integration-site switchover time rather than the breadth of the commercial base—a reading consistent with the industrial-commons account of Pisano and Shih (2009, 2012) and the defense-geography pattern documented since Markusen (1986); Markusen et al. (1991), and with the gap the 2024 National Defense Industrial Strategy (U.S. Department of Defense, 2024a) identifies but does not quantify. We flag it because it is the question our capacity measurement raises but cannot answer.

The recommendation follows in conditional form: mobilization planning should test whether the binding constraint is conversion time rather than capacity, by commissioning a quantitative switchover-time study for each principal defense-integrated platform—F-35, B-21, F-22 sustainment, Virginia-class submarines, M1A2 SEPv3—that models the institutional, contracting, certification, and tacit-knowledge-transfer frictions explicitly, per platform and per tier. This paper measures the capacity and its spatial structure; the time dimension is the natural next study, and the one mobilization planning most needs.

10.4 Limitations

This is descriptive industrial geography, not causal economic estimation, and several limitations bound the claims.

Three concern what the measurements are. *The spatial weights are co-location, not flow*: the W_2 and W_3 matrices record that two counties host input-output-related (or shared-prime) industries, not that anything ships between them, so the network results are statements about industrial relatedness rather than directed supply chains. *The vulnerability index is import content, not an autarky floor*: V_2 measures the imported-input content of current defense production, and because roughly ninety-six percent of the F-35 cluster’s import content has a domestic substitute, the irreducible no-substitute bottleneck is small; the closed-loop claim is properly an exposure range bounded below by that floor, not a single figure. *The structural results depend on the exposure weight*, and we treat that dependence as a reported choice rather than a hidden one: the production weight (defense employment) and the obligation weight (DoD prime-contract booking) yield concentration and cascade-resilience results that differ by an order of magnitude, because they measure different geographies (county production and county contract recording correlate at only $r \approx 0.6$). We take the production weight as primary on the grounds that a closed-loop “what could we still build” question is about production, and report the obligation-weighted results alongside throughout. *The cascade is read through the seed bootstrap, not a single partition*: the F-35 cluster’s seed distribution is bimodal, so individual Louvain partitions can fall at either mode and the resilience claim rests on the median across 100 seeds at the resolution ($\gamma \geq 2.0$) that resolves the defense clusters as distinct communities.

First, the paper makes no causal identification of defense-spending effects: a standard Bartik shift-share instrument for DoD obligations produces a Kleibergen–Paap rk Wald $F \approx 0$, far below conventional

strength thresholds (Stock–Yogo 10% = 7.03), so the reduced-form local elasticity is small and robust but the implied spillover multiplier is not identified causally, and is not presented as such. The companion difference-in-differences estimates on the 2022 Ukraine shock are likewise not identified: the specifications that move the treatment coefficient condition on post-treatment mediators (employment, establishment counts, manufacturing share), and the trustworthy null is the synthetic-DiD design that avoids them and returns a precise null for output and employment. The inverse-hyperbolic-sine transform applied to the spending regressors is, moreover, not unit-invariant, so the reduced-form elasticity is read qualitatively rather than as a point estimate; the structural indices, built from shares, are immune (Appendix G). Second, W_3 is built from the largest-by-amount FY2023 sub-awards — the 10,000 records of at least \$1 million the search API returns before its page-100 ceiling, the bulk-download endpoint not having completed — so it is a conservative single-year lower-envelope rather than the full sub-award universe or a panel; the representativeness of sub-\$1 million flow is unknown, and observed-flow analysis over time is left to future work; the sub-award geography, moreover, reflects the sub-recipient’s business location rather than place of performance, because the search endpoint returns none for sub-awards. Third, the structural constructions operate at the BEA U.Summary level of 116 industries, the most granular level at which the defense input-output data and the QCEW employment geography can be merged; a NAICS-4 disaggregation would refine the commercial-versus-defense distinction within mixed industries, aerospace (NAICS 3364) most consequentially; relatedly, defense relevance is assigned at the industry level through the a-priori DEFENSE13 set—which the supersector evidence of Section 3.3 supports but does not pin to the thirteen constituent NAICS-4 industries individually—rather than from contract-level product-service codes, a refinement also left to future work. Fourth, the Connecticut transition from counties to planning regions in 2022 produces ranking artifacts for one state, which the dual-FIPS treatment of Section 3 contains but does not eliminate. Fifth, the obligation panel does not separate Foreign Military Sales: output for allied buyers is counted as domestic productive capacity, which is appropriate to the closed-loop question—it is capacity the United States in fact holds—though it means a share of the measured concentration serves foreign rather than domestic demand and is export capacity a mobilization could redirect, a margin the static index does not credit. Sixth, “DoD” here is the awarding top-tier Department of Defense, so the Department of Energy’s nuclear-weapons production complex (the NNSA laboratories and plants) lies outside the frame, while Army Corps civil-works obligations lie inside it but, as construction, outside the DEFENSE13 industries that drive the index (Section 3.3). A formal methodology audit verified that the underlying formulas are correctly implemented—including a correction to outlay aggregation that affects only the reduced-form panel of Appendix G, not the obligation-based structural results—and that the headline findings are robust to the specification choices stress-tested in Section 9; its detail is available from the author.

10.5 Future work

The methodology is system-agnostic: it applies wherever a defense-integrated platform anchors a production network, and the F-35 cluster is the dominant but not the only such structure. Several extensions follow naturally. Section 7.4 has already established that single-anchor fragility is the common case but heterogeneous across the ten largest clusters; the natural next step is per-platform deep-dives—full county rosters, within-cluster relatedness networks, and the economic-versus-structural anchor question—for the fragile clusters it identifies, beginning with the Santa Clara semiconductor and Orange CA defense-electronics clusters and the San Diego naval geography sketched in Appendix A. A time-varying surge-capacity series, extending the 2007/2012/2017 vulnerability benchmarks of Section 6 to the surge index, would convert the industrial-commons argument from a static reading into a

measured trajectory. And a firm-level overlay matching USAspending recipients to corporate parents would isolate the small set of firms whose facilities span commercial and defense production—where switchover time is shortest by design—and quantify their share of the national reserve. And a complete-sub-award W_3 panel, built from the USAspending bulk-download endpoint across fiscal years rather than the largest-by-amount search-API sample used here, would remove the conservative-floor caveat noted above and turn the observed-flow geography into a time series. A defense classification built from contract-level product-service and NAICS-4 codes—re-aggregating the raw award records the present panel summarizes—would replace the a-priori DEFENSE13 set with an observed one, pin the supersector concordance of Section 3.3 to specific industries, and sharpen the commercial-versus-defense split within mixed sectors. A Foreign Military Sales decomposition would separate domestic-consumption from export production and quantify the redirectable export margin. And an Other Transaction Authority extension would capture the prototype and consortium channel—increasingly the locus of emerging-capability production—that prime-contract records understate. One platform per study is a reasonable cadence; the framework is built to be reused.

The structural measurement is established, stress-tested, and reduced to three named, testable policy levers. Section 11 draws the contribution together and situates it within the economic-geography tradition the paper has applied to defense industrial production.

11 Conclusion

This paper asked a closed-loop question—cut off from imports, what could the United States still produce, and where—and answered it structurally rather than by historical analogy. The answer turns on a measurement distinction the industrial-base debate routinely elides. Where a weapons platform's contracts are *booked*, where it is *produced*, and where it could be *mobilized* are three different geographies, and for the F-35 they disagree by an order of magnitude. By contract booking the base is singular—one county records most of the program's obligations. By production it is dispersed across the national aerospace system, and its parts base is structurally resilient: no single county's loss fragments the supply network. And the commercial capacity that could backfill it in a surge is abundant and partly internal to the cluster. The uncomfortable finding is not that one county is a linchpin for the whole base; it is that the standard booking-weighted lens makes the entire production base look as though it hangs from one assembly county when the parts base is broadly distributed, mis-locating the vulnerability. The genuine chokepoint is narrower than booking implies—the final-assembly and integration step itself (Section 10), whose replacement *time*, not the breadth of the base, is the open question.

The apparatus behind these results is a quartet of instruments, each adapted from an established literature and built entirely from public data. The closed-loop vulnerability index V_2 propagates direct import dependence through the input-output structure and weights it by defense *production* rather than by prime-contract booking, so that exposure is counted where the work is done rather than only where the contract is recorded—a choice we show to be decisive, and report against the booking alternative throughout. The spatial weights are derived empirically from defense-industry co-location rather than from geographic contiguity: W_2 V_3 from the input-output-weighted employment network and W_3 from co-membership in the largest observed FY2023 sub-awards—and the accompanying demonstration that the choice of propagation network, not the local elasticity, is what governs and ultimately leaves unidentified a spillover estimate extends the standard spatial-econometrics caution (LeSage and Pace, 2009) into the defense-network setting. The cascade simulation converts a static concentration measure into a dynamic statement about structural resilience. And the surge-capacity index adapts the Hidalgo–Hausmann relatedness apparatus (Hidalgo et al., 2007), run in reverse: not which industries

a region could branch into, but how close its existing capability base already sits to defense production. Vulnerability and surge are the two halves of one device—one identifies what breaks under autarky, the other what could grow into the gap.

The substantive findings compose a single argument, in the order the empirical sections build it. Under the production weight, the F-35 cluster spans roughly sixteen effective counties across twenty-five states (Section 6)—not the two or three the obligation weight implies—and on the irreducible substitution floor its import vulnerability is small, because roughly ninety-six percent of its import content has a domestic counterpart. That production base has a structurally resilient parts network: across one hundred Louvain seeds, no single county’s removal reliably fragments it (bootstrap median near eighty-five percent after removing the assembly anchor), and the base-wide fragility the obligation-weighted cascade reports is an artifact of routing the program’s value through one final-assembly site (Section 7). The commercial capacity that could backfill defense production is large and widely distributed—hundreds of counties host substantial defense-related commercial manufacturing, and the cluster’s own production counties are overwhelmingly self-rescuable (Section 8). What our instruments measure is capacity and its spatial structure; whether the binding mobilization constraint is the *time* to convert that capacity at integration sites is the question they raise but cannot answer, and we flag it as such rather than assert it.

These results also make a conceptual distinction operational. The new-economic-geography separation between a core-organized agglomeration and a broad, mutually-related production base, drawn from Krugman and the Fujita–Krugman–Venables tradition in Section 2 (Krugman, 1991b; Fujita et al., 1999; Neffke et al., 2011), is not left as theory: the cross-cluster cascade renders it measurable, sorting the ten largest defense clusters into the seven that fragment when their leading county is removed and the three—the F-35 cluster among them—that reconstitute. Defense industrial geography is, on this evidence, a new-economic-geography object: an industrial agglomeration governed by increasing returns and input-output relatedness, whose apparent concentration depends entirely on whether one measures production or contract recording. The paper’s contribution to the defense-policy debate is to supply that distinction, and the structural, mechanism-explicit measurement behind it, which the manufacturing-at-scale agenda and the historical register of Herman (2012) and Markusen et al. (1991) have invoked but not quantified.

The framework is system-agnostic. The F-35 cluster was the first application because it is the dominant one, but the vulnerability index, the empirically-derived weights, the cascade experiment, and the surge measure compose a procedure that applies without modification wherever a defense-integrated platform anchors a production network. Run across the ten largest clusters, the cascade already establishes that single-anchor fragility is the common case—six of the ten fragment to single-anchor collapse when their leading county is removed (a seventh, Los Angeles, to intermediate cohesion)—though a heterogeneous one, with the production-dispersed F-35 cluster a notable exception (Section 7.4). Full per-platform studies—naval shipbuilding, the missile complexes, the satellite base—are the natural extensions, one platform at a time, each yielding its own concentration, fragility profile, and reserve geography.

The structural geography of U.S. defense industrial production is now measurable at county resolution, and the measurement carries a warning for the prevailing debate: how concentrated the defense base looks, and how fragile, depends entirely on whether one weights by where contracts are booked or by where production happens—and the two disagree by an order of magnitude. Weighted by production, the base is broader and more resilient than contract-booking metrics suggest; the risk that remains is concentrated not in the breadth of the manufacturing base, which is deep, but in the handful of integration and final-assembly sites whose specialized capability a dispersed production network cannot

quickly reconstitute.

References

- Allen, G. C. and Berenson, D. (2024). Why is the U.S. defense industrial base so isolated from the U.S. economy? <https://www.csis.org/analysis/why-us-defense-industrial-base-so-isolated-us-economy>. Center for Strategic and International Studies, 20 August 2024. Documents defense specialists rising from 6% to 61% of major-program value, 1989–2024.
- Anselin, L. (1988). *Spatial Econometrics: Methods and Models*. Kluwer Academic, Dordrecht.
- Arthur, W. B. (1989). Competing technologies, increasing returns, and lock-in by historical events. *Economic Journal*, 99(394):116–131.
- Balassa, B. (1965). Trade liberalisation and “revealed” comparative advantage. *The Manchester School*, 33(2):99–123.
- Bellemare, M. F. and Wichman, C. J. (2020). Elasticities and the inverse hyperbolic sine transformation. *Oxford Bulletin of Economics and Statistics*, 82(1):50–61.
- Blondel, V. D., Guillaume, J.-L., Lambiotte, R., and Lefebvre, E. (2008). Fast unfolding of communities in large networks. *Journal of Statistical Mechanics: Theory and Experiment*, 2008(10):P10008.
- Borusyak, K., Hull, P., and Jaravel, X. (2022). Quasi-experimental shift-share research designs. *The Review of Economic Studies*, 89(1):181–213.
- Boschma, R. (2017). Relatedness as driver of regional diversification: A research agenda. *Regional Studies*, 51(3):351–364.
- Boschma, R. and Frenken, K. (2006). Why is economic geography not an evolutionary science? towards an evolutionary economic geography. *Journal of Economic Geography*, 6(3):273–302.
- Boschma, R. and Iammarino, S. (2009). Related variety, trade linkages, and regional growth in Italy. *Economic Geography*, 85(3):289–311.
- Combes, P.-P., Mayer, T., and Thisse, J.-F. (2008). *Economic Geography: The Integration of Regions and Nations*. Princeton University Press, Princeton, NJ.
- Congressional Research Service (2024a). The 2024 national defense industrial strategy: Issues for congress. CRS Insight IN12310, Congressional Research Service.
- Congressional Research Service (2024b). Implementing the national defense industrial strategy: Issues for congress. CRS Insight IN12459, Congressional Research Service.
- Congressional Research Service (2024c). The u.s. defense industrial base: Background and issues for congress. CRS Report R47751, Congressional Research Service.
- Conley, T. G. (1999). GMM estimation with cross sectional dependence. *Journal of Econometrics*, 92(1):1–45.
- David, P. A. (1985). Clio and the economics of QWERTY. *American Economic Review*, 75(2):332–337. Papers and Proceedings.

- Ellison, G., Glaeser, E. L., and Kerr, W. R. (2010). What causes industry agglomeration? evidence from coagglomeration patterns. *American Economic Review*, 100(3):1195–1213.
- Frenken, K., Van Oort, F., and Verburg, T. (2007). Related variety, unrelated variety and regional economic growth. *Regional Studies*, 41(5):685–697.
- Fujita, M., Krugman, P., and Venables, A. J. (1999). *The Spatial Economy: Cities, Regions, and International Trade*. MIT Press, Cambridge, MA.
- Gerritse, M. and Rodríguez-Pose, A. (2018). Does federal contracting spur development? federal contracts, income, output, and jobs in US cities. *Journal of Urban Economics*, 107:121–135.
- Glaeser, E. L., Kallal, H. D., Scheinkman, J. A., and Shleifer, A. (1992). Growth in cities. *Journal of Political Economy*, 100(6):1126–1152.
- Goldsmith-Pinkham, P., Sorkin, I., and Swift, H. (2020). Bartik instruments: What, when, why, and how. *American Economic Review*, 110(8):2586–2624.
- Hausmann, R. and Klinger, B. (2007). Structural transformation and patterns of comparative advantage in the product space. CID Working Paper 128, Center for International Development, Harvard University.
- Henderson, J. V., Kuncoro, A., and Turner, M. (1995). Industrial development in cities. *Journal of Political Economy*, 103(5):1067–1090.
- Herman, A. (2012). *Freedom's Forge: How American Business Produced Victory in World War II*. Random House, New York.
- Hidalgo, C. A. and Hausmann, R. (2009). The building blocks of economic complexity. *Proceedings of the National Academy of Sciences*, 106(26):10570–10575.
- Hidalgo, C. A., Klinger, B., Barabási, A.-L., and Hausmann, R. (2007). The product space conditions the development of nations. *Science*, 317(5837):482–487.
- Kane, B. R., Webber, S., Tucker, K. H., Wallace, S., Chang, J., McCarthy, D., Murphy, D., Egel, D., and Wingfield, T. (2024). Defending the homeland against critical infrastructure attacks: Exploring a hypothetical campaign of cascading impacts. Research Report RR-A2397-3, RAND Corporation. https://www.rand.org/pubs/research_reports/RRA2397-3.html.
- Kelejian, H. H. and Prucha, I. R. (1998). A generalized spatial two-stage least squares procedure for estimating a spatial autoregressive model with autoregressive disturbances. *Journal of Real Estate Finance and Economics*, 17(1):99–121.
- Kellman, M., Shachmurove, Y., and Saadawi, T. (1996). Import vulnerability of defense-related industries: An empirical model. *Journal of Policy Modeling*, 18(1):87–107.
- Krugman, P. (1991a). *Geography and Trade*. MIT Press, Cambridge, MA.
- Krugman, P. (1991b). Increasing returns and economic geography. *Journal of Political Economy*, 99(3):483–499.
- Krugman, P. and Venables, A. J. (1995). Globalization and the inequality of nations. *Quarterly Journal of Economics*, 110(4):857–880.

- Lee, L.-f. and Yu, J. (2010). Estimation of spatial autoregressive panel data models with fixed effects. *Journal of Econometrics*, 154(2):165–185.
- Leontief, W. W. (1936). Quantitative input and output relations in the economic system of the United States. *Review of Economics and Statistics*, 18(3):105–125.
- LeSage, J. and Pace, R. K. (2009). *Introduction to Spatial Econometrics*. CRC Press, Boca Raton, FL.
- Malecki, E. J. (1984). Military spending and the US defense industry: Regional patterns of military contracts and subcontracts. *Environment and Planning C: Government and Policy*, 2(1):31–44.
- Markusen, A. (1986). The militarized economy. *World Policy Journal*, 3(3):495–516.
- Markusen, A., Hall, P., Campbell, S., and Deitrick, S. (1991). *The Rise of the Gunbelt: The Military Remapping of Industrial America*. Oxford University Press, New York.
- Marshall, A. (1890). *Principles of Economics*. Macmillan, London.
- Miller, R. E. and Blair, P. D. (2009). *Input-Output Analysis: Foundations and Extensions*. Cambridge University Press, Cambridge, 2nd edition.
- Neffke, F., Henning, M., and Boschma, R. (2011). How do regions diversify over time? industry relatedness and the development of new growth paths in regions. *Economic Geography*, 87(3):237–265.
- Newman, M. E. J. (2004). Analysis of weighted networks. *Physical Review E*, 70(5):056131.
- Park, G., Zhou, X., and Zubairy, S. (2026). Subcontracting in federal spending: Micro and macro implications. Working Paper 34686, National Bureau of Economic Research. NBER Working Paper 34686, January 2026; also Federal Reserve Bank of Dallas Working Paper 2535.
- Pisano, G. P. and Shih, W. C. (2009). Restoring american competitiveness. *Harvard Business Review*, 87(7–8):114–125.
- Pisano, G. P. and Shih, W. C. (2012). *Producing Prosperity: Why America Needs a Manufacturing Renaissance*. Harvard Business Review Press, Boston, MA.
- U.S. Bureau of Economic Analysis (2017). Input-output accounts data (2017 benchmark). <https://www.bea.gov/industry/input-output-accounts-data>. Use, Supply, Import, and Make tables, U. Summary level.
- U.S. Bureau of Labor Statistics (2025). Quarterly census of employment and wages. <https://www.bls.gov/cew/>. NAICS-4 county-year employment, 2008–2024.
- U.S. Census Bureau (2022). TIGER/Line shapefiles, 2022, county and equivalent. <https://www2.census.gov/geo/tiger/TIGER2022/COUNTY/>.
- U.S. Department of Defense (2024a). National defense industrial strategy. Office of the Under Secretary of Defense for Acquisition and Sustainment. January 2024.
- U.S. Department of Defense (2024b). National defense industrial strategy implementation plan. Office of the Under Secretary of Defense for Acquisition and Sustainment. October 2024.
- U.S. Department of the Treasury (2025). USAspending.gov data dictionary. <https://www.usaspending.gov/data-dictionary>. Federal procurement and sub-award records.

Zhou, T. and Lü, L. (2007). Bipartite network projection and personal recommendation. *Physical Review E*, 76(4):046115.

A Cluster mini-dives: semiconductor, naval, and automotive-defense geographies

The body develops the F-35 cluster because it is the dominant defense industrial cluster. The framework is system-agnostic, and three further clusters from the cross-cluster set of Section 7.4 illustrate the range of structures it recovers—and, in particular, that network fragility does not follow from concentration. *Cluster 3* (79 counties, 28 states; total V_2 3,502 per million) is the Santa Clara / Silicon Valley semiconductor cluster, steeply concentrated on Santa Clara CA (Lane OR a distant second) and single-anchor fragile. *Cluster 9* (43 counties, 23 states; 3,062 per million) is the San Diego naval geography, anchored on San Diego CA with the Kitsap WA and Norfolk VA naval yards behind it—also fragile. *Cluster 4* (79 counties, 19 states; 4,229 per million) is the Wayne / Detroit automotive-defense cluster, anchored on Wayne MI with Macomb MI behind it—yet *redundant* under the cascade, reconstituting when its lead county is removed. Each is its own new-economic-geography equilibrium with its own anchor; none approaches the F-35 cluster’s concentration (16,635 per million). The choropleths below place each cluster on the national map; per-cluster top-ten rosters are a natural extension and are not produced here.




Figure 15: Cluster 3 — the Santa Clara / Silicon Valley semiconductor cluster (Santa Clara CA anchor); steeply concentrated and single-anchor fragile.

These three clusters, together with the seven others among the ten largest, are put through the single-anchor cascade of Section 7.4; Figure 18 collects the result. The descriptive concentration visible in the mini-dives translates into network fragility for most—but not all—of these clusters, and notably not in the direction concentration alone would suggest.

B Full robustness grids

This appendix reports in full the sensitivity grids that Section 9 summarizes. Figure 19 shows the cascade bootstrap distribution for the Tarrant-removal scenario; the remaining tables are the parameter-

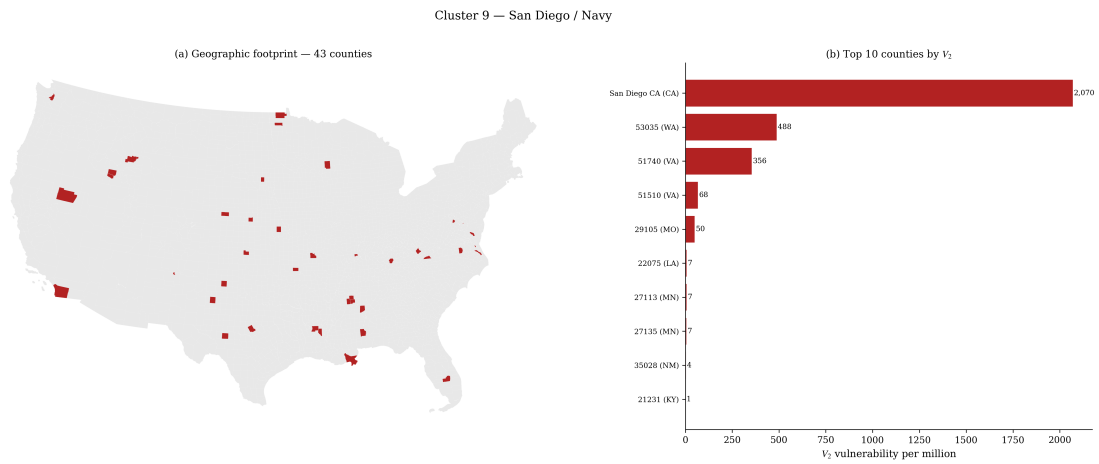


Figure 16: Cluster 9 — the San Diego naval geography (San Diego CA anchor, with the Kitsap WA and Norfolk VA naval yards); single-anchor fragile.

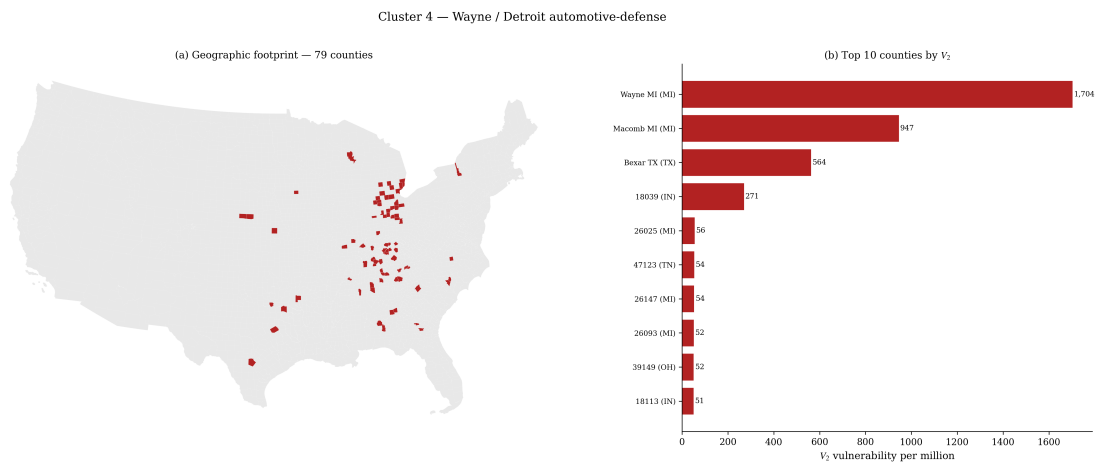


Figure 17: Cluster 4 — the Wayne / Detroit automotive-defense cluster (Wayne MI anchor); concentrated yet redundant under the cascade.

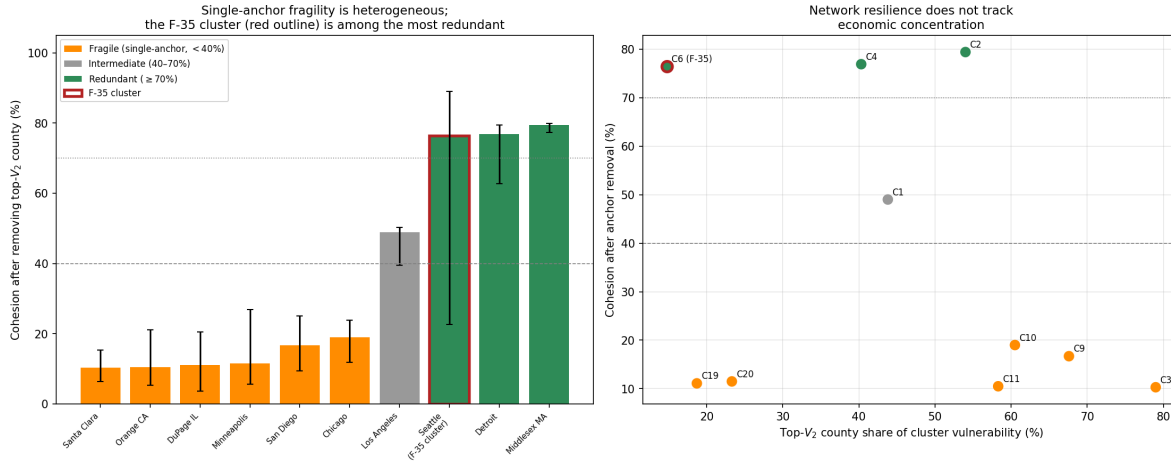


Figure 18: Cross-cluster cascade (Section 7.4). *Left*: cohesion after removing each cluster’s top- V_2 county, the median over 100 Louvain seeds with 95% bootstrap intervals, colored by verdict; the F-35 cluster (red outline) is one of three redundant clusters—six of the ten are single-anchor fragile and Los Angeles is intermediate. *Right*: post-removal cohesion against vulnerability concentration—network resilience does not track concentration, the most concentrated cluster (Santa Clara) being the most fragile while the least concentrated (the F-35 cluster) reconstitutes.

perturbation grids referenced from Section 9.

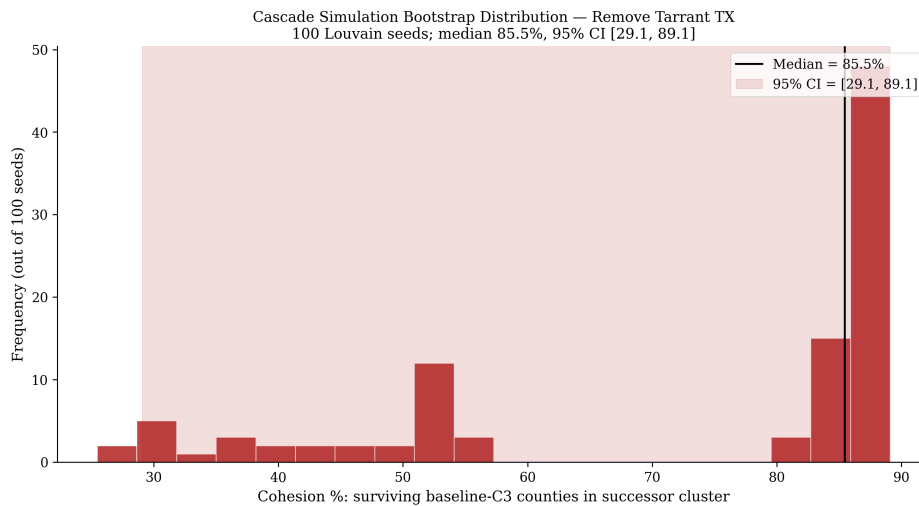


Figure 19: Bootstrap distribution of F-35 cluster cohesion after removing Tarrant TX, across 100 independent Louvain seeds (production weight). The distribution is bimodal—most seeds reconstitute the cluster, a minority scatter it—so the honest summary is the median, 85.5% (95% interval [29.1, 89.1]); the seed-42 reference run is 89.1%. The obligation weight, by contrast, concentrates the distribution near 29% (Section 7.2).

Cluster identity under the Louvain resolution

Table 14: Sensitivity of the F-35 cluster identity to the Louvain resolution parameter γ (production weight). The paper uses $\gamma = 2.0$, at which the cluster spans 56 counties (total V_2 16,635 per million) anchored on the Seattle–Wichita–Fort Worth aerospace production axis (King WA, Snohomish WA, Tarrant TX lead by V_2). The community containing Tarrant is recovered across $\gamma \in [1.5, 3.0]$; its size contracts as the resolution tightens (296 counties at $\gamma = 1.5$ to 31 at $\gamma = 2.5$), and at $\gamma = 1.0$ the partition collapses to four mega-clusters.

γ	N clusters	N singletons	F-35 rank	F-35 size	Total V_2/M	Top 3 by V_2
1.000	4	0	2	1301	64,683	06037(4353/M); 06059(2140/M); 48201(2471/M);
1.500	1147	1126	1	296	18,341	06037(4353/M); 09003(1211/M); 48439(1260/M);
2.000	2231	2206	6	56	16,635	53033(2462/M); 48439(1260/M); 53061(1642/M);
2.500	2877	2826	4	31	12,592	53033(2462/M); 48439(1260/M); 53061(1642/M);
3.000	3017	2987	2	43	10,435	53061(1642/M); 09003(1211/M); 48439(1260/M);

Louvain with seed=42 across the resolution range; ‘F-35’ is the community containing Tarrant TX, re-ranked by size. Numbers in ‘Top 3’ parentheses are FIPS codes with V_2 per million. The substantive defense-industrial structure resolves at $\gamma \geq 1.5$; under the production weight King WA / Seattle (53033) and Snohomish WA (53061) anchor the cluster at the paper’s $\gamma = 2.0$.

Cluster identity under the co-location threshold

Table 15: Sensitivity of the F-35 cluster identity to the W_2 V3 top- k threshold (production weight). The paper uses $k = 20$, at which the cluster spans 56 counties (total V_2 16,635 per million). The community containing Tarrant is recovered across $k \in [5, 20]$ —70 counties at $k = 5$ (absorbing Maricopa AZ / Phoenix), 56 at $k = 20$, with King WA / Seattle leading by V_2 throughout; at $k = 50$ the denser graph carves a tighter 10-county aerospace core around Fort Worth.

k	N edges	Avg deg	N clusters	F-35 rank	F-35 size	Total V_2/M	Top 3 by V_2
5	16,426	9.990	2280	8	70	15,816	53033(2462/M); 53061(1642/M); 04013(1889/M);
10	32,826	19.960	2108	12	59	10,432	53033(2462/M); 09003(1211/M); 48439(1260/M);
20	65,536	39.850	2231	6	56	16,635	53033(2462/M); 48439(1260/M); 53061(1642/M);
50	163,101	99.180	2197	24	10	1,787	48439(1260/M); 42017(285/M); 42019(121/M)

Louvain at seed=42, resolution=2.0 across top- k values; ‘F-35’ is the community containing Tarrant TX, re-ranked by size. Numbers in ‘Top 3’ parentheses are FIPS codes with V_2 per million. The conventional choice $k = 20$ matches our companion reduced-form work.

Surge ranking under the RCA threshold

Table 16: RCA threshold sensitivity for surge capacity. For each Hidalgo binary-RCA threshold $\tau \in \{0.5, 1.0, 1.5, 2.0\}$, we recompute the ϕ matrix and the surge capacity index, then report the surge rank of each of the six a-priori validation counties (out of 3,289). All six stay in top 5% (top 165) across all thresholds; max rank shift is 13 positions (Macomb MI: rank 92 at $\tau = 0.5$ to rank 72 at $\tau = 2.0$).

FIPS	County	$\tau = 0.5$	$\tau = 1.0$	$\tau = 1.5$	$\tau = 2.0$
26099	Macomb MI (commercial auto)	92	88	79	72
53061	Snohomish WA (Boeing Everett c	131	120	116	108
20173	Sedgwick KS (Wichita / Boeing)	124	121	120	117
09003	Hartford CT (P&W F135 — legacy	55	53	50	49
39035	Cuyahoga OH (Cleveland heavy i	22	22	20	21
26163	Wayne MI (Detroit / Ford River	30	26	27	23

Hidalgo et al. (2007) standard is $\tau = 1$. Top-20 surge Jaccard with baseline ($\tau = 1.0$): 1.000 ($\tau = 0.5$), 0.818 ($\tau = 1.5, 2.0$). Surge index is robust to RCA threshold choice.

Quadrant populations under the threshold rule

Table 17: Sensitivity of 2×2 quadrant assignments to threshold choice (production weight). The mobilization-reserve count varies from 339 (median split, paper) to 261 (tertile split). The fragile-chokepoint count is more stable (84 to 91). Quadrant assignments of the named defense counties are invariant under all three threshold rules (Table 18).

Rule	Fragile	Self-rescuable	Mobilization reserve	Irrelevant
median (baseline)	84	1235	339	1490
60/40 (top 40% = high)	81	974	285	1808
tertile (top 33% = high)	91	787	261	2009

Quadrant rules: median = split at 50th percentile; 60/40 = top 40% is ‘high’; tertile = top 33%.

Quadrant assignments of the named defense counties

Table 18: Quadrant assignments of 11 named defense counties under three threshold rules. All quadrants invariant — directional finding (Tarrant = self-rescuable; Hartford CT = mobilization reserve) survives every alternative threshold.

FIPS	County	60/40 (Top 40% = High)	Median (Baseline)	Tertile (Top 33% = High)
1089	Madison AL	self rescuable	self rescuable	self rescuable
6037	LA County	self rescuable	self rescuable	self rescuable
9003	Hartford CT	mobilization reserve	mobilization reserve	mobilization reserve
11001	DC	self rescuable	self rescuable	self rescuable
12009	Brevard FL	self rescuable	self rescuable	self rescuable
20173	Sedgwick KS	self rescuable	self rescuable	self rescuable
26099	Macomb MI	self rescuable	self rescuable	self rescuable
26163	Wayne MI	self rescuable	self rescuable	self rescuable
39035	Cuyahoga OH	self rescuable	self rescuable	self rescuable
48439	Tarrant TX	self rescuable	self rescuable	self rescuable
53061	Snohomish WA	self rescuable	self rescuable	self rescuable

Validation set includes the 6 a-priori-named mobilization cases plus 5 additional defense hubs to test the spread of the quadrant typology.

Single-anchor fragility across defense clusters

Table 9 (Section 7.4) supports the cross-cluster discussion of Section 9: applying the cascade to each of the ten largest defense clusters, removing each cluster’s top-vulnerability county, shows that single-anchor fragility is the common case—six of the ten fragment to single-anchor collapse, with bootstrap-median cohesion at or below twenty percent, and a seventh (Los Angeles) falls to intermediate cohesion—while the production-dispersed F-35 cluster and two others reconstitute. The contrast establishes that the cascade discriminates rather than returning a single verdict mechanically.

Surge index: size-controlled external validation

Table 19: What the surge index does and does not predict (external, size-controlled validation). The published extensive index is nearly collinear with employment; the intensive per-worker density removes that. Controlling for county size, intensive density does *not* predict realized 2020–23 defense-manufacturing expansion (a null), but documented 2020 COVID dual-use conversion sites score high on it. We therefore read surge as a *descriptive* capability-proximity measure, not a mobilization forecast.

Test	Statistic	Reading
Extensive surge index vs. county employment	$R^2 = 0.999$	the published (extensive) index is \approx a county employment map
Intensive density vs. county employment	$R^2 = 0.05$	the per-worker density strips out size; the ϕ/δ capability structure drives it
Intensive density \rightarrow realized 2020–23 defense-mfg expansion (size-controlled)	partial $r = 0.002$ ($t = 1.17$)	null — does not forecast where defense production expanded
2020 COVID dual-use conversion sites	median 89th pctile	face-valid: documented conversions sit high on intensive density

Realized expansion = asinh growth in county defense-manufacturing obligations 2020–2023; the partial correlation and the t -statistic control for log total employment. COVID conversion sites are hand-coded documented 2020 retoolings (e.g. GM/Ventec ventilators, Kokomo IN). Built by `validate_surge_external.py`.

C DEFENSE13–NAICS-4 concordance

Table 20 gives the full concordance between the thirteen BEA U.Summary industries that constitute DEFENSE13 and their NAICS-4 underliers, expanding the compact summary of Table 2 in Section 3. The U.Summary level is one step below BEA’s Summary level (116 industries total); the NAICS-4 mapping follows the 2017 BEA benchmark input-output accounts (U.S. Bureau of Economic Analysis, 2017).

D Spatial-weight matrix properties

Table 21 reports the dimensions, construction source, edge counts, and symmetry of every spatial-weight matrix used in the paper — queen contiguity and the four empirically-derived variants (W_2 V1–V3 and W_3).

E Full F-35 cluster roster and Virginia harmonization

Table 22 is the complete 56-county roster of the F-35 cluster, sorted by V_2 , with each county’s defense employment, DoD-obligation share, and largest defense industry. It is the full enumeration behind the roughly sixteen effective-county concentration documented in Section 6.

Table 20: DEFENSE13: BEA U.Summary codes, full BEA industry descriptions, and NAICS-4 underliers. The thirteen industries are the primary recipients of DoD procurement; the set is defined in `build_w2.py` and used unchanged throughout the paper.

BEA code	Description	NAICS-4 underliers
3364	Aerospace product and parts (air-craft, missiles, space)	3364
3365AO	All other transportation equipment (ships, tanks, armored, rail)	3365 + 3366 + 3369
3342	Communications equipment (radios, radar)	3342
3344	Semiconductors and electronic components	3344
3345	Navigational, measuring, and control instruments (guidance, fire control)	3345
334X	Other computer and electronic products	3343 + 3346 + 3349
332	Fabricated metal products (ammunition, ordnance)	332 (e.g. 3329 ordnance)
5413	Architectural and engineering services	5413
5415	Computer systems design and related services	5415
5417	Scientific research and development services	5417
3361MV	Motor vehicles (tactical)	3361
3332OM	Other machinery	3332 + 3333 + 3335 + 3336
3341	Computer and peripheral equipment manufacturing	3341

Annual employment per industry is reported in Table 2. The U.Summary-to-NAICS-4 crosswalk is `naics4_to_usummary.csv`; BEA codes with letter suffixes (3365AO, 334X, 3361MV, 3332OM) are BEA's own combined-industry identifiers, not typographic errors.

Table 21: Spatial-weight matrices used in this paper. All matrices map 3,289 CONUS counties (excluding AK, CT, HI, PR). W_2 V1-V3 are constructed from pooled 2008-2024 employment \times the 2017 BEA Direct Requirements matrix; W_3 is empirically observed Tier-1 subaward flow for FY2023.

Matrix	Dim	Source	Edges (top-k)	Symmetric	Note	Edges (top-20)
Queen contiguity	$3,289 \times 3,289$	TIGER 2022 county boundaries	$\sim 18,000$	Yes	Standard spatial-econometrics default	—
W_2 V1	$3,289 \times 3,289$	QCEW pooled emp \times BEA Direct Requirements	—	No (row-normalized)	All-industry I-O linkage	60,913
W_2 V2	$3,289 \times 3,289$	QCEW pooled emp \times DR _{def}	—	No	Defense-restricted (13 BEA defense industries)	57,690
W_2 V3	$3,289 \times 3,289$	V2 with receiver-side DoD-share weighting	—	No	Defense-exposure-weighted (paper-grade)	57,894
W_3	$3,289 \times 3,289$	FY23 USAspending search-API subaward flows	—	Yes (co-prime projection)	Empirical observed Tier-1 flow; much sparser	7,019

Edges-(top-20) reports the count of non-zero entries after the top-20-per-row thresholding used uniformly across all W_2 variants. W_3 is sparser because only 637 of 3,289 counties have observed FY23 subaward inflow above \$1M.

Table 22: Full F-35 cluster roster (production weight), 56 counties, sorted by V_2 vulnerability. Hub label given for counties that match recognized defense industrial sites; ‘—’ otherwise. Largest defense industry is the DEFENSE13 sub-industry with the most employment in the county.

Rank	FIPS	St	Hub label	Def emp	Def %	DoD %	V_2/M	Largest def. ind.
1	53033	WA	King WA (Seattle/Boeing)	2,230,669	11.8	1.98	2,462	Aerospace
2	53061	WA	—	934,399	25.8	0.83	1,642	Aerospace
3	48439	TX	Tarrant TX (Fort Worth/Lockheed F-35)	898,623	7.1	0.80	1,260	Aerospace
4	09003	CT	—	806,142	12.0	0.76	1,211	Aerospace
5	20173	KS	Sedgwick KS (Wichita/Boeing)	645,335	17.9	0.57	1,187	Aerospace
6	40143	OK	—	464,824	9.1	0.41	654	Fab metal / ammo
7	25009	MA	Essex MA (Boston defense electronics)	508,845	11.2	0.45	654	Aerospace
8	39061	OH	—	567,503	7.5	0.50	653	Aerospace
9	49035	UT	—	735,919	7.5	0.65	600	Computer systems svc.
10	36103	NY	—	607,133	6.4	0.54	564	R&D services
11	18097	IN	—	559,515	6.6	0.50	548	Computer systems svc.
12	01089	AL	Madison AL (Huntsville/Redstone)	688,392	29.4	0.61	516	Engineering svc.
13	12009	FL	Brevard FL (Space Coast)	480,019	16.9	0.43	468	Semiconductors
14	06071	CA	—	376,534	3.6	0.33	445	Fab metal / ammo
15	09009	CT	—	331,660	6.5	0.32	397	Fab metal / ammo
16	12099	FL	—	354,094	4.3	0.31	319	Engineering svc.
17	39085	OH	—	184,765	14.1	0.16	277	Fab metal / ammo
18	01097	AL	—	197,540	8.7	0.18	250	Ships/tanks/rail
19	42045	PA	—	152,667	4.9	0.14	207	Aerospace
20	06083	CA	—	215,663	8.1	0.19	187	Engineering svc.
21	53053	WA	—	141,632	3.7	0.13	176	Aerospace
22	39017	OH	—	120,969	6.0	0.11	174	Fab metal / ammo
23	48355	TX	—	123,851	5.7	0.11	170	Aerospace
24	09007	CT	—	96,099	13.6	0.09	163	Aerospace
25	53063	WA	—	146,423	5.1	0.13	157	Fab metal / ammo
26	18141	IN	—	96,518	6.1	0.09	130	Fab metal / ammo
27	48309	TX	—	82,192	5.9	0.07	127	Aerospace
28	29183	MO	—	99,026	5.3	0.09	117	Aerospace
29	49057	UT	—	87,645	7.9	0.08	116	Fab metal / ammo
30	49011	UT	—	115,265	7.8	0.10	113	Engineering svc.
31	12091	FL	—	123,282	12.2	0.11	105	Engineering svc.
32	09005	CT	—	70,636	9.0	0.06	99	Fab metal / ammo
33	21015	KY	—	64,913	6.4	0.06	99	Other machinery
34	26065	MI	—	101,331	7.0	0.09	95	Computer systems svc.
35	39109	OH	—	45,761	10.0	0.04	72	Other machinery
36	48485	TX	—	32,498	5.9	0.03	51	Fab metal / ammo
37	54033	WV	—	26,421	6.8	0.02	32	Aerospace
38	32510	NV	—	21,045	8.7	0.02	29	Fab metal / ammo
39	13313	GA	—	49,102	6.8	0.04	27	Computer systems svc.
40	05051	AR	—	15,632	3.5	0.01	22	Aerospace
41	20191	KS	—	7,912	14.2	0.01	14	Aerospace
42	26029	MI	—	7,657	9.4	0.01	10	Instruments
43	28087	MS	—	6,976	3.1	0.01	9	Aerospace
44	20015	KS	—	5,053	3.3	0.00	6	Engineering svc.
45	16017	ID	—	5,266	4.0	0.00	6	Engineering svc.
46	19045	IA	—	2,673	1.3	0.00	4	Fab metal / ammo
47	40041	OK	—	2,472	3.6	0.00	4	Fab metal / ammo

continued on next page

22, continued

Rank	FIPS	St	Hub label	Def emp	Def %	DoD %	V_2/M	Largest def. ind.
48	27161	MN	—	2,766	7.9	0.00	2	Communications eq.
49	53039	WA	—	2,036	3.6	0.00	2	Engineering svc.
50	01039	AL	—	1,436	1.5	0.00	2	Engineering svc.
51	21165	KY	—	17	2.8	0.00	0	Computer systems svc.
52	09110	CT	—	50,692	11.9	0.00	0	Aerospace
53	09160	CT	—	1,484	6.3	0.00	0	Fab metal / ammo
54	09130	CT	—	6,375	12.9	0.00	0	Aerospace
55	09140	CT	—	13,422	10.1	0.00	0	Fab metal / ammo
56	09170	CT	—	12,417	5.6	0.00	0	R&D services

Sorted by V_2 desc. CONUS-only; Hartford CT shown as legacy 09003 since the new planning regions (09110, etc.) have sparse QCEW history.

Virginia geographic harmonization

Virginia’s independent cities are not counties; the BEA input-output accounts combine each with an adjacent county under a synthetic combined FIPS, while the QCEW employment geography uses the raw Census city and county codes. For the structural constructions the panel carries the BEA-combined code; for the DoD-exposure (δ) weighting the combined share is split equally across the raw Census constituents so that the V_3 weighting reaches every defense-active Virginia unit. Table 23 is the full crosswalk applied (`build_w2.py`).

F Top-100 county rankings

Table 24 lists the 100 most vulnerable counties by V_2 (Leontief-propagated closed-loop vulnerability), and Table 25 the 100 highest by dual-use surge capacity. Together they are the full ranking behind the headline concentration and reserve-geography findings of Sections 6 and 8.

Table 24: Top 100 US counties by V_2 (Leontief-propagated defense industrial vulnerability).

Rank	FIPS	St	Hub	V_2/M	Def %	DoD %
1	06037	CA	Los Angeles CA	4,352.6	5.8	3.39
2	06085	CA	Santa Clara CA (Silicon Valley)	2,766.0	26.2	3.53
3	48201	TX	—	2,471.3	8.0	2.43
4	53033	WA	King WA (Seattle / Boeing)	2,462.2	11.8	1.98
5	06059	CA	Orange CA	2,140.0	9.2	1.92
6	17031	IL	—	2,077.8	5.4	1.85
7	06073	CA	San Diego CA	2,070.1	11.3	2.04
8	04013	AZ	Maricopa AZ (Phoenix)	1,889.0	6.8	1.71
9	25017	MA	Middlesex MA	1,740.7	18.2	2.18
10	26163	MI	Wayne MI (Detroit / Ford River Rouge)	1,704.0	10.9	0.98
11	53061	WA	Snohomish WA (Boeing Everett commercial)	1,642.0	25.8	0.83
12	48113	TX	Dallas TX	1,502.9	7.7	1.68

continued on next page

24, continued

Rank	FIPS	St	Hub	V_2/M	Def %	DoD %
13	48439	TX	Tarrant TX (Fort Worth / Lockheed F-35)	1,259.5	7.1	0.80
14	09003	CT	Hartford CT (P&W F135 — legacy FIPS)	1,211.1	12.0	0.76
15	26125	MI	—	1,188.5	13.3	1.27
16	20173	KS	Sedgwick KS (Wichita / Boeing)	1,187.5	17.9	0.57
17	06001	CA	—	1,052.9	12.2	1.17
18	27053	MN	—	1,001.9	8.6	1.02
19	26099	MI	Macomb MI (commercial auto)	946.8	14.4	0.58
20	39035	OH	Cuyahoga OH (Cleveland heavy industry)	941.5	7.7	0.73
21	17043	IL	—	836.3	10.3	0.83
22	48453	TX	—	759.7	13.1	1.10
23	36061	NY	—	718.3	4.8	1.46
24	51059	VA	Fairfax VA (defense IT)	704.4	20.5	1.50
25	40143	OK	—	654.1	9.1	0.41
26	25009	MA	Essex MA (Boston defense electronics)	653.8	11.2	0.45
27	39061	OH	—	653.2	7.5	0.50
28	41067	OR	—	633.5	17.5	0.62
29	49035	UT	—	600.0	7.5	0.65
30	42003	PA	Allegheny PA (Pittsburgh)	592.0	7.1	0.65
31	36055	NY	—	570.3	9.2	0.45
32	36103	NY	—	563.7	6.4	0.54
33	48029	TX	Bexar TX (San Antonio)	563.7	4.6	0.49
34	18097	IN	—	547.5	6.6	0.50
35	42091	PA	—	541.2	9.7	0.63
36	55079	WI	—	530.0	6.2	0.39
37	55133	WI	—	528.9	11.9	0.37
38	01089	AL	Madison AL (Huntsville)	515.5	29.4	0.61
39	26081	MI	—	503.4	6.8	0.34
40	53035	WA	—	488.2	24.9	0.25
41	12009	FL	Brevard FL (Space Coast)	468.5	16.9	0.43
42	06075	CA	—	464.8	11.9	0.97
43	06081	CA	—	461.9	13.7	0.66
44	37119	NC	—	449.9	5.7	0.48
45	06071	CA	—	445.4	3.6	0.33
46	12086	FL	—	444.4	3.2	0.45
47	24031	MD	—	439.2	12.1	0.69
48	39049	OH	—	435.5	5.5	0.50
49	36029	NY	—	421.8	6.4	0.37
50	29189	MO	St Louis County MO (Boeing)	416.8	5.5	0.43
51	37183	NC	—	413.0	9.5	0.62
52	12095	FL	—	406.4	5.0	0.50
53	17201	IL	—	398.1	14.5	0.21
54	09009	CT	—	396.9	6.5	0.32
55	48085	TX	—	392.6	11.1	0.57
56	25027	MA	—	392.1	8.3	0.33
57	11001	DC	DC (Pentagon)	382.0	9.2	0.70
58	06065	CA	—	365.3	3.1	0.27
59	06111	CA	—	363.9	7.8	0.32
60	27003	MN	—	363.7	17.6	0.25
61	12011	FL	—	357.0	3.8	0.37

continued on next page

24, continued

Rank	FIPS	St	Hub	V_2/M	Def %	DoD %
62	51740	VA	—	355.6	37.3	0.17
63	39113	OH	—	355.5	8.2	0.26
64	39153	OH	—	350.0	7.6	0.25
65	12103	FL	—	348.7	5.8	0.31
66	42133	PA	—	346.7	9.7	0.21
67	12057	FL	—	343.6	5.5	0.46
68	17097	IL	—	341.1	6.6	0.29
69	08013	CO	—	334.0	20.4	0.44
70	33011	NH	—	330.9	13.9	0.33
71	40109	OK	—	327.4	4.9	0.27
72	29095	MO	—	325.8	8.0	0.36
73	12099	FL	—	319.2	4.3	0.31
74	09001	CT	—	316.6	5.7	0.30
75	35001	NM	—	310.2	10.0	0.38
76	27123	MN	—	309.4	7.7	0.29
77	21111	KY	—	305.4	5.0	0.27
78	12031	FL	Duval FL (Jacksonville Navy)	303.5	5.2	0.33
79	34023	NJ	—	303.0	8.8	0.45
80	41051	OR	—	301.8	5.7	0.34
81	45045	SC	—	301.5	7.6	0.24
82	13121	GA	—	292.9	5.5	0.55
83	15003	HI	—	291.7	5.1	0.28
84	25025	MA	—	287.9	6.1	0.50
85	42017	PA	—	285.5	6.7	0.23
86	39085	OH	—	276.5	14.1	0.16
87	34003	NJ	—	271.8	5.5	0.31
88	37063	NC	—	271.8	19.1	0.37
89	06067	CA	—	271.7	5.3	0.38
90	18039	IN	—	271.1	9.6	0.15
91	17089	IL	—	270.9	7.4	0.18
92	34027	NJ	—	265.8	8.4	0.31
93	13135	GA	—	264.2	7.3	0.33
94	55025	WI	—	260.2	7.7	0.28
95	36059	NY	—	250.3	3.0	0.23
96	01097	AL	—	250.1	8.7	0.18
97	39095	OH	—	247.8	5.8	0.14
98	20091	KS	—	247.8	7.7	0.34
99	25021	MA	—	244.3	6.1	0.27
100	26139	MI	—	238.1	11.4	0.15

Companion to Table 4 in the main paper. Lower-ranked counties are progressively lower-vulnerability but still defense-relevant.

Table 25: Top 100 US counties by dual-use surge capacity.

Rank	FIPS	St	Hub	Surge	Top ind.	V_2/M	Quadrant
1	06037	CA	Los Angeles CA	8,290,285	722	4,353	self rescuable
2	17031	IL	—	5,071,020	61	2,078	self rescuable
3	48201	TX	—	4,480,466	23	2,471	self rescuable
4	36061	NY	—	4,408,686	523	718	self rescuable
5	04013	AZ	Maricopa AZ (Phoenix)	3,806,622	5613	1,889	self rescuable

continued on next page

25, continued

Rank	FIPS	St	Hub	Surge	Top ind.	V ₂ /M	Quadrant
6	48113	TX	Dallas TX	3,346,373	5613	1,503	self rescuable
7	06059	CA	Orange CA	3,150,164	722	2,140	self rescuable
8	06073	CA	San Diego CA	2,678,732	722	2,070	self rescuable
9	53033	WA	King WA (Seattle / Boeing)	2,549,881	722	2,462	self rescuable
10	06085	CA	Santa Clara CA (Silicon Valley)	2,025,160	5415	2,766	self rescuable
11	12086	FL	—	2,023,074	722	444	self rescuable
12	25017	MA	Middlesex MA	1,855,601	61	1,741	self rescuable
13	27053	MN	—	1,843,793	55	1,002	self rescuable
14	32003	NV	Clark NV (Nellis)	1,716,696	721	197	self rescuable
15	48439	TX	Tarrant TX (Fort Worth / Lockheed F-35)	1,640,046	722	1,260	self rescuable
16	48029	TX	Bexar TX (San Antonio)	1,573,563	722	564	self rescuable
17	13121	GA	—	1,567,463	722	293	self rescuable
18	26125	MI	—	1,533,803	5413	1,188	self rescuable
19	12095	FL	—	1,506,350	713	406	self rescuable
20	12011	FL	—	1,493,397	722	357	self rescuable
21	06001	CA	—	1,460,805	61	1,053	self rescuable
22	39035	OH	Cuyahoga OH (Cleveland heavy industry)	1,437,941	622	941	self rescuable
23	39049	OH	—	1,408,736	722	435	self rescuable
24	06071	CA	—	1,406,736	61	445	self rescuable
25	42003	PA	Allegheny PA (Pittsburgh)	1,380,428	61	592	self rescuable
26	26163	MI	Wayne MI (Detroit / Ford River Rouge)	1,331,743	622	1,704	self rescuable
27	49035	UT	—	1,316,709	23	600	self rescuable
28	48453	TX	—	1,313,455	722	760	self rescuable
29	37119	NC	—	1,309,992	722	450	self rescuable
30	06065	CA	—	1,291,994	23	365	self rescuable
31	17043	IL	—	1,288,086	5613	836	self rescuable
32	12057	FL	—	1,283,614	722	344	self rescuable
33	36103	NY	—	1,274,548	61	564	self rescuable
34	36047	NY	—	1,271,460	6215OH	121	self rescuable
35	06075	CA	—	1,253,701	5415	465	self rescuable
36	51059	VA	Fairfax VA (defense IT)	1,236,278	5415	704	self rescuable
37	29189	MO	St Louis County MO (Boeing)	1,198,609	55	417	self rescuable
38	25025	MA	—	1,191,265	622	288	self rescuable
39	42101	PA	—	1,172,236	622	221	self rescuable
40	18097	IN	—	1,164,203	5613	548	self rescuable
41	11001	DC	DC (Pentagon)	1,155,093	813	382	self rescuable
42	36059	NY	—	1,140,984	61	250	self rescuable
43	36081	NY	—	1,129,066	23	136	self rescuable
44	12099	FL	—	1,111,950	722	319	self rescuable
45	06067	CA	—	1,068,976	722	272	self rescuable
46	39061	OH	—	1,031,352	55	653	self rescuable
47	42091	PA	—	1,028,246	722	541	self rescuable
48	37183	NC	—	1,012,565	722	413	self rescuable
49	12031	FL	Duval FL (Jacksonville Navy)	963,906	722	304	self rescuable
50	55079	WI	—	945,050	622	530	self rescuable
51	47157	TN	—	941,475	5613	226	self rescuable
52	08031	CO	Denver CO (Lockheed)	933,772	722	225	self rescuable
53	09003	CT	Hartford CT (P&W F135 — legacy FIPS)	921,751	524	1,211	self rescuable
54	41051	OR	—	919,257	722	302	self rescuable
55	47037	TN	—	883,467	722	160	self rescuable

continued on next page

25, continued

Rank	FIPS	St	Hub	Surge	Top ind.	V ₂ /M	Quadrant
56	34003	NJ	—	880,771	61	272	self rescuable
57	24031	MD	—	878,667	5415	439	self rescuable
58	36029	NY	—	875,046	722	422	self rescuable
59	21111	KY	—	852,243	722	305	self rescuable
60	12103	FL	—	825,569	722	349	self rescuable
61	34023	NJ	—	823,528	5613	303	self rescuable
62	40109	OK	—	815,666	722	327	self rescuable
63	36119	NY	—	799,301	61	165	self rescuable
64	15003	HI	—	797,415	722	292	self rescuable
65	48085	TX	—	785,864	5415	393	self rescuable
66	26081	MI	—	779,119	5613	503	self rescuable
67	09001	CT	—	757,677	722	317	self rescuable
68	06081	CA	—	738,899	5415	462	self rescuable
69	36055	NY	—	722,106	61	570	self rescuable
70	13135	GA	—	718,142	722	264	self rescuable
71	20091	KS	—	689,394	722	248	self rescuable
72	25021	MA	—	686,018	722	244	self rescuable
73	17097	IL	—	685,630	61	341	self rescuable
74	29095	MO	—	682,591	722	326	self rescuable
75	40143	OK	—	679,579	722	654	self rescuable
76	24005	MD	—	672,069	23	163	self rescuable
77	09009	CT	—	671,419	61	397	self rescuable
78	06013	CA	—	660,331	722	154	self rescuable
79	31055	NE	—	659,327	722	123	self rescuable
80	13067	GA	Cobb GA (Lockheed Marietta)	658,824	23	178	self rescuable
81	01073	AL	—	640,078	722	177	self rescuable
82	06019	CA	—	636,602	61	113	self rescuable
83	04019	AZ	Pima AZ (Tucson Raytheon)	630,000	722	203	self rescuable
84	08005	CO	—	619,231	23	227	self rescuable
85	26099	MI	Macomb MI (commercial auto)	618,567	3362BP	947	self rescuable
86	25027	MA	—	618,076	722	392	self rescuable
87	25009	MA	Essex MA (Boston defense electronics)	600,384	722	654	self rescuable
88	27123	MN	—	590,313	55	309	self rescuable
89	06111	CA	—	588,081	722	364	self rescuable
90	35001	NM	—	579,460	722	310	self rescuable
91	34027	NJ	—	572,830	61	266	self rescuable
92	55025	WI	—	568,655	722	260	self rescuable
93	19153	IA	—	565,408	524	121	self rescuable
94	34013	NJ	—	563,198	722	147	self rescuable
95	24510	MD	—	551,789	622	110	self rescuable
96	41067	OR	—	541,468	3344	634	self rescuable
97	44007	RI	—	540,416	61	159	self rescuable
98	39153	OH	—	525,642	55	350	self rescuable
99	42017	PA	—	519,677	722	285	self rescuable
100	48141	TX	El Paso TX (Fort Bliss)	516,659	722	76	self rescuable

Companion to Table 10 in the main paper. Top industries are mostly large-metro services and manufacturing — high surge counties are concentrated in major commercial/industrial centers.

Table 23: Virginia BEA-combined FIPS to raw Census constituents. The combined share is redistributed equally across constituents in the DoD-exposure weighting; all other Virginia units enter the analysis under their raw Census FIPS unchanged.

BEA-combined FIPS	Raw Census constituents
51901	51003, 51540
51903	51005, 51580
51907	51015, 51790, 51820
51911	51031, 51680
51913	51035, 51640
51918	51053, 51570
51919	51059, 51600, 51610
51921	51069, 51840
51923	51081, 51595
51929	51089, 51690
51931	51095, 51830
51933	51121, 51750
51939	51143, 51590
51941	51149, 51670, 51730
51942	51153, 51683, 51685
51944	51161, 51775
51945	51163, 51530, 51678
51947	51165, 51660
51949	51175, 51620
51951	51177, 51630
51953	51191, 51520
51955	51195, 51720
51958	51199, 51735

Twenty-three BEA-combined Virginia FIPS; all remaining Virginia counties and independent cities enter under their raw Census FIPS.

G Reduced-form supporting analysis

The paper’s contribution is the descriptive structural geography of Sections 6–8; it is not a causal-effects study. A companion reduced-form analysis—a lean spatial Durbin model of county output on DoD obligations, and a difference-in-differences on the 2022 defense-spending shock—was nonetheless run as supporting, not headline, evidence. It is weak by construction and not relied upon: the local elasticity is small, and a Bartik shift-share instrument is too weak—a Kleibergen–Paap rk Wald F of essentially zero (≈ 0.001 , against the Stock–Yogo 10% threshold of 7.03), and even the weak-instrument-robust Anderson–Rubin test rejects nothing ($p \approx 0.43$)—to identify a spillover multiplier, so no causal claim is made. The weakness is structural, not incidental: a variance decomposition of the constructed instrument attributes 71% of its variation to the common national (shift) component shared by all counties and effectively zero to persistent cross-county (share) differences, leaving the remainder as county-by-year idiosyncrasy. Because the specification’s county and year fixed effects absorb exactly those first two margins, almost none of the instrument’s systematic variation survives demeaning—a construction-level reason, distinct from the economics, that the first stage is dead. Once a set of in-panel confounders is added, the difference-in-differences effect attenuates to zero and the placebo and pre-trend checks clear—but several of those controls (county employment, establishment counts, manufacturing share, and county-specific trends extrapolated through the treatment window) are themselves post-treatment outcomes of a defense-spending shock, so the attenuation cannot distinguish a true null from over-control. We therefore treat the design as *not identified* rather than as a clean null, and lead the null reading through a synthetic-DiD specification that avoids such controls, which returns a precise null for manufacturing output and employment (ATT ≈ 0.00 and $+0.01$, neither significant) and only a small negative wage effect (ATT ≈ -0.01); the baseline two-way fixed-effects estimate, for its part, fails its own pre-trend and 2018-placebo tests. The estimates are consistent with—but do not establish—the structural argument; the corrected specifications are documented with the project’s Stata scripts. The one reduced-form result reported in full here is the spatial-weights sensitivity of Table 26: across five \mathbf{W} specifications the direct elasticity is invariant (about \$0.0002 per dollar) while the implied spillover multiplier is not point-identified—it ranges over two orders of magnitude across the matrices, flips sign when Virginia’s combined-FIPS counties are correctly matched into the panel, and rests for W_3 on a network that leaves most counties isolated—so it is the propagation-network assumption, not the local coefficient, that drives and ultimately fails to identify any spillover estimate. The full SDM and difference-in-differences point-estimate tables are reported separately.

A final caveat concerns the inverse-hyperbolic-sine transform applied to the spending regressors. Because $\operatorname{asinh}(cx) \approx \ln(2c) + \ln(x)$ for large x , the asinh coefficient is not invariant to the units in which spending is measured (Bellemare and Wichman, 2020); a within-2SLS sensitivity replication confirms that re-scaling DoD obligations from dollars to billions moves the estimated coefficient by roughly a factor of forty, while leaving it economically small at every scale. The transform is nonetheless the appropriate choice for this regressor: net obligations are zero for about twenty-eight percent of county-years and negative for some (de-obligations exceeding new awards), values on which the logarithmic alternatives are undefined. The reduced-form elasticity is therefore read qualitatively—as directional corroboration of the structural argument—rather than as an identified point estimate. The structural indices of Sections 6–8, built from employment and obligation *shares*, are scale-free by construction and carry none of this sensitivity; the full scale ladder is reported in the replication script `test_asinh_scale.py`.

Table 26: SDM estimates across five spatial-weights specifications, three outcomes, Model 1 (Tier B with controls). The direct (local) β is stable across \mathbf{W} choices for each outcome, but the implied total effect (spillover multiplier) is reported only where the spatial autoregressive parameter is stable ($|\rho| < 1$); where it is not, the multiplier formula is undefined and the cell is left blank. Among the ρ -stable cells the multiplier is small and network-dependent—a slight negative under the defense-restricted I-O matrix and essentially zero under queen contiguity and observed sub-awards for output and income—confirming that the spillover is governed by the network choice and is not point-identified, while the local elasticity is invariant. The education rows are shown for completeness only: their ρ sits at the unit-root boundary, where the near-singular $1 - \rho$ renders the reported total uninterpretable, compounded by ACS five-year pooling.

Outcome	W	Direct β	ρ	θ	Total effect	R^2
GDP	Queen contiguity	0.00024	+0.159	-0.0000	+0.0002	0.174
GDP	W ₂ V1 (general I-O)	0.00024	-3.064	-0.0529	— (ρ unst.)	0.114
GDP	W ₂ V2 (defense I-O)	0.00024	+0.136	-0.0114	-0.0129	0.109
GDP	W ₂ V3 (defense-exposure)	0.00024	-3.462	+0.0656	— (ρ unst.)	0.110
GDP	W ₃ (observed subawards)	0.00022	+0.229	+0.0012	+0.0018	0.109
GDP per capita	Queen contiguity	0.00021	+0.471	-0.0000	+0.0004	0.152
GDP per capita	W ₂ V1 (general I-O)	0.00023	-2.720	-0.0676	— (ρ unst.)	0.010
GDP per capita	W ₂ V2 (defense I-O)	0.00024	+0.002	-0.0132	-0.0130	0.005
GDP per capita	W ₂ V3 (defense-exposure)	0.00024	-1.386	+0.0737	— (ρ unst.)	0.006
GDP per capita	W ₃ (observed subawards)	0.00022	+0.158	+0.0013	+0.0018	0.004
Education	Queen contiguity	0.00250	+0.972	-0.0039	-0.0501	0.005
Education	W ₂ V1 (general I-O)	0.00326	+1.188	+1.6426	— (ρ unst.)	0.066
Education	W ₂ V2 (defense I-O)	0.00322	+0.979	+0.1570	+7.7610	0.067
Education	W ₂ V3 (defense-exposure)	0.00322	+4.531	-0.9994	— (ρ unst.)	0.067
Education	W ₃ (observed subawards)	0.00245	+0.143	+0.0222	+0.0287	0.075

Within-2SLS estimation with two-way fixed effects (Lee & Yu 2010); Kelejian-Prucha 1998 IV strategy with W^2X as instruments for the endogenous Wy spatial lag. Direct β = local coefficient on $\text{asinh}(\text{DoD obligations})$. ρ = spatial autoregressive coefficient. θ = WX coefficient. Total = $(\beta + \theta)/(1 - \rho)$. Cross-outcome balanced panel: 3,070 counties \times 13 years = 39,910 obs.