ElementMindX: Offline Supplier-Substitution Ranking for Natural-Language Trade-Shock Decision Support

Anonymous Authors[1]

Abstract

Trade disruptions, whether from tariffs, export bans, or production shocks force analysts to make high-stakes sourcing decisions under deep uncertainty. We present ELEMENTMINDX, an interactive decision-support tool that translates natural-language scenario descriptions into counterfactual simulations over global trade graphs. While the deployed system supports materials ranging from silicon and tantalum to energy products, this paper evaluates the supplier-ranking module on copper. We approach supplier substitution as an offline learning-to-rank task: using historical exposure events in observational trade panels, we train models to rank plausible alternative suppliers based on observed next-period market-share gains. ELEMENTMINDX then pairs these learned rankings with a transparent heuristic allocator to visualize volume rerouting and estimate shortages. In a pilot study of 1,277 copper importer-shock groups from 1992–2022, the learned rankers improve over global-share and incumbent-share baselines on Hit@1 and NDCG@10 across three copper HS6 markets.

1. Motivation and Fit

Supply-chain analysts often need rapid answers to questions such as: *Which importers are exposed if a supplier cuts exports? Which alternative exporters are plausible substitutes? How much disrupted flow may remain unmet?* Classical equilibrium, gravity (Anderson and van Wincoop, 2003), global value-chain (Antràs and Chor, 2022), and input-output network models (Acemoglu et al., 2012) are valuable, but are often too slow or assumption-heavy for interactive scenario exploration. Empirical disruption studies show how shocks propagate through production networks (Carvalho et al., 2021), and critical-mineral supply risks are central to clean-energy technologies (IEA, 2021). Purely heuristic dashboards are also risky when they rank suppliers only by global scale or incumbency.

ELEMENTMINDX (ElementMindX, 2026).targets offline decision-making from logged data (Levine et al., 2020). Online exploration is not realistic: governments cannot impose random tariffs to explore substitution, and firms cannot randomize supplier failures. Historical trade panels also lack randomized logging propensities, so we do not claim inverse-propensity off-policy evaluation. Instead, we use the single-step decision structure of contextual bandits (Swaminathan and Joachims, 2015) only to organize the problem: contexts are importer-shock states, candidate exporters are actions, and observed positive share gains provide relevance. The implemented task is supervised learning-to-rank (Joachims, 2002; Liu, 2009; Burges, 2010) from observational data.

Contributions. First, we frame trade-shock response as an *offline supplier discovery* problem, ranking plausible alternative suppliers rather than predicting exact, unobservable private contracts. Second, we implement a modular shock engine that separates deterministic disruption mechanics (edge reduction and shortage accounting) from the machine-learning task (ranking) and the downstream decision-support visualization (heuristic allocation). Third, we empirically validate this approach using historical data from three distinct copper markets. Finally, we show how the deployed system architecture scales to a broader set of critical commodities through HS6 registry routing and commodity-specific panels.

2. System and Shock Engine

ELEMENTMINDX receives a scenario, converts it to a typed `ShockSpec`, executes a shock primitive, and returns headline metrics, winners/losers, edge deltas, charts, and explanations. The production schema exposes `tariff`, `exporter_shock`, and `edge_reduction`. The deployed interface presents material-specific browsing and execution paths for copper, silicon, tantalum, neodymium proxies, and energy products. The registry routes HS6 codes

[1]Anonymous Institution, Anonymous City, Anonymous Region, Anonymous Country. Correspondence to: Anonymous Author <anon.email@domain.com>.

Preliminary work. Under review by the International Conference on Machine Learning (ICML). Do not distribute.

to commodity-specific panels: copper {260300, 740311, 740400}; energy/gasoline {270900, 271012, 271019, 271112, 271121}; neodymium proxy {284610, 284690}; silicon {280461, 280469, 720221, 720229}; and tantalum {261590, 810320, 810330, 810391}. The code-reviewed free-text parser path is strongest for copper; non-copper scenarios can be executed through element-specific UI/HS6 routing.

Figure 1 shows the deployed multi-commodity portal and an executed non-copper scenario.

Let $v_{i,x,h,t} \geq 0$ denote annual trade from exporter x to importer i for HS6 product h in year t. A shock with severity σ removes

$$L_{i,x,h,t} = \sigma v_{i,x,h,t}, \qquad L = R + S, \tag{1}$$

where R is rerouted volume and S is shortage. Dependence share is

$$d_{i,x,h,t} = \frac{v_{i,x,h,t}}{\sum_{x'} v_{i,x',h,t}}. \tag{2}$$

The shortage fraction uses monotone bins: 0.10 for $d < 0.25$, 0.20 for $0.25 \leq d < 0.50$, 0.30 for $0.50 \leq d < 0.75$, 0.40 for $d \geq 0.75$, and 0.25 if missing, multiplied by an HS6 scale and clipped to $[0, 0.95]$. The current non-unit scale is 1.15 for copper ores/concentrates (HS6 260300); other supported HS6 codes use 1.0.

For exporter shocks and edge reductions, the candidate pool is the union of incumbents to importer i, top-10 global exporters for h, top-5 exporters in the importer's continent, and top-10 peer suppliers serving other importers in the same continent. The shocked exporter is removed and the pool is capped at 30. Each per-HS6 ranker scores candidates; the top k are converted to allocation weights via

$$w_c = \frac{\exp(f_h(c)/\tau)}{\sum_{c'} \exp(f_h(c')/\tau)}. \tag{3}$$

Unless overridden, the backend uses `min_edge`=100,000, `topk_alloc`=10, τ = 1.0, capacity cap `cap_frac`=0.05, new-edge cap `new_edge_cap`=\$50M, and `lookback_years`=3. These caps are visualization guardrails; they are not part of the supervised ranking evaluation.

3. Offline Supplier-Substitution Ranker

The experiments use processed annual bilateral HS6 trade panels derived from UN Comtrade data (UN Comtrade, 2024), with `period`, `hs6`, `importer_iso3`, `exporter_iso3`, and `trade_value_usd` fields. The copper panel contains 85,119 nonzero annual bilateral rows from 1988-2023 across HS6 `260300`, `740311`, and `740400`; it contains 241 importers and 199 exporters. We report panel coverage directly from processed files; upstream extraction, reporter/mirror-flow convention, HS concordance, and re-export handling are reproducibility items for the extended version.

The balanced copper event file contains 150 observed year-over-year exporter supply-drop events, with 50 per HS6. Recomputing from the uploaded copper panel shows these are consistent with the top-50 exporter-HS6-year drops per HS6 by absolute lost global exporter value among year-over-year global exporter declines of at least 40%. After importer exposure filtering and candidate construction, the supplier-ranking table contains 27,924 candidate rows, 1,277 importer-shock groups, and 133 event IDs across 1992-2022.

For an observed supply-drop event involving shocked exporter x, we measure candidate substitution using next-period import-share change among non-shocked suppliers. Let

$$V^{(-x)}_{i,h,u} = \sum_{c \neq x} v_{i,c,h,u}, \quad s^{(-x)}_{i,c,h,u} = v_{i,c,h,u}/V^{(-x)}_{i,h,u} \tag{4}$$

for $u \in \{t, t+1\}$. The signed outcome is $\Delta s_{i,c,h,t} = s^{(-x)}_{i,c,h,t+1} - s^{(-x)}_{i,c,h,t}$ and the ranking relevance is

$$r_{i,c,h,t} = \max\{0, \Delta s_{i,c,h,t}\}. \tag{5}$$

Features include incumbency, supplied-recently, same-continent, pre-share, candidate global share, severity, baseline exposure, importer total, pre-value, candidate global value, share gap, and log transforms. The saved copper models use XGBoost (Chen and Guestrin, 2016) with a `rank:pairwise` objective.

4. Copper Empirical Summary

We evaluate supplier discovery using ranking metrics standard in information retrieval and learning-to-rank: Hit@1, Hit@3, and NDCG@10 (Liu, 2009; Järvelin and Kekäläinen, 2002). Baselines are Global Share (GS), ranking by candidate global HS6 export share, and Incumbent Share (IS), ranking by pre-shock share with the importer. For group g, let $r_{g,c}$ be relevance and $r^*_g = \max_c r_{g,c}$. We define

$$\text{Hit@}k(g) = \mathbf{1}[r^*_g > 0 \wedge \max_{c \in \text{Top}_k(g)} r_{g,c} = r^*_g], \tag{6}$$

counting groups with no positive-gain candidate as zero-hit groups. NDCG@10 uses continuous relevance values with discounted rank positions, rewarding methods that place higher-gain suppliers earlier in the shortlist (Järvelin and Kekäläinen, 2002).

Table 1 provides pilot aggregate evidence: learned per-HS6 rankers improve over both baselines on Hit@1 and

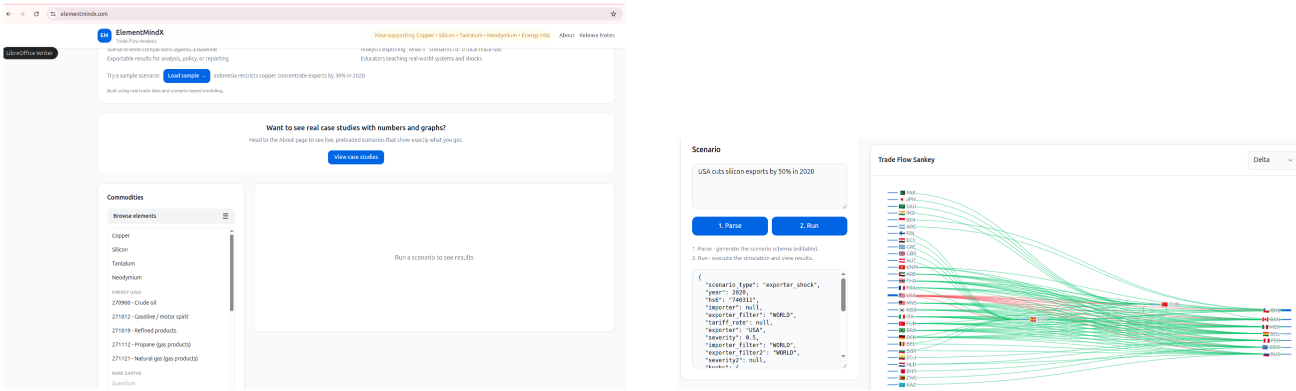

(a) Portal overview with the commodity browser and support banner.

(b) Silicon exporter-shock scenario with parsed specification and Sankey output.

Figure 1. Screenshots from the working ElementMindX portal. The deployed interface supports multi-commodity browsing and scenario execution; this short submission reports completed empirical ranking results for copper.

Table 1. Copper supplier-discovery pilot. The learned ranker improves Hit@1 and NDCG@10 over global-share and incumbent-share baselines across all three copper HS6 codes; Incumbent Share is slightly higher on Hit@3 for HS6 740311.

	Hit@1			Hit@3			NDCG@10		
HS6	GS	IS	Ours	GS	IS	Ours	GS	IS	Ours
260300	.249	.254	**.359**	.580	.575	**.590**	.544	.562	**.664**
740311	.162	.249	**.443**	.404	**.614**	.602	.476	.598	**.682**
740400	.195	.211	**.338**	.523	.522	**.600**	.533	.551	**.659**

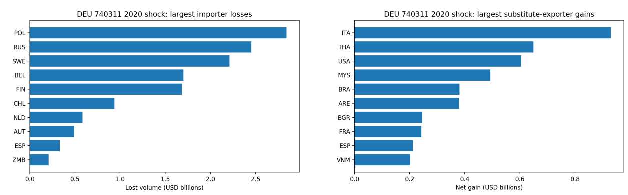

(a) Importer losses.

(b) Exporter gains.

Figure 2. Germany–HS740311–2020 case-study plots. Values are rounded to billions of USD.

NDCG@10 for all three copper markets, while also improving Hit@3 for two of the three HS6 codes. The available artifact provides aggregate point estimates but not per-group prediction scores, so we do not report learned-ranker bootstrap intervals in this short submission. The extended version will release per-group prediction logs, split metadata, and confidence intervals under a fully specified protocol.

5. Case Study: Germany Refined Copper Shock

To demonstrate the end-to-end decision workflow, we simulate a 50% exporter shock for Germany (DEU) on refined copper cathodes, HS6 `740311`, in 2020 using the current engine and parser-default candidate pruning. The scenario affects 27 importer episodes. Baseline exposed trade is 28.90B USD, so the shock removes 14.45B USD. The conserved run allocates 5.48B USD to alternative suppliers and records 8.97B USD as first-year shortage. Top substitute exporters by net gain are Italy (0.93B USD), Thailand (0.65B), the United States (0.60B), Malaysia (0.49B), and Brazil (0.38B). The case illustrates how the learned ranker supplies a shortlist, while the allocator turns this shortlist into auditable flow and shortage summaries.

6. Discussion and Limitations

It is important to note that ElementMindX is not a calibrated equilibrium model. It learns supplier shortlists from observational trade data, using heuristic allocation only to make those ranked recommendations visually actionable. This design choice is intentional: public trade panels do not reveal private contracts, inventories, spot-market access, or unrealized supplier capacity. Consequently, our dependence bins, capacity caps, and tariff elasticities serve as transparent, auditable guardrails rather than learned parameters.

Furthermore, the engine simulates static, one-period counterfactuals; it does not currently model dynamic price feedback, multi-year contracts, or shipping route capacities.

Finally, we intentionally separate validated empirical evidence from broader system availability. While the deployed portal supports multi-commodity scenario execution, this paper scopes its formal ranking evaluation to copper. Generating the corresponding event files and baselines for non-copper elements is left for future work. Ultimately, ElementMindX demonstrates a practical approach to offline decision support: we use machine learning only for the supplier-discovery task that observational logs can reliably supervise, while handling capacity and flow through explicit, user-auditable assumptions.

7. Conclusion

ElementMindX frames trade-shock response as offline supplier discovery: learn a ranked shortlist from historical trade panels, then use explicit heuristic allocation for interpretable visualization. The copper empirical module shows that learned rankers improve over global-share and incumbent-share baselines, while the multi-commodity registry and trained per-HS6 bundles provide concrete system support for broader HS6 execution.

References

ElementMindX. ElementMindX: AI-powered trade-shock and critical-material supply-chain decision support. https://elementmindx.com/, accessed 2026.

Acemoglu, D., Carvalho, V. M., Ozdaglar, A., and Tahbaz-Salehi, A. The network origins of aggregate fluctuations. *Econometrica*, 80(5):1977–2016, 2012.

Anderson, J. E. and van Wincoop, E. Gravity with gravitas: A solution to the border puzzle. *American Economic Review*, 93(1):170–192, 2003.

United Nations Statistics Division. UN Comtrade Database. https://comtradeplus.un.org/, accessed 2024.

Antràs, P. and Chor, D. Global value chains. In *Handbook of International Economics*, volume 5, pp. 297–376. Elsevier, 2022.

Burges, C. J. C. From RankNet to LambdaRank to LambdaMART: An overview. *Microsoft Research Technical Report*, MSR-TR-2010-82, 2010.

Carvalho, V. M., Nirei, M., Saito, Y. U., and Tahbaz-Salehi, A. Supply chain disruptions: Evidence from the Great East Japan Earthquake. *The Quarterly Journal of Economics*, 136(2):1255–1321, 2021.

Chen, T. and Guestrin, C. XGBoost: A scalable tree boosting system. In *KDD*, pp. 785–794, 2016.

Hidalgo, C. A., Klinger, B., Barabási, A.-L., and Hausmann, R. The product space conditions the development of nations. *Science*, 317(5837):482–487, 2007.

International Energy Agency. The role of critical minerals in clean energy transitions. IEA, Paris, 2021.

Järvelin, K. and Kekäläinen, J. Cumulated gain-based evaluation of IR techniques. *ACM Transactions on Information Systems*, 20(4):422–446, 2002.

Joachims, T. Optimizing search engines using clickthrough data. In *KDD*, pp. 133–142, 2002.

Levine, S., Kumar, A., Tucker, G., and Fu, J. Offline reinforcement learning: Tutorial, review, and perspectives on open problems. *arXiv:2005.01643*, 2020.

Liu, T.-Y. Learning to rank for information retrieval. *Foundations and Trends in Information Retrieval*, 3(3):225–331, 2009.

Swaminathan, A. and Joachims, T. Counterfactual risk minimization: Learning from logged bandit feedback. In *ICML*, pp. 814–823, 2015.

A. Additional Portal Screenshots

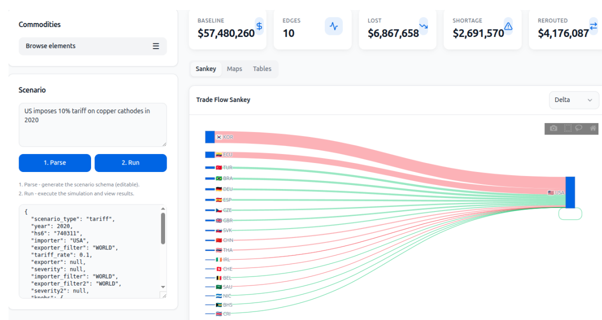

(a) Sankey view.

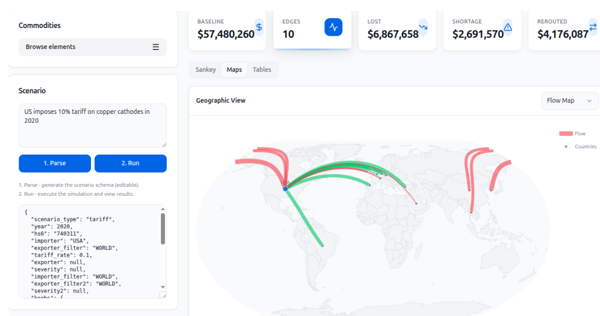

(b) Map view.

Figure 3. Additional portal views for a copper tariff scenario, illustrating KPI reporting and multiple result modalities.