АКАДЕМІЧНІ ВІЗІЇ
Випуск 48/2026

# Кількісне оцінювання резильєнтності малого виробництва персоналізованої продукції в умовах складених транскордонних збоїв: індекс складеної резильєнтності (CDR) та ранжування заходів пом'якшення за витратною ефективністю

Є.В. Шоломій[1]


| Опубліковано | Секція | УДК |
|---|---|---|
| 29.10.2025 | Соціальні і поведінкові науки | 005.94:339.9:330.341.2 |

DOI: https://doi.org/10.5281/zenodo.22053963

Анотація. Мікропідприємство, яке виготовляє персоналізовану продукцію в одній країні та продає її в іншій, стикається зі збоями різної природи: відключеннями електроенергії, зривами транскордонної логістики, затримками постачальників, закриттям каналу збуту на маркетплейсі та тарифними шоками. Якісна література з резильєнтності визначає, які здатності мають значення, але не вказує, який саме захід слід придбати першим за обмеженого бюджету; кількісна література пропонує формальні моделі розподілу ресурсів, проте передбачає наявність багаторівневої мережі, придатної для оцінювання історії збоїв та аналітичної спроможності, якої мікропідприємство не має. У статті запропоновано модель складеної резильєнтності (Compound Disruption Resilience, CDR), що заповнює цю прогалину. Модель приймає річний маржинальний дохід за нумерарій, завдяки чому різнорідні шоки стають порівнюваними та підсумовуються в єдиний функціонал очікуваних втрат; вона розділяє діагностику, здійснювану через композитний індекс резильєнтності RI, який зводиться до аудитованої тотожності — зваженого за експозицією середнього ефективності заходів, — і припис, здійснюваний через коефіцієнт ефективності пом'якшення MER, що ранжує заходи за очікуваними річними відверненими втратами на одиницю річних витрат на захід, із порогом упровадження MER ≥ 1. Модель відкалібровано на операції УФ-друку, яка виробляє в Україні для ринку Сполучених Штатів. Чотири з п'яти ймовірностей збоїв є практичними апріорними оцінками автора, а не частотами, обчисленими з журналу подій; аналіз чутливості кількісно оцінює їхній вплив. Повний портфель із п'яти заходів досягає RI = 0,716 і відвертає 41 588 дол. США очікуваних річних втрат за 17 900 дол. США річних витрат; застосовуючи власний поріг упровадження моделі, оператор натомість зупиняється на 13 700 дол. США та RI = 0,648, відвертаючи 37 661 дол. США, оскільки п'ятий захід порогу не проходить. Центральний результат полягає в тому, що впорядкування за MER розходиться як із ранжуванням за непом'якшеним впливом, так і з ранжуванням за абсолютною величиною відвернених втрат: режим із найбільшим впливом посідає останнє місце і не проходить поріг. Спільне моделювання методом Монте-Карло за всіма п'ятьма ймовірностями розміщує RI у 90-відсотковому інтервалі [0,694; 0,742] за повного носія 0,672–0,762 і залишає найвище ранжований захід незмінним у 77%

[1] незалежний дослідник, сфера - виробництво персоналізованої продукції та транскордонна електронна комерція; orcid.org/0009-0000-3916-5564; e-mail: evgeniysholomiy@gmail.com

реалізацій. Внесок роботи полягає не у твердженні, що ймовірності збоїв виміряно, а в демонстрації того, що для цього рішення та за цієї калібровки таке вимірювання не є необхідним.

Ключові слова: резильєнтність ланцюга постачання; ризик збоїв; мікропідприємство; очікувані втрати; портфель заходів пом'якшення; коефіцієнт «вигоди–витрати»; транскордонна електронна комерція; тарифний шок; друк на замовлення; аналіз чутливості

# QUANTIFYING RESILIENCE IN SMALL PERSONALIZED-PRODUCT MANUFACTURING UNDER COMPOUND CROSS-BORDER DISRUPTION: THE COMPOUND DISRUPTION RESILIENCE (CDR) INDEX AND A COST-EFFICIENCY RANKING OF MITIGATIONS

Abstract. A micro-enterprise that manufactures personalized products in one country and sells them in another faces disruptions differing in kind: electricity outages, cross-border logistics failures, supplier delays, marketplace channel closure, and tariff shocks. The qualitative resilience literature identifies which capabilities matter but not which mitigation to buy first under a constrained budget; the quantitative literature offers formal allocation models but presupposes a multi-echelon network, an estimable disruption history, and analytical capacity a micro-enterprise does not possess. This study proposes the Compound Disruption Resilience (CDR) model to occupy that gap. It adopts annual contribution margin as a numeraire, rendering heterogeneous shocks commensurable and summable into a single expected-loss functional E[L]; and it separates diagnosis, via a composite Resilience Index RI that reduces to an auditable identity — the exposure-weighted mean of mitigation effectiveness — from prescription, via a Mitigation Efficiency Ratio MER ranking mitigations by expected annual loss averted per unit of annual mitigation cost, with adoption gated at MER ≥ 1. The model is calibrated on a UV-printing operation producing in Ukraine for the United States market. Four of the five disruption probabilities are the author's practitioner priors, not frequencies estimated from an event log; the sensitivity analysis quantifies their influence. The full five-mitigation portfolio attains RI = 0.716, averting USD 41,588 of expected annual loss for USD 17,900 of annual spend; applying the model's own adoption gate, the operator instead stops at USD 13,700 and RI = 0.648, averting USD 37,661, because the fifth mitigation fails the gate. The central result is that the MER ordering diverges from both the ranking by unmitigated impact and that by absolute loss averted: the mode with the largest impact ranks last and fails the gate. A joint Monte Carlo over all five probabilities places RI in a 90% interval of [0.694, 0.742] (full support 0.672–0.762) and leaves the top-ranked mitigation unchanged in 77% of draws. The contribution is not a claim to have measured disruption probabilities, but a demonstration that for this decision, on this calibration, such measurement is not required.

Keywords: supply chain resilience; disruption risk; micro-enterprise; expected loss; mitigation portfolio; benefit-cost ratio; cross-border e-commerce; tariff shock; print-on-demand; sensitivity analysis

## Introduction

Resilience is normally studied at the scale at which it can be measured: a corporation with a multi-echelon supplier network, a disruption event history, and an analytics function capable of estimating from it. A substantial class of producers falls outside that scale. A micro-enterprise that prints personalized accessories on imported blanks in one country and sells them through marketplaces in another is a manufacturing operation in every economically meaningful sense — it owns equipment, holds inventory, incurs conversion cost, and can be stopped — yet it has one site, a handful of suppliers, no disruption log, and no analyst.

Such an operation is nonetheless exposed to disruption of an unusually compound kind. The operation motivating this study produces in Ukraine and sells into the United States. Its electricity supply is subject to wartime interruption; its finished goods travel a single consolidated cross-border lane; its blanks arrive from China on a second lane; its revenue reaches customers through platforms it does not own; and its landed cost is set in part by a regulatory regime that changed materially on 29 August 2025, when Executive Order 14324 suspended duty-free de minimis treatment for all countries [17]. These exposures are not variants of one another. An outage destroys output; a tariff compresses margin; a channel closure severs demand. They differ in kind, share no common physical unit, and can occur in the same year.

The operator's question is not whether resilience is desirable, but which of several candidate mitigations — a generator, an inventory buffer abroad, a second supplier, a second sales channel, a customs-consolidation programme — to fund first, given a small protection budget and no precise knowledge of any underlying probability. This is an allocation question, and a quantitative one, answerable neither by a capability checklist nor by a model requiring data the firm does not have.

This study proposes the Compound Disruption Resilience (CDR) model, with four objectives: to specify an expected-loss functional rendering heterogeneous modes commensurable through a single numeraire; to derive an auditable composite Resilience Index; to derive a ranking statistic — the Mitigation Efficiency Ratio — converting the index into a funding sequence with an explicit adoption gate; and to establish whether the resulting decision survives the admitted subjectivity of its inputs.

## Literature review

### Qualitative frameworks

Christopher and Peck established the canonical definition of supply chain resilience — the ability to return to an original or more desirable state after disruption — and organized it around resilience by design, collaboration, agility, and a risk-management culture [2]. Sheffi and Rice framed the operational choice as one between redundancy and flexibility, redundancy buying protection at a carrying cost [3]; the cost-effectiveness ratio developed below formalizes that trade-off. Pettit, Fiksel and Croxton advanced a framework in which resilience emerges from the balance of vulnerabilities and capabilities [5], later operationalized as the SCRAM assessment tool [9]. Craighead and colleagues separated disruption severity from the mitigation capabilities that attenuate it [4] — the precedent for separating impact from effectiveness here. This work is indispensable for identifying what matters, but yields no computable rule for allocating a constrained budget: SCRAM produces a capability-vulnerability balance, not a ranked funding sequence with a cost gate.

### Quantitative models of resilience investment

The quantitative strand is rigorous and correspondingly demanding of data. Ivanov, Sokolov and Dolgui introduced the ripple effect — propagation of a local disruption through a network — and framed disruption management as a trade-off among efficiency, flexibility and resilience [10]. Sawik formulated resilient supply-portfolio selection as a mixed-integer program, in which interaction among suppliers makes the optimum irreducible to a simple ordering [8]. Sahebjamnia, Torabi and Mansouri integrated business continuity and disaster recovery planning into one optimization model [11]. Ivanov and Dolgui extended the field from resilience to viability and survivability, motivated by the long-duration, simultaneous disruptions of COVID-19 [15] — a framing that justifies modelling several heterogeneous modes at once. Hosseini, Ivanov and Dolgui's review maps this toolkit and thereby delineates the gap addressed here: the methods are built for multi-echelon networks and for firms with disruption data and optimization capacity [14].

### Small firms and empirical analogues

Herbane argued that small-business crisis exposure is structurally distinct and under-researched [6], and Sullivan-Taylor and Branicki showed that prescriptions designed for large organizations do not transfer to SMEs, which differ in resource base, formalization and risk perception [7]. The strand warrants separate treatment but remains descriptive.

Three empirical literatures provide qualitative analogues for modes whose parameters nonetheless remain assumed. Allcott, Collard-Wexler and O'Connell estimated plant-level effects of electricity shortages in Indian manufacturing [12] — the closest independently estimated analogue to the outage mode, though two features of their results bear on how it may be used. Their shortage effect is of the order of 5–10% of producer surplus for the average plant, whereas $d_1 = 0.25$ is assumed below on the basis of the wartime context; their estimate does not support that magnitude, so [12] is an analogue rather than corroboration. They further report significant economies of scale in generator costs, implying that self-supply economics for a single-site micro-enterprise are likely worse, not better, than for the plants they study — a consideration cutting against the base case rather than for it. Zhu and Liu documented that a platform owner may enter the product spaces of its complementors, depressing their subsequent growth and innovation incentives [13]; this evidences platform-dependence risk — that a platform owner can act against the interests of sellers dependent on it — without speaking to termination of access. Fajgelbaum and colleagues estimated near-complete pass-through of United States import tariffs to domestic prices [16], licensing treatment of a duty as a direct margin loss to the importer rather than one shared with suppliers.

### Research gap

Qualitative frameworks [2, 3, 5, 9] identify which capabilities matter but not which to buy first; quantitative models [8, 10, 11, 14] presuppose networks and data; the SME strand [6, 7] establishes difference without providing a rule. There is accordingly no reproducible, low-data decision model that lets a micro-enterprise facing simultaneous heterogeneous disruptions rank candidate mitigations by loss reduction per unit cost and defend that ranking when the underlying probabilities are admittedly subjective. The CDR model is offered to fill that gap.

## Methodology

### Model boundary and the numeraire

The model takes the production decision as given and asks how to keep producing through shocks. It is single-echelon and firm-level — the operation, not a network, with no propagation of loss between tiers — and static and annual.

The central modelling choice is the numeraire. Modes $k_1$–$k_4$ destroy output; mode $k_5$ compresses margin. These are not the same kind of harm and share no physical unit, but they share a financial one: both reduce annual contribution. Expressing every mode as a fraction of annual baseline contribution destroyed is what makes an outage and a tariff commensurable, and hence what makes a single expected loss possible. The cost is that harms which are not cash flows are invisible to the model.

Unit economics are the author's own operating data for the representative product basket — phone case 40%, T-shirt 25%, mug 20%, MacBook case 15%. The blended per-order quantities are

$$\bar{p} = 0.40 \times 29.99 + 0.25 \times 34.99 + 0.20 \times 24.99 + 0.15 \times 54.99 = 33.99 \text{ USD/order}$$

$$\bar{c} = 0.40 \times 6.50 + 0.25 \times 9.20 + 0.20 \times 6.30 + 0.15 \times 13.00 = 8.11 \text{ USD/order}$$

The blended marketplace fee applies a 70% Etsy / 30% Amazon Handmade channel split to the respective platform takes:

$$\text{fee\_rate} = 0.70 \times 0.108 + 0.30 \times 0.15 = 0.1206 \Rightarrow \bar{f} = 0.1206 \times 33.99 = 4.10 \text{ USD/order}$$

With blended cross-border logistics $\bar{l}$ = 5.00 USD/order, contribution per order is

$$\bar{\pi} = \bar{p} - \bar{c} - \bar{f} - \bar{l} = 33.99 - 8.11 - 4.10 - 5.00 = 16.78 \text{ USD/order} \quad \text{(F1)}$$

$$\Pi_0 = \bar{\pi} \times Q_a = 16.78 \times 6{,}000 = 100{,}685 \text{ USD/yr}$$

where $Q_a$ = 6,000 orders/yr (500 per month) is the base operating point, matching the 500-order scenario of the author's companion study, and sitting well inside the author's observed single-printer throughput of at least 1,000 orders/month. All intermediate quantities are computed at full precision and rounded once for display; $\bar{\pi}$ = 16.7808 and $\Pi_0$ = 100,684.84 before display rounding. This convention applies to every figure reported below.

Three divergences from other studies in the series, which use the same basket, should be stated so the difference is not read as inconsistency. First, fee basis: a companion study on capacity sequencing applies an Etsy-only schedule (9.5% × $\bar{p}$ + 0.45 = 3.68, effectively 10.8%), this paper the 70/30 Etsy/Amazon blend above (12.06% × $\bar{p}$ = 4.10). Second, declared value for duty: that study uses in-house cost C = 8.11, this paper $\bar{c}$ plus the international leg, 8.71. Third, duty placement: that study books a 10% duty inside the margin (0.81/order), reaching a contribution margin of 16.39/order, whereas this paper excludes the duty from $\bar{\pi}$ and models it as shock $k_5$. $\Pi_0$ = USD 100,685 is therefore the pre-29-August-2025 counterfactual baseline against which the tariff shock is measured, not current contribution — which is why the figures differ, and why one figure serving both would make the tariff shock either invisible or double-counted.

### The expected-loss functional and the composite index

Each mode k is characterized by the risk triplet in the sense of Kaplan and Garrick [1] — what can go wrong, how likely, with what consequence — reduced to an annual probability $p_k$ and an unmitigated impact $d_k$ expressed as the fraction of $\Pi_0$ destroyed. Following the severity/mitigation separation of Craighead and colleagues [4], a mitigation removes a fraction $m_k$ of that impact at annualized cost $c_k$. Expected normalized annual contribution loss, and the unmitigated counterfactual against which resilience is measured, are

$$E[L] = \sum_k p_k \times d_k \times (1 - m_k) \quad \text{(F2)}$$

$$E[L_0] = \sum_k p_k \times d_k \quad \text{(F3)}$$

The composite Compound Disruption Resilience Index is the share of unmitigated exposure removed:

$$RI = 1 - \frac{E[L]}{E[L_0]} = \frac{\sum_k p_k \times d_k \times m_k}{\sum_k p_k \times d_k} \quad \text{(F4)}$$

The right-hand identity matters: RI is the exposure-weighted mean of the $m_k$, and is therefore auditable from the parameter table.

Because (F2) treats modes as independent and losses as additive, it is a first-order upper bound. A co-occurrence-corrected form, bounded at unity under independence, is reported alongside it:

$$E[L]^{mult} = 1 - \prod_k (1 - p_k \times d_k \times (1 - m_k)) \quad \text{(F5)}$$

$$E[L_0]^{mult} = 1 - \prod_k (1 - p_k \times d_k), \qquad RI^{mult} = 1 - \frac{E[L]^{mult}}{E[L_0]^{mult}}$$

**The ranking statistic and the decision rules**

The index diagnoses; it does not prescribe. Prescription requires cost. The Mitigation Efficiency Ratio is the USD of expected annual loss averted per USD of annual mitigation spend:

$$MER_k = \frac{p_k \times d_k \times m_k \times \Pi_0}{c_k} \quad \text{(F6)}$$

MER is a dimensionless annualized benefit-cost ratio. Its gate has three equivalent forms:

$$MER_k \geq 1 \Leftrightarrow p_k \geq p_k^* = \frac{c_k}{d_k \times m_k \times \Pi_0} \Leftrightarrow c_k \leq c_k^{\max} = p_k \times d_k \times m_k \times \Pi_0 \quad \text{(F7)}$$

Two nested rules follow. Rule 1 (diagnosis) maps RI to bands: RI < 0.40 fragile; 0.40 ≤ RI < 0.60 partially resilient; 0.60 ≤ RI < 0.80 resilient; RI ≥ 0.80 hardened. Rule 2 (selection) is operative: (a) adopt k only if $MER_k \geq 1$; (b) rank survivors by descending $MER_k$ — this ordering, not that by $d_k$ or $p_k \cdot d_k$, is the funding sequence; (c) under an annual budget B, fund greedily down the MER order:

$$\text{adopt in descending } MER \text{ while } \sum_k c_k \leq B \text{ and } MER_k \geq 1 \quad \text{(F8)}$$

and (d) stop at the first $MER_k < 1$ even if budget remains. The MER ordering is the exact funding priority, and greedy funding attains the optimum wherever the budget does not bind mid-sequence — including every budget reported in Table 5, each of which equals a cumulative greedy cost, and any B ≥ 13,700, at which all gate-passing mitigations are funded. Because the mitigations are indivisible, the general budget-constrained problem is a 0-1 knapsack and (F8) is then a heuristic: at B = 11,900, for instance, (F8) funds the outage, supplier and tariff mitigations for USD 29,204 averted, whereas outage, supplier and logistics averts USD 33,755. What separates this setting from Sawik's supplier portfolios [8] is therefore separability rather than exactness — the mitigations do not interact, so the objective decomposes by mode and no mixed-integer program is needed to evaluate a portfolio, only to optimize one under a budget that binds mid-sequence. Two adjacent mitigations a and b exchange rank at

$$p_a^{rev} = \frac{MER_b \times c_a}{d_a \times m_a \times \Pi_0} \quad \text{(F9)}$$

**Bridge to the transition decision**

The author's companion study defined the Print-to-In-house Transition Index and reported break-even Q* = 351 orders/month at capital expenditure I = USD 25,000 and maximum payback $P_{max}$ = 12 months. That model is deterministic and shock-free. The present model deflates its reported output without re-deriving any of its mathematics.

Only output-destroying modes belong in the deflator. PTI's numerator is the monthly saving $M = \bar{S} \times Q$, which scales with the volume actually produced: modes $k_1$–$k_4$ destroy output and hence reduce Q and M proportionally, whereas mode $k_5$ compresses margin per unit and leaves Q — and therefore M — unchanged. A duty does not reduce the per-unit saving from internalizing production, so its expected loss does not deflate PTI. The deflator is accordingly restricted:

$$E[L]^{out} = \sum_{k=1..4} p_k \times d_k \times (1-m_k), \qquad E[L_0]^{out} = \sum_{k=1..4} p_k \times d_k$$

$$PTI_R = PTI \times (1 - E[L]^{out}), \qquad Q_R^* = Q^*/(1 - E[L]^{out}) \quad \text{(F10)}$$

The bridge runs strictly one way: PTI's per-unit savings, monthly saving and break-even expression are imported as reported constants.

## Parameters and their status

Table 1 reports the five modes, with every assumed cell flagged. The probabilities $p_1$–$p_4$ are the author's practitioner priors elicited from operating experience, not frequencies estimated from a disruption event log: no such log exists for a single micro-enterprise over a period long enough to estimate a 10%-probability event. They are model parameters, and the sensitivity analyses below quantify their influence.

Table 1
*Disruption modes and base-case parameters*

| k | Disruption mode | $p_k$ | Status of $p_k$ | $d_k$ | Status of $d_k$ | $m_k$ | $c_k$ (USD/yr) | Mitigation |
|---|---|---|---|---|---|---|---|---|
| 1 | Electricity outage at the production site | 0.90 | Assumed | 0.250 | Assumed | 0.85 | 3,500 | Inverter/battery bank + generator; amortized capital + fuel |
| 2 | Cross-border logistics disruption (UA→US lane) | 0.40 | Assumed | 0.300 | Assumed | 0.70 | 6,000 | US micro-warehouse buffer + second lane/carrier; rent + working-capital carry |
| 3 | Supplier (blank) delivery delay (CN→UA lane) | 0.50 | Assumed | 0.150 | Assumed | 0.80 | 1,800 | Blank safety stock + second supplier; carrying cost |
| 4 | Marketplace account/channel closure | 0.10 | Assumed | 0.600 | Assumed | 0.65 | 4,200 | Dual channel (Etsy + Amazon Handmade) + owned storefront; listing/ops overhead |
| 5 | Tariff/regulatory shock (US de minimis suspension) | 1.00 | Observed [17] | 0.097 | Derived (one observed and two assumed inputs) | 0.40 | 2,400 | Master-box consolidation + HTS classification + US-side final packaging; broker retainer + compliance |

All $m_k$ and all $c_k$ are assumed. Mode $k_5$'s probability is observed: Executive Order 14324, signed 30 July 2025 and effective for articles entered for consumption on or after 29 August

2025, suspended duty-free de minimis treatment for all countries [17]. The tariff shock is realised, not prospective, so $p_5 = 1.00$. The model degenerates to a deterministic cost when $p_k = 1$, which lets a realised shock and a prospective risk be ranked on one scale.

Mode $k_5$'s impact is derived by arithmetic rather than assumed directly, but two of its three inputs are themselves assumptions. Declared value per unit is in-house cost plus the international leg, $\bar{c} + 0.60 = 8.71$ USD, where the USD 0.60 international leg is the author's observed master-box economics (USD 72 per box holding 120 units). The duty rate is taken as 10% of declared value; the applicable rate is tariff-classification-dependent, and 10% is an assumption, not an observed rate for this operation's goods. Customs brokerage at USD 90 per formal entry — also an assumption — spread over the 120-unit master box, is 0.750 USD/unit. Duty at 10% is 0.871 USD/unit. Hence

$$d_5 = \frac{0.871 + 0.750}{16.78} = \frac{1.621}{16.78} = 0.0966 \approx 0.097$$

Treating the duty as a direct margin loss rather than one shared with suppliers is licensed by the near-complete pass-through estimated by Fajgelbaum and colleagues [16].

### Computation and sensitivity design

All quantities were computed in Python from the inputs above; every figure in the Results is reproducible from Table 1 and (F1)–(F10), up to Monte Carlo error where simulation is used. Because $p_k$ is the least observable parameter class, three analyses address it: a one-way perturbation of each $p_k$ by ±50% (capped at 1.0); a joint Monte Carlo drawing all five probabilities simultaneously from $p_k \sim \text{Uniform}[0.5 \cdot p_k, \min(1, 1.5 \cdot p_k)]$ over 20,000 draws, implemented with numpy.random.default_rng(7), five independent uniforms drawn per iteration in mode order $k_1..k_5$; and the exact rank-reversal thresholds from (F9). A one-way analysis of $m_k$ at ±0.20 absolute, capped at 1.0, is reported for comparison, and (F5) is evaluated as a functional-form robustness check.

## Results

### Expected loss and the composite index

Table 2
*Expected annual contribution loss by disruption mode ($\Pi_0$ = USD 100,685/yr)*

| k | Mode | $p_k \times d_k$ | Residual $p_k \times d_k \times (1 - m_k)$ | Unmitigated loss (USD/yr) | Residual loss (USD/yr) | Loss averted (USD/yr) |
|---|---|---|---|---|---|---|
| 1 | Electricity outage | 0.2250 | 0.0338 | 22,654 | 3,398 | 19,256 |
| 2 | Cross-border logistics | 0.1200 | 0.0360 | 12,082 | 3,625 | 8,458 |
| 3 | Supplier delay | 0.0750 | 0.0150 | 7,551 | 1,510 | 6,041 |
| 4 | Channel closure | 0.0600 | 0.0210 | 6,041 | 2,114 | 3,927 |
| 5 | Tariff shock | 0.0970 | 0.0582 | 9,766 | 5,860 | 3,907 |
| | Total | 0.5770 | 0.1640 | 58,095 | 16,507 | 41,588 |

Unmitigated expected annual loss is $E[L_0] = 0.5770$ of baseline contribution, or USD 58,095/yr; with all five mitigations in place it falls to $E[L] = 0.1640$, or USD 16,507/yr. Total annual mitigation cost is USD 17,900. Net benefit is

$$\text{Net} = (0.5770 - 0.1640) \times 100{,}685 - 17{,}900 = 41{,}588 - 17{,}900 = 23{,}688 \text{ USD/yr}$$

and the portfolio benefit-cost ratio is 41,588 / 17,900 = 2.32. The composite index is

$$RI = 1 - \frac{0.1640}{0.5770} = 0.716$$

placing the operation in the resilient band. By the identity in (F4), 0.716 is the exposure-weighted mean of the five $m_k$, weighted by $p_k \cdot d_k$. This is the full five-mitigation set; the gate applied below removes one of them.

### The MER ranking

Table 3
*Mitigation Efficiency Ratio ranking (base case)*

| Rank | Mode | Loss averted (USD/yr) | Annual cost (USD/yr) | MER | Gate (MER ≥ 1) |
|---|---|---|---|---|---|
| 1 | Electricity outage | 19,256 | 3,500 | 5.50 | Pass |
| 2 | Supplier delay | 6,041 | 1,800 | 3.36 | Pass |
| 3 | Tariff shock | 3,907 | 2,400 | 1.63 | Pass |
| 4 | Cross-border logistics | 8,458 | 6,000 | 1.41 | Pass |
| 5 | Channel closure | 3,927 | 4,200 | 0.93 | Fail |

Three orderings of the same five mitigations are available, and all three differ. By raw unmitigated impact $d_k$: channel closure (0.60) > logistics (0.30) > outage (0.25) > supplier (0.15) > tariff (0.097). By absolute loss averted: outage (19,256) > logistics (8,458) > supplier (6,041) > channel (3,927) > tariff (3,907). By MER: outage > supplier > tariff > logistics > channel. The mode with the largest impact ranks last; the mode averting the second-largest absolute loss ranks fourth.

### The adoption gate

Table 4
*Adoption gate: break-even probability and cost headroom*

| Mode | $p_k^*$ | $p_k$ (base) | Verdict | $c_k^{max}$ (USD/yr) | $c_k$ (USD/yr) | Headroom (USD/yr) |
|---|---|---|---|---|---|---|
| Electricity outage | 0.164 | 0.90 | Pass | 19,256 | 3,500 | +15,756 |
| Cross-border logistics | 0.284 | 0.40 | Pass | 8,458 | 6,000 | +2,458 |
| Supplier delay | 0.149 | 0.50 | Pass | 6,041 | 1,800 | +4,241 |
| Channel closure | 0.107 | 0.10 | Fail | 3,927 | 4,200 | −273 |
| Tariff shock | 0.614 | 1.00 | Pass | 3,907 | 2,400 | +1,507 |

The dual-channel mitigation fails the gate at base parameters: MER = 0.93, headroom −273 USD/yr, and a break-even probability of 0.107 sitting within 7% of the assumed $p_4$ = 0.10. The operation does in fact maintain dual channels. The model's verdict is therefore not that the dual channel should be abandoned, but that its presence is not justified by disruption avoidance alone at these parameters; its justification is its demand-side revenue contribution, which lies outside this model's boundary. Applying Rule 2 to Table 3, the gate-disciplined portfolio comprises the four passing mitigations: annual cost USD 13,700, expected annual loss averted USD 37,661, and RI = 0.648.

### The budget frontier

Table 5

*Budget-constrained resilience frontier (greedy over the MER order)*

| Budget B (USD/yr) | Last mitigation funded | RI | Residual E[L] (USD/yr) |
|---|---|---|---|
| 3,500 | Electricity outage | 0.331 | 38,839 |
| 5,300 | Supplier delay | 0.435 | 32,798 |
| 7,700 | Tariff shock | 0.503 | 28,892 |
| 13,700 | Cross-border logistics | 0.648 | 20,434 |
| 17,900 | Channel closure | 0.716 | 16,507 |

The frontier is markedly concave: 30% of full spend (USD 5,300 of 17,900) buys 61% of the index attainable at full spend (0.435 of 0.716). The final increment is the one failing the gate, so a budget-disciplined operator applying Rule 2(d) stops at B = 13,700 and RI = 0.648.

### Sensitivity to the assumed parameters

Table 6

*One-way sensitivity of RI to disruption probability (±50%, capped at 1.0) and to mitigation effectiveness (±0.20 absolute, capped at 1.0)*

| Mode | RI at $0.5 \cdot p_k$ | RI at $1.5 \cdot p_k$ | Range ($p_k$) | RI at $m_k - 0.20$ | RI at $m_k + 0.20$ | Range ($m_k$) |
|---|---|---|---|---|---|---|
| Electricity outage | 0.683 | 0.721 | 0.038 | 0.638 | 0.774 | 0.136 |
| Cross-border logistics | 0.718 | 0.714 | 0.003 | 0.674 | 0.757 | 0.083 |
| Supplier delay | 0.710 | 0.721 | 0.011 | 0.690 | 0.742 | 0.052 |
| Channel closure | 0.719 | 0.713 | 0.007 | 0.695 | 0.737 | 0.042 |
| Tariff shock | 0.745 | 0.716 | 0.029 | 0.682 | 0.749 | 0.067 |

The cap at 1.0 binds only for the outage row, where $m_1 + 0.20 = 1.05$ would otherwise imply averting more than the impact; uncapped, that cell would read 0.794 and the range 0.156.

No single ±50% perturbation of any $p_k$ moves RI by more than 0.038 from the base value of 0.716, whereas a ±0.20 shift in a single $m_k$ moves it by up to 0.136. The index is materially more sensitive to mitigation effectiveness than to disruption probability, and $m_k$ is at least partly under the operator's control whereas $p_k$ is not.

Table 7

*Joint Monte Carlo over all five probabilities ($p_k \sim \mathrm{Uniform}[0.5 \cdot p_k, \min(1, 1.5 \cdot p_k)]$; 20,000 draws; numpy.random.default_rng(7), five independent uniforms per iteration in mode order $k_1..k_5$)*

| Statistic | RI |
|---|---|
| Mean | 0.719 |
| Median | 0.719 |
| 5th percentile | 0.694 |
| 95th percentile | 0.742 |
| Minimum | 0.672 |
| Maximum | 0.762 |

Rank stability is the decision-relevant question. Across the 20,000 draws the top-ranked mitigation is the outage mitigation in 77.2% of draws and the supplier mitigation in 22.8%; no other ever attains rank 1. Rank 2 is held by the supplier mitigation in 75.1% of draws, the outage mitigation in 22.8%, and the logistics mitigation in 2.1%; the channel-closure mitigation never enters the top two. It fails the MER gate in 57.0% of draws, indicating that its base-case failure is marginal rather than robust.

**Rank-reversal thresholds**

Applying (F9) to adjacent pairs gives the conditions under which the recommendation would change. The outage mitigation falls below the supplier mitigation only if $p_1 < 0.549$ — a 39% drop from the assumed 0.90. The supplier mitigation falls below the tariff mitigation only if $p_3 < 0.242$, against an assumed 0.50. The logistics mitigation falls below channel closure only if $p_2 < 0.265$, against an assumed 0.40. The tariff mitigation would fall below logistics if $p_5 < 0.866$, but $p_5$ is observed at 1.00, so this branch is closed.

**Functional-form robustness**

Evaluating (F5) gives $E[L_0]^{mult} = 0.4645$ (USD 46,770/yr), $E[L]^{mult} = 0.1541$ (USD 15,511/yr) and $RI^{mult} = 0.668$. The additive form overstates $E[L_0]$ by 24.2% and $E[L]$ by 6.4%; the index differs by 4.7 percentage points and remains inside the same resilient band. The additive form is a first-order upper bound valid while $\Sigma\, p_k \cdot d_k$ remains well below 1; at $E[L_0] = 0.577$ it is already visibly conservative, which is why both forms are reported.

**The bridge to the transition decision**

Table 8

*Resilience-adjusted transition index (PTI values are the reported constants of the author's companion study; I = USD 25,000; $P_{max}$ = 12 months; deflator restricted to output-destroying modes $k_1$–$k_4$)*

| Orders/month | PTI | $PTI_R$ (mitigated, $E[L]^{out} = 0.106$) | $PTI_R$ (unmitigated, $E[L_0]^{out} = 0.480$) |
|---|---|---|---|
| 50 | 0.14 | 0.13 | 0.07 |
| 100 | 0.29 | 0.26 | 0.15 |
| 300 | 0.86 | 0.77 | 0.45 |
| 500 | 1.43 | 1.28 | 0.74 |
| 1,000 | 2.85 | 2.55 | 1.48 |

(F10) is applied to the index values as reported in the companion study; applying it to full-precision PTI values instead shifts the 500-order figure to 1.27 and leaves every band assignment unchanged.

Applying (F10) to the reported break-even volume gives $Q^*_R = 351 / (1 - 0.106) = 393$ orders/month under mitigation, an increase of 11.8%; without mitigation, $Q^*_R = 351 / (1 - 0.480) = 675$ orders/month, an increase of 92.3%. At the base volume of 500 orders/month, mitigated exposure moves the index from 1.43 to 1.28 — still inside PTI's transition to in-house band, but with the headroom above the 1.25 boundary reduced from 0.18 to 0.03. Unmitigated, the index falls to 0.74, which lies in the retain POD band, but by a margin of only 0.006.

## Discussion

**The ranking is the contribution**

The headline finding is not the value of RI but the divergence of the three orderings in Table 3. A practitioner asked to protect a business against five threats will, absent a model, protect the biggest threat first. Here the biggest threat by unmitigated impact is channel closure — plausibly so, since losing the primary marketplace account would remove most of the year's contribution, and the platform-dependence risk evidenced by Zhu and Liu is real [13]. Yet the dual-channel mitigation ranks last and fails the gate, because the mode is improbable and the mitigation expensive relative to what it averts. Conversely the outage mitigation protects against a mode of moderate impact but near-certain occurrence, and averts USD 5.50 per dollar spent; its annual cost $c_1$ = USD 3,500 is assumed, and Allcott and colleagues' economies-of-scale finding suggests it may be optimistic for a single site [12].

Ranking by absolute loss averted also fails: it places the logistics mitigation second when it ranks fourth, because averting USD 8,458 for USD 6,000 is a poorer use of a scarce budget than averting USD 6,041 for USD 1,800. Only when the denominator enters does the funding sequence become correct — which is why the model separates index from ranking: an index alone, however auditable, cannot allocate a budget.

### On the status of the parameters

Four of five probabilities are the author's priors, and the question is whether a decision resting on them can be defended. The sensitivity results answer it in a specific and limited way. Perturbing all five probabilities simultaneously by up to ±50% leaves RI within a range narrower than a single decision band. More importantly, the funding decision is stable where it matters: the mitigation to fund first is one of exactly two candidates across 20,000 draws, and the mitigation to reject never rises into the top two. The rank-reversal thresholds locate the boundary exactly — the ordering would change only if $p_1$ fell below 0.549 or $p_2$ below 0.265, outside the plausible range for a wartime site. The claim is therefore not that the probabilities are correct, but that for this decision, on this calibration, they do not need to be.

Two caveats bound this. The model is markedly more sensitive to $m_k$ than to $p_k$, and $m_k$ is assumed and not independently validated: the parameter that matters most is the one least examined here. And the channel-closure gate failure is marginal in both senses — 7% from break-even at base parameters, failing in 57.0% of draws. The model discriminates, but does not adjudicate that case.

### Relation to the resilience literature

Against SCRAM [9], the CDR model trades breadth for decidability: SCRAM is a qualitative multi-item instrument producing a capability-vulnerability balance, whereas RI is a single computed ratio and MER an explicit cost gate. SCRAM assesses far more than five modes, and CDR is no substitute for it as an assessment. Against Sawik [8] the difference is structural: an MIP is required because supplier portfolios interact, whereas CDR's mitigations are separable, so the objective decomposes by mode — a simplification that follows from the setting and would not survive transfer to a network. Against the ripple-effect literature [10, 15], CDR is single-echelon by construction and models no propagation: a boundary rather than an oversight, and the first thing that would have to change at larger scale. Relative to the SME strand [6, 7], the model supplies the computable rule that strand identifies the need for; relative to the business-continuity models of Sahebjamnia and colleagues [11], the difference is again analytical capacity, an integrated BCM/DRP optimization presupposing a planning function a micro-enterprise does not have.

### Relation to the author's companion and sibling work

The bridge in Table 8 produces a result neither model yields alone. The companion study asked whether to produce at all and reported, at 500 orders/month, that the transition is justified; that model is deterministic and shock-free. Deflating its reported output by expected output loss leaves that verdict standing under mitigation, at 1.28 against a 1.25 boundary, but consumes most of the headroom that supported it. Unmitigated, the index falls to 0.74, marginally on the other side of the 0.75 boundary; given a margin of 0.006 and elicited probabilities, this is better read as the in-house case being brought to the edge of reversal than as a demonstrated reversal. On this calibration, the mitigation portfolio is a condition that preserves the in-house case rather than an overhead upon it. The bridge runs one way only, and none of PTI's mathematics is re-derived here.

Boundaries against the sibling models should also be stated. The model does not represent labour, task time or the design workflow, and so does not overlap the author's work on AI leverage, whose numeraire is operator-hours. It is product-agnostic by construction: it collapses the portfolio into a single blended contribution and never resolves loss to the product level, which is the product-complexity model's unit of analysis. The operation's observed 7–8% scrap rate is not a disruption mode — it is a steady-state process loss already booked into $\bar{c}$, and does not enter E[L]. The equipment is a single node that is either up or degraded, with nothing represented about balance, queueing or utilization, which is the equipment-utilization model's territory. One adjacency must be disclaimed rather than exploited: a second UV printer would incidentally provide outage redundancy, but redundancy-as-capacity is a capacity-expansion decision belonging to the sequencing model — the author's queueing-grounded Production Scalability Index [18], which times the acquisition of additional capacity rather than pricing its resilience side-effect — and is not priced here.

## Limitations

Several limitations qualify these findings. The probabilities $p_1$–$p_4$ are elicited priors, not estimated frequencies; the sensitivity analyses bound their consequences but do not substitute for data. The effectivenesses $m_k$ are assumed and unvalidated, and dominate the index's sensitivity. The costs $c_k$ are annualized estimates, bounded by the author's observed cost structure where it constrains them and by judgement elsewhere. Three inputs entering the headline figures are assumptions and are flagged as such: the 10% duty rate and the USD 90 per-entry brokerage, both entering $d_5$, and the 70/30 channel split entering $\bar{f}$. The blended logistics cost $\bar{l}$ = USD 5.00 is the author's observed phone-case figure used as a portfolio-wide proxy; larger blanks carry more, making $\bar{\pi}$, $\Pi_0$ and every absolute USD figure mildly optimistic.

Structurally, the model is static, annual and deterministic in its parameters: no seasonality, no demand volatility, no recovery dynamics — resilience is expected annual loss avoided, not a time-to-recover trajectory. Independence and additivity are first-order approximations; (F5) corrects for co-occurrence, but genuine correlation, and loss propagation [10], are not modelled. An outage and a logistics failure sharing a wartime cause are positively correlated in reality; a copula or simulation approach, in the spirit of the methods surveyed by Hosseini and colleagues [14], is the first target for future work. Mitigations are assumed indivisible, mode-specific and non-interacting, whereas in practice the US micro-warehouse buffers both the logistics mode and, partly, the outage mode; this overlap is not credited, making the reported RI conservative.

The bridge carries its own limitation: it deflates PTI's reported index for expected output loss only. A duty also bears on the in-house regime's per-unit saving $\bar{S}$ directly, since it is the in-house regime that ships across the border; that channel is not modelled, because re-deriving PTI's savings term lies outside this study's boundary. $PTI_R$ as reported is therefore an upper bound in that respect, and the band assignments at 500 orders/month should be read with that in mind.

The calibration is single-operation: one firm, one technology, one origin-destination pair. Parameter values are not claimed to generalize; the model structure is what is offered as transferable, subject to validation elsewhere. Finally, RI is invariant to $Q_a$, but every absolute USD figure and every MER scales linearly with $\Pi_0$ and hence with volume — so the adoption gate is volume-dependent, and a mitigation failing at 500 orders/month may pass at 1,000.

The reputation cost of the roughly 10% of dissatisfied buyers observed on personalized, non-returnable goods is not modelled, being no cash flow and absent from $\Pi_0$.

Conclusions

This study proposed the Compound Disruption Resilience model for a micro-enterprise facing simultaneous heterogeneous disruptions. Its contribution is threefold: contribution margin as a numeraire, which makes shocks differing in kind commensurable and summable into one expected-loss functional; a composite index that reduces to an auditable identity rather than an unfalsifiable score; and a ranking statistic with an explicit cost gate that turns the index into a funding sequence.

Applied to a UV-printing operation producing in Ukraine for the United States market, the full five-mitigation portfolio attains RI = 0.716, averting USD 41,588 of expected annual loss for USD 17,900 of spend; the gate-disciplined portfolio the model actually prescribes stops at USD 13,700 and RI = 0.648, averting USD 37,661. The principal finding is that the cost-efficiency ordering diverges from both the ranking by unmitigated impact and that by absolute loss averted: the mode with the largest impact ranks last and fails the gate. In this calibration, ranking by impact misallocates the budget; whether the divergence generalizes is subject to validation on further operations. Four of five probabilities are elicited priors, yet perturbing all five simultaneously by ±50% leaves both the index and the funding decision materially unchanged — so the claim is not to have measured disruption probabilities, but that for this decision, on this calibration, measurement is not required.

Bridging to the author's companion study shows that mitigated exposure leaves the in-house verdict at the base volume intact while consuming most of its headroom, and that unmitigated exposure brings it marginally to the edge of reversal — so on this calibration the mitigation portfolio is better read as a condition preserving the in-house case than as an overhead upon it. Future work should model correlation between modes explicitly, since independence is the most consequential simplification; validate the mitigation-effectiveness parameters empirically, since they dominate the index's sensitivity; and test the structure across further operations.

**Література**

1. Kaplan S., Garrick B. J. On the quantitative definition of risk. Risk Analysis. 1981. Vol. 1, № 1. P. 11–27. DOI: 10.1111/j.1539-6924.1981.tb01350.x
2. Christopher M., Peck H. Building the resilient supply chain. The International Journal of Logistics Management. 2004. Vol. 15, № 2. P. 1–14. DOI: 10.1108/09574090410700275
3. Sheffi Y., Rice J. B. Jr. A supply chain view of the resilient enterprise. MIT Sloan Management Review. 2005. Vol. 47, № 1. P. 41–48.
4. Craighead C. W., Blackhurst J., Rungtusanatham M. J., Handfield R. B. The severity of supply chain disruptions: design characteristics and mitigation capabilities. Decision Sciences. 2007. Vol. 38, № 1. P. 131–156. DOI: 10.1111/j.1540-5915.2007.00151.x
5. Pettit T. J., Fiksel J., Croxton K. L. Ensuring supply chain resilience: development of a conceptual framework. Journal of Business Logistics. 2010. Vol. 31, № 1. P. 1–21. DOI: 10.1002/j.2158-1592.2010.tb00125.x
6. Herbane B. Small business research: time for a crisis-based view. International Small Business Journal. 2010. Vol. 28, № 1. P. 43–64. DOI: 10.1177/0266242609350804
7. Sullivan-Taylor B., Branicki L. Creating resilient SMEs: why one size might not fit all. International Journal of Production Research. 2011. Vol. 49, № 18. P. 5565–5579. DOI: 10.1080/00207543.2011.563837
8. Sawik T. Selection of resilient supply portfolio under disruption risks. Omega. 2013. Vol. 41, № 2. P. 259–269. DOI: 10.1016/j.omega.2012.05.003

9. Pettit T. J., Croxton K. L., Fiksel J. Ensuring supply chain resilience: development and implementation of an assessment tool. Journal of Business Logistics. 2013. Vol. 34, № 1. P. 46–76. DOI: 10.1111/jbl.12009
10. Ivanov D., Sokolov B., Dolgui A. The ripple effect in supply chains: trade-off «efficiency-flexibility-resilience» in supply chain disruption management. International Journal of Production Research. 2014. Vol. 52, № 7. P. 2154–2172. DOI: 10.1080/00207543.2013.858836
11. Sahebjamnia N., Torabi S. A., Mansouri S. A. Integrated business continuity and disaster recovery planning: towards organizational resilience. European Journal of Operational Research. 2015. Vol. 242, № 1. P. 261–273. DOI: 10.1016/j.ejor.2014.09.055
12. Allcott H., Collard-Wexler A., O'Connell S. D. How do electricity shortages affect industry? Evidence from India. American Economic Review. 2016. Vol. 106, № 3. P. 587–624. DOI: 10.1257/aer.20140389
13. Zhu F., Liu Q. Competing with complementors: an empirical look at Amazon.com. Strategic Management Journal. 2018. Vol. 39, № 10. P. 2618–2642. DOI: 10.1002/smj.2932
14. Hosseini S., Ivanov D., Dolgui A. Review of quantitative methods for supply chain resilience analysis. Transportation Research Part E: Logistics and Transportation Review. 2019. Vol. 125. P. 285–307. DOI: 10.1016/j.tre.2019.03.001
15. Ivanov D., Dolgui A. Viability of intertwined supply networks: extending the supply chain resilience angles towards survivability. A position paper motivated by COVID-19 outbreak. International Journal of Production Research. 2020. Vol. 58, № 10. P. 2904–2915. DOI: 10.1080/00207543.2020.1750727
16. Fajgelbaum P. D., Goldberg P. K., Kennedy P. J., Khandelwal A. K. The return to protectionism. The Quarterly Journal of Economics. 2020. Vol. 135, № 1. P. 1–55. DOI: 10.1093/qje/qjz036
17. Executive Order 14324 of July 30, 2025. Suspending Duty-Free De Minimis Treatment for All Countries. Federal Register. 2025. Vol. 90. P. 37775–37780. URL: https://www.federalregister.gov/d/2025-14943 (дата звернення: 05.10.2025).
18. Sholomii Y. V. A queueing-grounded decision model for timing the acquisition of additional production capacity in small personalized-product manufacturers: the Production Scalability Index (PSI). Актуальні проблеми економіки. 2023. № 10, т. 2 (268/2). С. 136–148. DOI: 10.32752/1993-6788-2023-2-268-136-148. URL: https://eco-science.net/wp-content/uploads/2023/10/10-2.23._topic_Yevhen-Sholomii-136-148-1.pdf