International Journal of Web of Multidisciplinary Studies
(Peer-Reviewed, Open Access, Fully Refereed International Journal)
website: https://ijwos.com
Volume 3 Issue 9, September 2026
E-ISSN : 3049-2424
DOI: https://doi.org/10.71366/ijwos03092680879

The Portability Congestion Paradox: An Economic Model of AI Resilience and Scarce Recovery Capacity

Kwan Hong TAN[1]
[1] *Singapore University of Social Sciences, 463 Clementi Road, Singapore 599494.*

Article Info

Article History:
Published: 22 Sept 2026

Publication Issue:
Volume 3, Issue 9
September-2026

Page Number:
344-358

Corresponding Author:
Kwan Hong TAN

Abstract:

Technical portability is increasingly proposed as a response to dependence on concentrated artificial intelligence infrastructure. Yet a workload that can switch providers may still lack somewhere to run when many firms switch simultaneously. This paper develops a welfare model separating private portability preparation from shared recovery capacity. Firms rationally prepare transferable workloads but take the probability of obtaining backup service as given. Under proportional rationing, private investment follows the average recovery return, whereas efficient investment follows the marginal increase in total restored output. The resulting congestion externality produces excessive preparation whenever backup capacity binds. A portability subsidy can then raise resource costs without restoring additional output. The paper derives the exact welfare loss, distinguishes subsidies from genuine reductions in engineering costs, and shows how capacity reservations can implement efficient joint investment. Extensions cover a finite number of firms, uncertain capacity, ordinary switching benefits, and heterogeneous recovery values. Reproducible numerical illustrations demonstrate that expanding backup supply without correcting allocation incentives can also reduce net welfare. The contribution is a conditional economic theory of recovery competition, rather than a general argument against interoperability. Effective resilience policy must evaluate the compatibility of workloads, the capacity surviving the relevant disruption, and enforceable claims on that capacity together.

Keywords: artificial intelligence, capacity investment, congestion externalities, operational resilience, portability, welfare economics

1. Introduction

Artificial intelligence creates value when it can be integrated into dependable production. Evidence from customer support shows that access to generative assistance can improve worker productivity, although gains vary across workers [1]. As firms redesign workflows around such assistance, service interruption becomes an economic problem extending beyond the performance of an individual model. An organization may have alternative software interfaces and tested migration procedures yet remain unable to obtain the computing resources, qualified staff, or validated alternative services needed during a widespread disruption.

The Financial Stability Board identifies third-party dependencies and provider concentration among the vulnerabilities associated with artificial intelligence adoption [2]. Its subsequent monitoring report places particular attention on measuring those dependencies and the remaining information gaps [3]. These concerns motivate a distinction between two forms of preparedness. Technical portability makes a workload eligible for transfer. Recovery capacity makes that transfer operationally useful within a specified deadline. The two are complements, but they need not be supplied or priced together.

Consider a hypothetical group of businesses using the same primary inference service. Each business develops adapters, validates a substitute model, synchronizes relevant data, and rehearses its fallback procedure. These preparations are valuable when a small subset of businesses must switch. During a common disruption, however, all may seek the same spare accelerators or the same alternative service. If admission is rationed proportionally to eligible demand, making additional workloads portable can increase a firm's share of emergency service while reducing the share available to others. The firm receives a private recovery benefit even after additional preparation ceases to increase aggregate recovered output.

This paper asks three questions. When does individually rational portability exceed its efficient level? How do subsidies, genuine cost reductions, and backup-capacity expansion differ in their effects? What allocation arrangement can align private preparation with the social value of recovery? The analysis treats the outage probability and the underlying use of artificial intelligence as given. It therefore isolates a recovery-allocation mechanism without assuming that adoption is inherently excessive or that common failures necessarily arise from incorrect model outputs.

The central result is an average-versus-marginal return wedge. Firms value the average probability that a portable workload receives service. A planner values the additional output recovered when aggregate preparation increases. When backup supply is exhausted, the former remains positive while the latter becomes zero. A tractable benchmark yields closed-form equilibrium investment, efficient investment, welfare loss, and subsidy effects. A reservation market then supplies a constructive benchmark for jointly funding preparation and capacity. The extension to uncertain capacity replaces the sharp bottleneck with a general result expressed through the distribution of capacity surviving an outage.

The contribution is deliberately narrower than a new theory of common resources. It applies established economic reasoning to an understudied separation between eligibility to switch and the capacity required to make switching useful. The distinctive policy implication is that a measured increase in portability can accompany unchanged recovery and lower net welfare. Conversely, reservations can improve welfare without increasing the number of workloads technically eligible to move. These conclusions depend on the allocation rule, the relevant disruption state, and the absence or presence of other switching benefits. They do not establish the empirical prevalence of congestion or justify preserving commercial barriers to switching.

2. Literature Review

2.1 Complementary investment and the meaning of resilience

The productivity J-curve explains why general-purpose technologies require complementary intangible investment whose costs and benefits appear at different times [4]. This helps distinguish purchasing an artificial intelligence service from building an organization capable of using it effectively. Portability preparation is one such investment, but the present model adds a cross-firm constraint: the return to an intangible complement depends on shared emergency resources. More preparation need not produce more realized resilience when those resources have already been exhausted.

The United Kingdom Competition and Markets Authority's 2025 cloud investigation documents both commercial and technical obstacles to switching and multi-cloud use [5]. Its findings concern competition and customer choice; they do not demonstrate the recovery externality modeled here. Nevertheless, they establish why portability is an economically meaningful choice rather than a costless feature. The present paper concerns simultaneous emergency use after preparation has occurred. Routine migration across providers and common-shock failover should not be treated as equivalent events.

2.2 Common resources and interdependent protection

Gordon's analysis of common-property resources provides a foundational account of the gap between privately rewarded effort and efficient resource use [6]. Here, the rival resource is emergency service within a recovery deadline, and the costly activity is preparing eligible workloads. Unlike a renewable stock, recovery capacity is not depleted permanently by the modeled event. The analogy concerns dissipation of surplus through competing claims on a fixed service flow.

Kunreuther and Heal show that investment in protection can depend on the actions of other exposed agents, with implications for insurance and coordinating mechanisms [7]. Their interdependent-security problem concerns how protection choices alter exposure to loss. In the present mechanism, the outage occurs independently of preparation; interaction arises during allocation of recovery. This difference matters for policy. A subsidy that supports protection with positive spillovers need not be appropriate for an investment that expands claims on an already congested fallback resource.

Acemoglu, Ozdaglar, and Tahbaz-Salehi establish that financial interconnections can absorb smaller shocks yet transmit sufficiently large ones [8]. Their result cautions against assessing system resilience solely from isolated relationships. The present paper does not reproduce a financial contagion network. Instead, a common shock places otherwise separate firms in competition for a recovery pool. The relevant dependency can exist even without direct production links or financial claims among the firms.

2.3 Capacity markets and artificial intelligence governance

Joskow and Tirole analyze the conditions under which competitive electricity markets support efficient reliability and capacity investment [9]. Their work provides a close economic precedent for connecting scarcity allocation to investment incentives. This paper uses a simplified reservation market as an implementation benchmark. Its novelty is not that capacity can be priced, but that separate promotion of portability can expand emergency claims while leaving the underlying constraint untouched.

Engineering research on multi-cloud environments identifies vulnerabilities involving architecture, interfaces, authentication, automation, and management differences [10]. Accordingly, the model's capacity variable represents usable service after compatibility and dependency constraints, not installed hardware alone. The National Institute of Standards and Technology's risk-management framework similarly supports assessing risk across the system lifecycle [11]. Neither source supplies the numerical parameters used below; their role is to motivate careful measurement of operational rather than nominal substitutes.

Three related contributions by Tan help locate the present mechanism. The autonomy externality concerns the difference between private delegation incentives and losses transmitted through correlated failures [12]. Cognitive monoculture concerns common error in aggregated human and artificial intelligence judgments [13]. The governance half-life concerns deterioration in the effectiveness of controls over time [14]. The current paper studies a different decision margin: expenditure on accessing finite recovery resources after a disruption. It holds the triggering shock fixed and derives inefficiency even when all firms correctly understand that shock. Table 1 summarizes this distinction.

Table 1. Positioning of the recovery-capacity mechanism

Question	Related literature	Mechanism in this paper
Productivity complement	Brynjolfsson, Rock, and Syverson [4]	Portability preparation is a complementary intangible whose return depends on shared fallback capacity.
Switching and cloud competition	Competition and Markets Authority [5]	Technical and commercial switching barriers motivate portability investment;

		simultaneous recovery is the separate object.
Common resources and protection	Gordon [6]; Kunreuther and Heal [7]	Rival emergency service creates an allocation externality among otherwise separate firms.
Networked systemic risk	Acemoglu, Ozdaglar, and Tahbaz-Salehi [8]	A common shock can make a recovery pool fragile even without direct inter-firm links.
AI governance	Tan [12]–[14]; NIST [11]	The analysis isolates an economic recovery margin that complements control and dependency assessment.

3. Method

3.1 Economic environment and timing

The study is analytical, with numerical illustrations generated from the stated model. There are no observations of firms, estimated treatment effects, or inferred market parameters. A unit mass of risk-neutral firms each operates one unit of a divisible workload. A common disruption occurs with probability $p \in (0,1)$. Restoring a unit of workload during the event preserves value $v > 0$. Define $b = pv$, the expected benefit of a unit of successful recovery. Values and costs refer to the same planning horizon. Losses without recovery and ordinary production income are constants and are omitted from welfare comparisons.

Before uncertainty resolves, firm i chooses the portable share $x_i \in [0,1]$ and pays real preparation cost $cx_i^2/2$, where $c > 0$. Preparation can represent interface adaptation, data replication, validation, and trained fallback procedures. The convex cost captures the increasing difficulty of preparing less transferable tasks. The maintained benchmark restriction is $0 < b/c < 1$, so full portability is not privately optimal when capacity is abundant. Aggregate portable workload and the benefit-cost ratio are

$$X = \int_0^1 x_i \, \mathrm{d}i, \qquad a = \frac{b}{c}, \qquad 0 < a < 1. \tag{1}$$

Recovery capacity $K \in (0,1]$ is the workload share that can actually be served during the common disruption and within the specified deadline. This is residual capacity after surviving providers meet prior commitments. It is initially fixed and common knowledge. Once the outage occurs, technically eligible workloads compete for this pool. Under the benchmark allocation rule, each receives the same service fraction

$$q(X,K) = \min\left\{1, \frac{K}{X}\right\}, \qquad R(X,K) = Xq(X,K) = \min\{X,K\}. \tag{2}$$

Set $q(0,K) = 1$. The allocation is feasible because aggregate restored workload is at most capacity. Proportional rationing can be interpreted as equal treatment of eligible workload units; it is not an empirical claim about a particular provider's scheduling policy. Firms cannot inflate claims beyond their genuinely prepared workloads, and preparation alone does not create binding rights to a specific quantity of backup service.

3.2 Private and social objectives

An individual firm is negligible relative to the pool and takes its service fraction as given. Its expected payoff from preparation is

$$\pi_i = bq(X,K)x_i - \frac{c}{2}x_i^2. \tag{3}$$

All firms understand the congestion they face. The externality arises because an individual ignores its effect on others, not because it wrongly assumes guaranteed recovery. The social objective counts restored output and real preparation costs. For fixed aggregate preparation, identical convex costs make equal preparation across firms efficient by Jensen's inequality. Consequently, both the symmetric equilibrium and the planner's solution can be expressed using X:

$$W(X,K) = b\min\{X,K\} - \frac{c}{2}X^2. \quad (4)$$

Capacity cost is omitted only while comparing allocations of a fixed pool. Section 4.4 introduces it explicitly. Taxes, reservation payments, and subsidies are transfers within the economy; they are not themselves destroyed resources. The benchmark assumes lump-sum financing and no administrative cost. Thus the welfare argument against a poorly targeted subsidy does not rely on a separate fiscal distortion.

3.3 Analytical and numerical strategy

The analysis first solves the atomistic equilibrium and the planner's problem. It then introduces a proportional subsidy to preparation expenditure, a reservation mechanism, and endogenous capacity. Extensions examine finite strategic players, uncertain surviving capacity, ordinary switching benefits, and heterogeneous values. The closed forms are checked against direct scalar optimization, including kink points where capacity begins to bind. Numerical examples are dimensionless illustrations rather than calibrations to a country, firm, or industry. Table 2 defines the principal notation.

Table 2. Principal notation

Symbol	Interpretation	Domain
p	Probability of common disruption	(0, 1)
v	Value of restoring one workload unit	Positive
$b=pv$	Expected recovery benefit	Positive
c	Preparation-cost curvature	Positive
x_i	Firm i portable workload share	[0, 1]
X	Aggregate portable share	[0, 1]
K	Usable recovery capacity	(0, 1]
$q(X,K)$	Service fraction under proportional rationing	[0, 1]
$R(X,K)$	Recovered workload share	[0, 1]
d	Capacity-cost curvature	Positive
σ	Preparation subsidy rate	[0, 1)

4. Results

4.1 Equilibrium preparation and excess investment

Proposition 1. Under the benchmark assumptions, the unique atomistic equilibrium and the efficient aggregate preparation are

$$X^{\mathrm{D}} = \begin{cases} a, & K \geq a, \\ \sqrt{aK}, & 0 < K < a, \end{cases} \qquad X^{\mathrm{S}} = \min\{a, K\}. \tag{5}$$

Derivation. The private best response is $x_i = bq/c$, which is below one under the maintained parameter restriction. With no congestion, consistency requires $X = a \leq K$. With congestion, substituting $q = K/X$ into the best response gives $X^2 = aK$, valid precisely when $K < a$. The private marginal return is nonincreasing in aggregate preparation while marginal cost increases, so there is a unique crossing. For the planner, welfare rises with derivative $b - cX$ below capacity and falls with derivative $-cX$ above it. The optimum is therefore the smaller of a and K.

When capacity is scarce, both allocations restore the same workload share K, but $X^{\mathrm{D}} > K$. Each private marginal investment competes for service already assigned elsewhere. The excess is a real cost even though no investment is fraudulent and every firm behaves rationally. Figure 1 shows the two regions and the disappearance of the wedge when the fallback pool is sufficiently large.

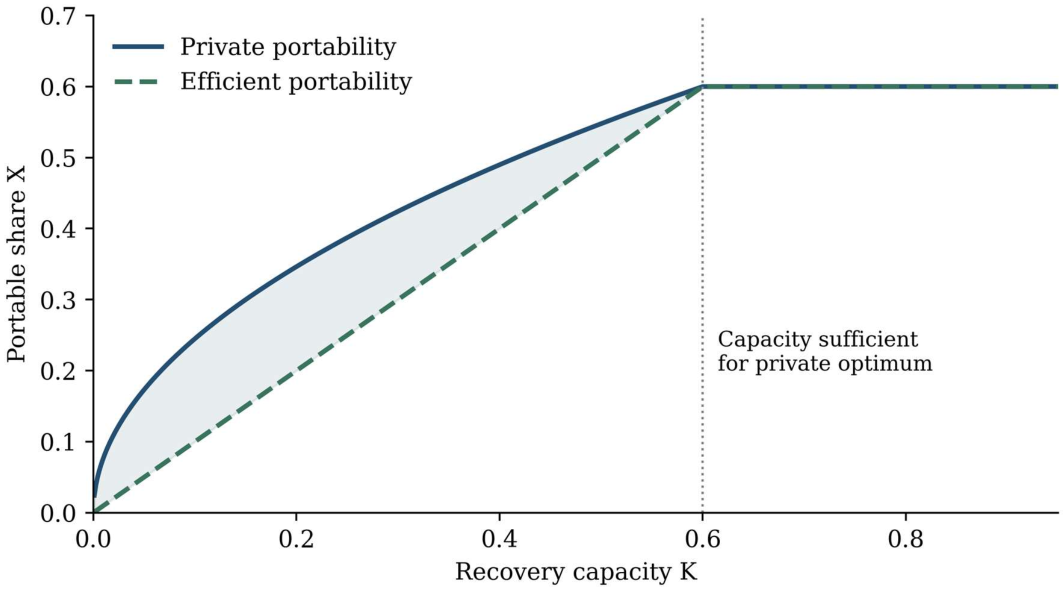

Fig. 1. Private and efficient portability under fixed recovery capacity.

The exact welfare loss for $K < a$ is

$$\Delta W = W(X^{\mathrm{S}}, K) - W(X^{\mathrm{D}}, K) = \frac{bK - cK^2}{2} > 0. \tag{6}$$

The fraction of efficient fixed-capacity welfare dissipated is

$$\frac{\Delta W}{W(X^{\mathrm{S}}, K)} = \frac{a - K}{2a - K}. \tag{7}$$

This ratio approaches one half as capacity approaches zero and falls to zero as capacity approaches a. The absolute loss, however, approaches zero when capacity approaches zero. High relative inefficiency in a tiny pool should not be confused with a large aggregate economic loss. Equation (6) is maximized at $K = a/2$, where the loss equals $b^2/(8c)$. Thus the largest absolute preparation waste arises at an intermediate level of fallback capacity.

4.2 A subsidy can increase preparedness while reducing welfare

Suppose a policy reimburses fraction $\sigma \in [0,1)$ of preparation expenditure. The firm faces effective marginal cost $(1 - \sigma)cx_i$, while society still uses real resources costing $cx_i^2/2$. The complete symmetric solution, including the upper bound on portability, is

$$X_\sigma^{\mathrm{D}} = \min\left\{1, \frac{a}{1-\sigma}, \sqrt{\frac{aK}{1-\sigma}}\right\}. \tag{8}$$

Proposition 2. If $K < a$ and the portability upper bound is slack, a larger subsidy raises preparation, leaves restored workload unchanged, and strictly lowers resource welfare. In particular,

$$\frac{\mathrm{d}X_\sigma^{\mathrm{D}}}{\mathrm{d}\sigma} = \frac{X_\sigma^{\mathrm{D}}}{2(1-\sigma)} > 0, \qquad \frac{\mathrm{d}W}{\mathrm{d}\sigma} = -\frac{bK}{2(1-\sigma)^2} < 0. \tag{9}$$

The slack-bound condition is $\sigma < 1 - aK$. To obtain the welfare derivative, substitute $c(X_\sigma^{\mathrm{D}})^2/2 = bK/[2(1-\sigma)]$ into (4). Additional eligible demand captures a larger private claim on the pool without increasing the pool's total output. Once all workloads are portable, further reimbursement leaves real investment and benchmark resource welfare unchanged, although fiscal transfers continue to rise. Figure 2 illustrates the interior mechanism.

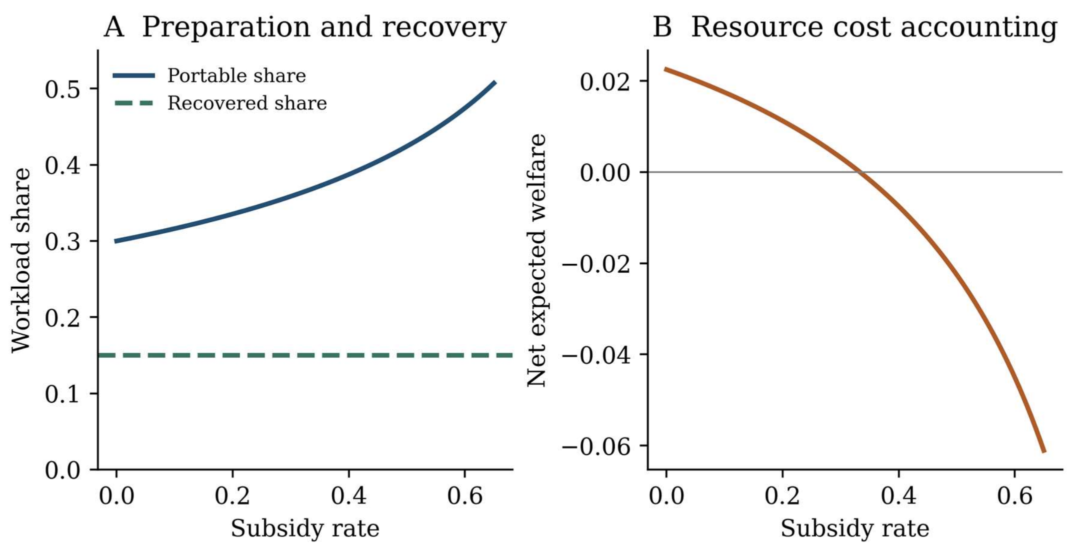

Fig. 2. Portability subsidy: preparation rises while recovery is fixed.

This result does not equate a subsidy with technological progress. If engineering improvements reduce the real cost coefficient from c to c', both private and social cost schedules change. As long as capacity remains binding and the equilibrium is interior, $c'(X^{\mathrm{D}})^2/2 = bK/2$. Preparation expands but aggregate preparation expenditure remains constant, so fixed-capacity welfare is unchanged. With efficient reservations, a genuine cost reduction raises welfare. Normal-time competition, innovation, or interoperability benefits may further strengthen its value. The negative subsidy result therefore applies to a price wedge that leaves the resource technology unchanged, not to every policy that makes switching easier.

4.3 Reservations align eligibility with scarce service

For fixed scarce capacity, an admission price of

$$\tau_K = b - cK > 0 \tag{10}$$

per unit of ex ante preparation implements $X = K$ in the atomistic benchmark. At that allocation, firms receive full service and choose $x_i = (b - \tau_K)/c = K$. Revenue can be returned lump sum without changing marginal incentives. A fixed fee tuned to the correct capacity is a mathematical implementation, but it requires knowledge of costs and expected benefits.

A quantity-backed reservation is more operationally informative. Firms buy enforceable claims on recovery service before the outage. Each prepared workload is paired with a reservation, and outstanding reservations cannot exceed the usable pool. Firms then choose the amount of a bundle consisting of technical preparation and guaranteed service in the modeled event. Competitive clearing of the fixed supply at price τ_K produces the same allocation. This construction requires credible delivery and exclusivity. A financial promise that pays compensation after failed delivery is not automatically a physical reservation.

4.4 Joint investment and the cost of unpriced access

Let creating recovery capacity cost $dK^2/2$, where $d > 0$. Capacity is installed before firms prepare workloads and before the shock. The joint planner maximizes

$$\mathcal{W}(X, K) = b\min\{X, K\} - \frac{c}{2}X^2 - \frac{d}{2}K^2. \tag{11}$$

Proposition 3. The efficient allocation and a competitive reservation price are

$$X^* = K^* = \frac{b}{c+d}, \qquad \tau^* = dK^*, \qquad \mathcal{W}^* = \frac{b^2}{2(c+d)}. \tag{12}$$

Derivation. Any allocation with $X > K$ can reduce preparation without reducing recovery, and any allocation with $K > X$ can reduce capacity without reducing recovery. Therefore $X = K$ at the optimum. Maximizing $bX - (c + d)X^2/2$ yields (12). With reservations, user demand satisfies $cx = b - \tau$, while competitive capacity supply satisfies $dK = \tau$. Market clearing equates the two quantities. Transfers cancel from social welfare. Convex aggregate supply costs can support inframarginal producer rents; efficiency does not require those rents to be treated as a resource loss.

The competitive reservation market is an implementation benchmark, not a prediction about present cloud markets. It assumes credible service commitments, sufficiently competitive capacity provision, no common failure of the reserved pool in the modeled state, and no double sale of the same capacity. Market power or delivery uncertainty can invalidate the price-taking implementation even though the planner's benchmark remains useful.

Now consider a public or collective capacity purchaser that finances capacity but leaves emergency access unpriced. It anticipates firms' equilibrium preparation. In the congested region, its welfare becomes

$$\mathcal{W}^{\mathrm{D}}(K) = \frac{bK}{2} - \frac{d}{2}K^2. \tag{13}$$

If $d > c/2$, the best capacity choice under this institutional constraint is $K^{\mathrm{D}} = b/(2d) < a$. It yields welfare $b^2/(8d)$. The first-best capacity is $b/(c + d)$. The constrained purchaser therefore installs more than the first best when $c/2 < d < c$, the same amount when $d = c$, and less when $d > c$. There is no universal capacity-underinvestment theorem. Unpriced preparation incentives distort both the amount of preparatory effort and the social return to the capacity that attracts it.

The boundary case matters. When $d \leq c/2$, the constrained optimum is $K^{\mathrm{D}} = a$: capacity beyond that point restores no additional workload and incurs additional cost. This follows from the negative derivative $-dK$ in the uncongested region. A capacity expansion may accordingly increase physical recovery while reducing net welfare when its construction cost and the induced preparation costs exceed its benefit.

5. Extensions and Boundary Conditions

5.1 A finite number of strategic firms

The atomistic model approximates a large pool. With N equal-sized firms, define $X = N^{-1}\sum_i x_i$ and physical capacity NK in firm-level workload units. An individual recognizes that it changes the denominator of the rationing rule. For $N \geq 2$, the symmetric equilibrium is

$$X_N^{\mathrm{D}} = \begin{cases} a, & K \geq a, \\ K, & a\left(1-\frac{1}{N}\right) \leq K < a, \\ \sqrt{aK\left(1-\frac{1}{N}\right)}, & 0 < K < a\left(1-\frac{1}{N}\right). \end{cases} \tag{14}$$

In the congested interior, differentiating a firm's own restored share produces the factor $1 - 1/N$. At the capacity kink, the derivative from below is $b - cK$ and from above is $b(1 - 1/N) - cK$, producing the middle region. The payoff is concave on each side and its derivative falls at the kink, establishing the stated best response. A single firm internalizes the entire pool and chooses $\min\{a, K\}$. As the number of independent claimants increases, the congested equilibrium converges to (5). Fragmented demand can therefore strengthen the externality even when total workload and total recovery capacity are held fixed.

5.2 Capacity uncertainty and a general marginal-return wedge

Let $\widetilde{K}$ denote the usable capacity conditional on the common outage. Its distribution may reflect the fact that some potential backups are also impaired. Assume a continuous distribution, positive finite mean, and a positive probability of nonzero capacity. Define

$$R(X) = \mathbb{E}\left[\min\{X, \widetilde{K}\}\right], \qquad q(X) = \frac{R(X)}{X}, \qquad R'(X) = \Pr\left(\widetilde{K} > X\right). \tag{15}$$

An interior atomistic equilibrium and an interior social optimum satisfy

$$cX^{\mathrm{D}} = b\,\frac{R(X^{\mathrm{D}})}{X^{\mathrm{D}}}, \qquad cX^{\mathrm{S}} = b\Pr\left(\widetilde{K} > X^{\mathrm{S}}\right). \tag{16}$$

Proposition 4. Private preparation weakly exceeds efficient preparation. It strictly exceeds it when $\mathbb{E}\left[\widetilde{K}\mathbf{1}_{\{\widetilde{K}<X^{\mathrm{S}}\}}\right] > 0$. A per-unit fee evaluated at the efficient allocation internalizes the wedge:

$$\tau(X) = b\left[\frac{R(X)}{X} - R'(X)\right] = \frac{b}{X}\mathbb{E}\left[\widetilde{K}\mathbf{1}_{\{\widetilde{K}<X\}}\right] \geq 0. \tag{17}$$

Here $\mathbf{1}$ denotes an indicator. The identity follows by separating capacity realizations below and above X. Average recovery includes a firm's chance of capturing already-congested capacity; marginal recovery counts only states in which another prepared unit can expand total restoration. Since $R(X)/X$ is nonincreasing and cX is strictly increasing, the private first-order condition has a unique solution. The planner's objective is strictly concave. Evaluating the private condition at the social solution establishes the ordering.

The sharp zero-extra-output result from fixed capacity becomes a marginal inefficiency result under uncertainty. More preparation can now restore additional output in some states, yet the gain is smaller than the private incentive suggests. With exponential surviving capacity of mean $\kappa > 0$,

$$R(X) = \kappa\left(1 - e^{-X/\kappa}\right), \qquad R'(X) = e^{-X/\kappa}. \tag{18}$$

This distribution is an illustrative smooth case, not an estimated outage law. Its support can exceed one, which is harmless because recovered workload is bounded by $X \leq 1$. Figure 3 shows that the wedge persists without a deterministic capacity threshold.

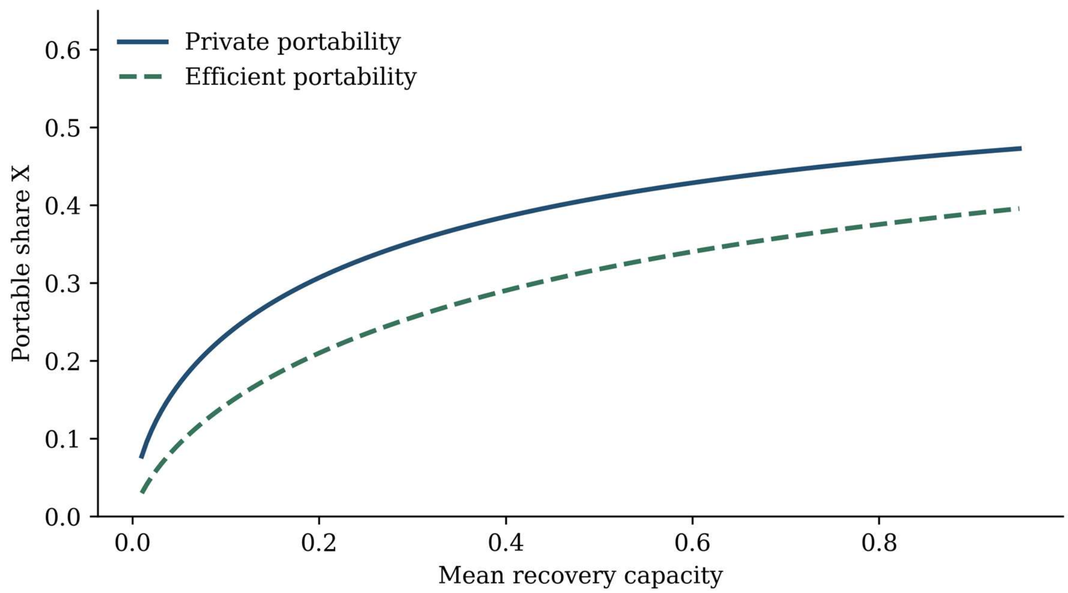

Fig. 3. Private and efficient preparation under uncertain capacity.

5.3 Ordinary benefits and heterogeneous recovery values

Suppose portability also creates ordinary switching value hx_i, where $h \geq 0$ and private and social values coincide. In a congested interior, the private condition becomes $cX^2 - hX - bK = 0$. The planner chooses $X^{\mathrm{S}} = \max\{K, h/c\}$ when $K < a$ and this quantity is below one. The private solution remains strictly larger because competition for recovery adds a positive benefit to every prepared unit. However, a policy that directly increases ordinary switching benefits or reduces genuine resource costs cannot be evaluated from the subsidy result alone. If ordinary benefits differ across firms or spill over to consumers, the full welfare comparison must include those effects.

Heterogeneity adds an allocation margin. Let firms differ in expected recovery value b_i, real cost coefficient c_i, and prepared share x_i. In a fully congested atomistic interior under proportional rationing,

$$x_i = \frac{b_i K}{c_i X}, \qquad X^2 = K \int_0^1 \frac{b_i}{c_i}\, \mathrm{d}i. \tag{19}$$

The socially efficient reservation allocation instead assigns recovered shares $y_i \in [0,1]$ to maximize net value subject to $\int y_i\, \mathrm{d}i \leq K$. Since extra preparation has no additional ordinary benefit in this extension, it sets $x_i = y_i$. The solution is

$$y_i^* = \min\left\{1, \max\left\{0, \frac{b_i - \lambda}{c_i}\right\}\right\}, \qquad \lambda \geq 0. \tag{20}$$

The scarcity multiplier λ clears the capacity constraint when it binds. Reservations therefore address both excessive aggregate preparation and the distribution of service toward workloads with the highest net marginal recovery value. If social losses exceed private willingness to pay for essential services, b_i must incorporate those external benefits; an unadjusted commercial auction need not deliver the social allocation.

6. Numerical Illustrations and Verification

6.1 Transparent parameterization and policy comparisons

The benchmark sets $b = 0.6$, $c = 1$, baseline capacity $K = 0.15$, and capacity-cost coefficient $d = 2$. These normalized values create a congested interior with moderate preparation. They have no empirical interpretation as market outage probabilities or monetary losses. A reader who wishes to split b into probability and event value may do so using any pair satisfying $pv = 0.6$; the benchmark depends only on their product.

Table 3 compares six arrangements using identical resource accounting. Baseline private preparation is 0.3000, twice the capacity that can actually restore service. A 25 percent subsidy raises the portable share to 0.3464 while recovered workload remains 0.1500. Net welfare, including the cost of creating the pool, falls from 0.0225 to 0.0075. A 50 percent subsidy increases preparation to 0.4243 and makes net welfare negative relative to the no-preparation, no-capacity reference. This is an internally generated counterexample, not a prediction of the effect of any existing programme.

Table 3. Normalized policy comparison

Policy	Capacity K	Portable X	Recovered R	Prep. cost	Capacity cost	Net welfare
Open access	0.1500	0.3000	0.1500	0.0450	0.0225	0.0225
25 percent portability subsidy	0.1500	0.3464	0.1500	0.0600	0.0225	0.0075
50 percent portability subsidy	0.1500	0.4243	0.1500	0.0900	0.0225	-0.0225
Reservations with fixed capacity	0.1500	0.1500	0.1500	0.0112	0.0225	0.0563
More capacity without reservations	0.2000	0.3464	0.2000	0.0600	0.0400	0.0200
Joint optimum with reservations	0.2000	0.2000	0.2000	0.0200	0.0400	0.0600

With capacity fixed, reservations deliver the same recovery using preparation of 0.1500, increasing net welfare to 0.05625. The joint optimum has preparation and capacity both equal to 0.2000 and welfare of 0.0600. Expanding capacity to 0.2000 while leaving access unpriced instead produces preparation of 0.3464 and welfare of 0.0200. The latter falls below the baseline despite greater restored output. Here the constrained open-access purchaser already starts at its optimum, because $b/(2d) = 0.15$. This makes the example a precise illustration of institutional complementarity, not evidence that physical backup investment is generally undesirable.

6.2 Strategic internalization and uncertain supply

At baseline capacity, two strategic firms choose preparation of 0.2121, ten firms choose 0.2846, and one hundred choose 0.2985. The efficient level remains 0.1500. These quantities refer to fixed aggregate workload and capacity, so increasing N changes strategic internalization rather than the scale of the economy.

Table 4 uses exponential surviving capacity and excludes capacity construction costs because its mean is fixed within each row. At mean capacity 0.1500, private preparation is approximately 0.2750 and efficient preparation 0.1803. The planner deliberately restores less expected workload, 0.1049 rather than 0.1260, because the extra preparation required for further recovery costs more than its output benefit. Efficient welfare is approximately 0.0467 compared with private welfare of 0.0378. Thus maximizing physical recovery and maximizing net economic value are different objectives even when all recovery is beneficial.

Table 4. Uncertain-capacity illustration

Mean capacity κ	Private X^D	Efficient X^s	Private R	Efficient R	Private W	Efficient W
0.05	0.1703	0.0931	0.0483	0.0422	0.0145	0.0210
0.15	0.2750	0.1803	0.1260	0.1049	0.0378	0.0467
0.30	0.3528	0.2558	0.2074	0.1721	0.0622	0.0706
0.60	0.4287	0.3403	0.3064	0.2597	0.0919	0.0979

6.3 Computational verification and reproducibility

The deterministic audit evaluates 741 parameter pairs: a takes 19 equally spaced values from 0.05 to 0.95, and K takes 39 equally spaced values from 0.025 to 0.975. Closed-form planner values are compared with bounded scalar optimization. The largest objective discrepancy is below 3.1×10^{-9}. For each pair, finite-player equilibria are checked at $N \in \{1,2,5,20,100\}$, producing 3,705 deviation checks. The search explicitly includes boundaries and the capacity kink. The greatest computed profitable deviation is below 1.7×10^{-16}, consistent with floating-point rounding.

The uncertain-capacity calculation uses the same 741 benefit-cost and mean-capacity pairs. Root solutions satisfy the private and social first-order conditions with maximum absolute residual below 1.3×10^{-12}. Direct scalar welfare optimization agrees with the social roots to an objective tolerance of 10^{-10}. In every evaluated case, private preparation exceeds efficient preparation. These checks verify implementation and illustrate the propositions; they do not substitute for proofs or establish empirical external validity. The accompanying replication files contain the complete script, parameter grids, numerical tables, and generated figures.

7. Discussion and Empirical Implications

7.1 Measure feasible recovery rather than nominal alternatives

A resilience assessment should distinguish workloads that are technically portable, the capacity available conditional on a specified disruption, and enforceable prior claims on that capacity. Counting alternative providers or completed migration tests cannot by itself identify the system's recovery constraint. Providers can share dependencies, and several firms can point to the same pool of spare capacity. The model therefore motivates a conditional recovery ratio $R(X)/X$ and a marginal recovery measure $R'(X)$. Their difference, not the count of alternative suppliers, determines the congestion wedge.

This distinction sharpens, rather than replaces, lifecycle governance. A previously valid recovery plan can lose effectiveness as dependencies or provider terms change, consistent with the governance half-life mechanism [14]. In practice, joint tests need to specify the failed service, surviving dependencies, restoration deadline, workload compatibility, and competing demand. A successful individual test establishes that switching is possible under its test conditions. It does not establish simultaneous deliverability across all organizations that rely on the same resources.

7.2 Testable predictions and identification

The model generates several falsifiable predictions. First, portability support should increase preparation more strongly than recovered output when residual capacity is tight. Second, a firm's share of recovered service may rise after it prepares more workloads while the total recovery of its pool changes little. Third, reservation reform should lower aggregate preparation costs per unit actually restored. Fourth, the effect of a common outage should depend on the overlap among fallback destinations, not only on overlap among primary providers. These predictions concern congestion and allocation, whereas the independent benefits of routine switching may follow different patterns.

A useful empirical unit is a recovery pool defined by a destination service, a disruption scenario, and a restoration deadline. Firm-level records would measure prepared workloads, pre-existing reservations, attempted transfers, served workloads, recovery time, and preparation costs. The analysis would also need residual capacity after pre-existing obligations. Provider-wide installed capacity is an inadequate proxy if it is unavailable to the affected firms during the tested event.

A possible design is a randomized phased offer of portability assistance combined with controlled simultaneous recovery exercises. Assignment should vary treatment saturation at the pool level because treated firms can displace untreated firms; an individual-level experiment that assumes no spillovers would miss the mechanism. Outcomes should include pool-wide recovered output and total resource expenditure, alongside firm-level service shares. A second randomized arm offering verified reservations could separate technical eligibility from guaranteed allocation. Such a study is proposed, not conducted in this paper.

For observational analysis, a policy change or contract reform that plausibly shifts preparation costs without directly changing backup supply could be informative. Identification would require evidence against simultaneous capacity expansion, differential outage exposure, and endogenous selection into fallback destinations. Common outages are infrequent and heterogeneous, so precision and transportability may be limited. Table 5 maps the main predictions to measurements and alternative explanations.

Table 5. Empirical predictions and identification risks

Prediction	Measurement	Main alternative explanation
Portability assistance raises preparation more than recovered output under tight capacity	Pool-level prepared workload and recovered output in simultaneous exercises	Capacity was expanded concurrently
Individual recovery shares increase while pool recovery is flat	Firm-level service and pool-level restoration	Provider prioritisation changed
Reservations lower preparation cost per unit restored	Matched reservation and open-access recovery tests	Reserved users differ systematically
Fallback overlap predicts common-shock loss	Dependency graph and destination overlap	Shock severity differs across pools
Saturation treatment effects are negative when capacity binds	Pool-level saturation experiment	Strategic selection into treatment

7.3 Policy scope and limitations

The model supports coupling any claim of greater resilience to a demonstrable increase in feasible recovery or a reduction in its total resource cost. Procurement could purchase tested bundles of portability and recoverable service instead of rewarding technical eligibility alone. A collective operator could also coordinate admission without creating a commercial reservation exchange. What matters economically is that private claims are reconciled with the resource constraint before competing preparation expenditures dissipate surplus.

Several assumptions bound this conclusion. The benchmark concerns a common disruption and a fixed recovery deadline, not everyday migration or a long-run process in which backup supply expands freely. Proportional rationing is one allocation rule. Strict priority, auctions, exclusive contracts, or random admission of whole firms can change investment incentives. Fixed setup costs and indivisible workloads can generate participation thresholds absent from the convex model. Endogenous provider choice and common dependencies between primary and backup suppliers require a richer network model. The exponential illustration holds the capacity distribution fixed and does not identify its dependence on the original shock.

Risk neutrality also matters. If preventing an interruption has discontinuous value, or a firm faces bankruptcy at a service threshold, recovery benefits need not be linear. Likewise, capacity investments may provide normal-time output, learning, or security benefits omitted here. Those benefits belong in a broader welfare objective. The restriction to one modeled planning horizon excludes dynamic entry and technological improvements that portability policies might stimulate. Consequently, the paper identifies a policy failure mode and an implementation benchmark, not a sufficient statistic for all interoperability policy.

Finally, efficient allocation is not synonymous with unrestricted willingness-to-pay allocation. Emergency public functions may carry social value beyond the purchasing firm's revenue. Their priority can be incorporated through social recovery weights or minimum service constraints. Establishing those weights requires institutional judgment and evidence outside this model. Transparent capacity accounting remains necessary under either commercial or public allocation.

8. Conclusion

AI portability and actual recoverability are distinct economic assets. When firms compete for finite backup service, private preparation is rewarded by its share of recovery, while society benefits from the additional output that preparation makes possible. That divergence can produce costly excess preparation even when firms correctly anticipate congestion and every prepared workload is technically valid.

The formal results distinguish three interventions often treated as interchangeable. A subsidy to unchanged preparation technology can deepen waste under binding capacity. A genuine reduction in engineering costs has different welfare consequences. Capacity expansion can also disappoint when it attracts further unpriced preparation. Competitive reservations provide a benchmark that coordinates the two investments, subject to credible delivery and appropriate pricing assumptions. The uncertain-capacity extension expresses the same mechanism through the difference between average and marginal recovery.

The next empirical task is to measure simultaneous recovery and competing claims, rather than infer resilience from nominal alternatives. Portability becomes a reliable contribution to economic resilience when the preparation of workloads, the surviving capacity they require, and the rights governing access are evaluated together.

Declarations

Author contributions: Kwan Hong TAN is the sole author and is responsible for the research framing, interpretation, and final manuscript. Generative artificial intelligence assistance is disclosed below.

Funding: No external funding was received for this study.

Conflicts of interest: The author declares no conflicts of interest relating to this study.

Ethics and consent: This theoretical study involves no human participants, personal data, or animal research. Ethics approval and participant consent are not applicable.

Data and code availability: All numerical results are generated from hypothetical inputs specified in the manuscript. The accompanying replication package includes the Python analysis script, machine-readable parameter outputs, verification results, and figure files. No proprietary or participant-level dataset is used.

References

[1] E. Brynjolfsson, D. Li, and L. Raymond, “Generative AI at Work,” *The Quarterly Journal of Economics*, vol. 140, no. 2, pp. 889–942, 2025, doi: 10.1093/qje/qjae044.

[2] Financial Stability Board, “The Financial Stability Implications of Artificial Intelligence,” Nov. 2024. Accessed: Sep. 19, 2026. [Online]. Available: https://www.fsb.org/2024/11/the-financial-stability-implications-of-artificial-intelligence/

[3] Financial Stability Board, “Monitoring Adoption of Artificial Intelligence and Related Vulnerabilities in the Financial Sector,” Oct. 2025. Accessed: Sep. 19, 2026. [Online]. Available: https://www.fsb.org/2025/10/monitoring-adoption-of-artificial-intelligence-and-related-vulnerabilities-in-the-financial-sector/

[4] E. Brynjolfsson, D. Rock, and C. Syverson, “The Productivity J-Curve: How Intangibles Complement General Purpose Technologies,” *American Economic Journal: Macroeconomics*, vol. 13, no. 1, pp. 333–372, 2021, doi: 10.1257/mac.20180386.

[5] Competition and Markets Authority, “Cloud Services Market Investigation: Summary of Final Decision,” Jul. 2025. Accessed: Sep. 19, 2026. [Online]. Available: https://assets.publishing.service.gov.uk/media/688b20e6ff8c05468cb7b120/summary_of_final_decision.pdf

[6] H. S. Gordon, “The Economic Theory of a Common-Property Resource: The Fishery,” *Journal of Political Economy*, vol. 62, no. 2, pp. 124–142, 1954, doi: 10.1086/257497.

[7] H. Kunreuther and G. Heal, “Interdependent Security,” *Journal of Risk and Uncertainty*, vol. 26, no. 2–3, pp. 231–249, 2003, doi: 10.1023/A:1024119208153.

[8] D. Acemoglu, A. Ozdaglar, and A. Tahbaz-Salehi, “Systemic Risk and Stability in Financial Networks,” *American Economic Review*, vol. 105, no. 2, pp. 564–608, 2015, doi: 10.1257/aer.20130456.

[9] P. Joskow and J. Tirole, “Reliability and competitive electricity markets,” *The RAND Journal of Economics*, vol. 38, no. 1, pp. 60–84, 2007, doi: 10.1111/j.1756-2171.2007.tb00044.x.

[10] M. Reece *et al.*, “Systemic Risk and Vulnerability Analysis of Multi-cloud Environments,” 2023, *arXiv*: arXiv:2306.01862. doi: 10.48550/arXiv.2306.01862.

[11] National Institute of Standards and Technology, “Artificial Intelligence Risk Management Framework (AI RMF 1.0),” NIST AI 100-1, Jan. 2023. doi: 10.6028/NIST.AI.100-1.

[12] K. H. Tan, “The Autonomy Externality: A Welfare-Economic Model of Agentic Artificial Intelligence, Correlated Failure, and Optimal Assurance Policy,” *International Journal of Economics (IJEC)*, vol. 5, no. 2, pp. 458–474, 2026, doi: 10.55299/ijec.v5i2.2022.

[13] K. H. Tan, “Cognitive Monoculture and the Individual-Team Reliance Paradox: A Formal Model of Shared AI Dependence in Organizational Decision Making,” *International Journal of Web of Multidisciplinary Studies*, vol. 3, no. 8, pp. 465–480, 2026, doi: 10.71366/ijwos03082685268.

[14] K. H. Tan, “The Governance Half-Life of Agentic Artificial Intelligence Controls: A Dynamic Assurance Framework for Control Decay, Residual Risk, and Revalidation,” *International Journal of Web of Multidisciplinary Studies*, vol. 3, no. 7, pp. 217–229, 2026, doi: 10.71366/ijwos03072654084.