IT Readiness for Joint Ventures and M&A

Preparing a global IT landscape for partnering or separation:
global programmes as cloaks, and tightened audits for due-diligence preparation.

Abstract

A large international group that may *partner*, contributing a division to a joint venture, being acquired, acquiring or merging, or *separate* one, carving it out to sell, spin off or list, faces a dilemma its IT organisation is rarely told of: the preparation that makes any such move clean takes years, yet preparation aimed visibly at one entity signals it before the board can speak. It dissolves when the preparation is made *symmetric*: one globalisation programme run across every entity, each step justified on its own, designing every entity to be separable and combinable up front, so that no observer can tell which entity will be partnered or separated, or whether any deal is in view. A programme built this way is *confidential by design*. We define an *entanglement index* across six dimensions and show that separation cost, transitional-service duration and the *deal friction*, the counterparty's discount, transitional-service cost and stranded cost, all rise convexly with it. The friction-minimising preparation under uncertainty is the symmetric programme well-run groups run anyway as the globalisation of shared IT: *readiness*, needing no deal; *confidentiality*, revealing none; *safeguard*, keeping vital capability in a core the group keeps; and a *standing due-diligence dossier*. The legal boundary, the Market Abuse Regulation insider lists, the works council's advice right on a transfer of control and disclosure duties, is kept, not evaded; a constructed group of four divisions traces the index before and after, to Day 1.

Keywords: joint venture; mergers and acquisitions; carve-out; IT separation and integration; entanglement; transitional service agreement; stranded cost; real options; signalling; shared services; due diligence; due-diligence dossier; Day 1; works council; Market Abuse Regulation; Netherlands

1. Introduction

A group reshapes its perimeter in two directions. It *partners*, combining a part of itself with another business: a *merger* fuses two groups, an *acquisition* brings a target inside the buyer, and a *joint venture* contributes a part of each of two parents into a jointly-controlled new company. Or it *separates* a part to stand on its own: sold to another owner, spun off, or listed. Both directions run on the same IT act. A *carve-out* is the work of cutting a part loose from the whole so that it can live elsewhere, under a buyer, as a listed company, or inside a joint venture; *integration* is the reverse, attaching a business to a core it did not grow in. A joint venture does both at once: each parent carves out its contribution and the two are combined into one entity that must run from Day 1. In a large international group with several divisions, thirty countries and an IT landscape assembled over decades of acquisitions, the IT part of any of these is usually the longest and least predictable part. The division to be moved shares a directory with the others, runs on servers in a regional data centre, keeps its customers in an ERP instance it shares with a sister division, buys its software under contracts signed in the group's name, and employs, somewhere on its payroll, the fourteen people who happen to know how the group's manufacturing systems work. None of that was designed; all of it has to be undone before the part can stand on its own or stand beside a partner's, and the counterparty, a buyer, a merger partner, a joint-venture partner, will price what cannot be undone by closing day. Practitioners draw a *deal perimeter*, the line that says what is in the deal and what is not, and speak of *NewCo* or *JVCo* for what is formed and *RemainCo* for what stays; in a landscape like this the perimeter runs through shared systems, and a *clean break* at *Day 1*, the first day of the new entity, is not available. The usual remedies are a *transitional service agreement*, a TSA, under which a parent keeps running services for the combined or departed business for a period, and a discount.

The IT organisation is usually told about a deal when it is announced. The reason is good: a deal that becomes known early moves everyone who hears of it. Staff in the division update their CVs, suppliers reopen terms, competitors call its customers, and, if the group is listed, the market trades on what it should not know. Boards therefore keep the intention within a small, ring-fenced deal team, working under a project codename, until they announce that talks are under way. The result is that the preparation that would make the deal clean, which takes years, cannot be started without revealing the intention, and so it is not started. The carve-out, or the integration, begins at announcement, under time pressure, in full view, and the counterparty pays, or credits, for an entangled business.

This paper argues that the dilemma is false, because the preparation can be made invisible by making it symmetric. If every entity in the group receives the same programme, and if every step of the programme has a rationale that stands on its own, then nothing about the programme tells an observer which entity will be in a deal, with whom, or whether any deal is in view. We show that symmetry is not a sacrifice made for secrecy but the economically correct choice under uncertainty about which deal will come, and that the programme which achieves it is one that well-run groups undertake for its own sake: the globalisation of

shared IT. One identity platform, one network and cloud landing zone, one ERP template or integration layer, one service desk, global contracts, and the compliance registers that regulation now requires per legal entity. Done with the right design rules, this programme is at once *readiness*, because a cleanly separable entity is also a cleanly combinable one; *confidentiality*, because its rationale is complete without any deal and its footprint is the same for every entity; *safeguard*, because it moves vital capabilities out of the divisions into a shared core the group keeps whichever division leaves or joins a venture; and the *standing due-diligence dossier*, because what it maintains as ordinary compliance hygiene is what a partner or acquirer's due diligence will ask for, so that the room can be filled in days without a visible scramble.

The paper asks: **how can the IT organisation of a global group prepare to partner any of its entities, in a joint venture, an acquisition or a merger, or to separate one, by sale, spin-off or listing, without signalling which or with whom, and what does that preparation consist of?** Section 2 places the question in the literature on IT in mergers, acquisitions, divestitures and joint ventures, on modularity, on real options and on signalling, and states the legal frame. Section 3 defines the entanglement index. Section 4 gives the model: separation cost, the deal friction, symmetric readiness as a real option, the leakage condition, the capability-location safeguard and the stranded-cost cliff. Sections 5 to 8 describe the programme as readiness, as dossier and as a timeline, and Section 9 works through the partnering side, joint ventures, acquisitions and mergers, in turn. Section 10 reads the Dutch situation, Section 11 traces a constructed group, and Sections 12 to 15 draw implications, answer objections, state limits and conclude.

The paper is written with the Dutch situation in view, for two reasons. Dutch-headquartered groups have, in the last decade, carried out some of Europe's largest deals in public view, mergers, carve-outs and contributions alike, and the record shows what combining a business takes. And the Netherlands gives the works council a statutory right to advise on an intended transfer of control, which a sale, a merger and a joint-venture contribution each are, at a moment when its advice can still make a difference, which fixes a point in time before which the intention is a board matter and after which it is not. Confidentiality, in this paper, is about what is done before that point and how; it is never about moving the point.

2. Background

2.1. IT in mergers, acquisitions, divestitures and joint ventures

The literature on acquisitions has long held that value is created or lost in integration rather than at signing. Haspeslagh and Jemison (1991) distinguished absorption, preservation and symbiosis as integration approaches and argued that the approach should follow the need for strategic interdependence and for organisational autonomy; Capron, Dussauge and Mitchell (1998) and Capron (1999) showed that acquirers redeploy resources in both directions and that the redeployment, not the deal, explains performance; Graebner, Heimeriks, Huy and Vaara (2017) reviewed the process literature since. Information systems entered the picture later. Robbins and Stylianou (1999) found that post-merger systems integration success depended on the acquirer's existing IS capabilities; Wijnhoven, Spil, Stegwee and Fa (2006) mapped IT integration strategies onto the Haspeslagh and Jemison typology; Mehta and Hirschheim (2007) theorised the IS integration decision as an alignment problem; Tanriverdi and Uysal (2011) showed that an acquirer's cross-business IT integration capability predicts the value it creates from acquisitions; Henningsson and Carlsson (2011) and Henningsson and Kettinger (2016) studied how integration deficiencies arise; and Henningsson, Yetton and Wynne (2018) reviewed the field. The consistent finding is that IT integration capability is built before the deal and is what makes the deal cheap.

Joint ventures and alliances have a theory of their own, and it turns on the same asset that this paper does. Kogut (1988) set out why firms form joint ventures rather than contract or merge, when the capability each needs is tacit and lodged in the other's organisation; Hennart (1988) gave the transaction-cost account of the equity joint venture as a way to pool assets that are hard to value and harder to transfer; Inkpen and Beamish (1997) showed that joint ventures are unstable precisely when one parent's contribution can be learnt and then dispensed with, so that bargaining power shifts with what each parent can still withhold; and Dyer and Singh (1998) located the advantage of such ties in relation-specific assets and the routines that govern them. The common thread is that a joint venture stands or falls on what each parent contributes and on whether that contribution can be cleanly defined, valued and run inside a jointly-controlled entity. An entity entangled with its parent is a contribution neither parent can value and the venture cannot run; a defined one is a contribution a partner can credit and JVCo can operate from Day 1.

Divestitures, the route a contribution and a sale share, have a literature too (Brauer, 2006; Moschieri & Mair, 2008; Feldman & McGrath, 2016), and Capron, Mitchell and Swaminathan (2001) showed that acquirers divest assets, their own as well as the target's, as a normal part of reconfiguration, so that a group is on both sides of the deal market at once. The IT side of the carve-out is thinner. Fähling, Leimeister, Yetton and Krcmar (2013) documented an IT carve-out at a multinational, and Yetton, Henningsson, Böhm, Leimeister and Krcmar (2022) showed that the complexity of the IT carve-out project affects the divestor's performance: the harder it is to cut the business loose, the worse the parent does. That finding is the empirical anchor of this paper, and it reads the same for a contribution to a joint venture as for a sale, because both begin with the same carve-out. We take it one step back: if carve-out complexity costs the parent, the parent should lower it before the deal, and the question becomes how to do so without saying why.

2.2. Modularity, asset specificity and real options

Three older ideas supply the model. Simon (1962) and Parnas (1972) described systems whose parts interact strongly within and weakly between, so that one part can be changed, removed, or recombined, without the others;

Baldwin and Clark (2000) showed that design rules which fix interfaces let modules be substituted and that this substitutability has an option value; Sanchez and Mahoney (1996) carried the argument from products to organisations. In this paper's terms, a combinable group is a modular one, and the interfaces are what make a division a module that can be lifted out or plugged in. Williamson (1979) explained why shared assets resist separation: an asset specific to a relationship has less value outside it, and the party that depends on it is held up. An IT capability that sits in one division and serves the whole group is such an asset, and the hold-up falls on whoever does not own it when the division leaves for a buyer or a venture. Myers (1977) and Dixit and Pindyck (1994) gave the vocabulary of real options: an investment that creates the right, but not the obligation, to act later at a known cost, whose value rises with uncertainty about whether and when the act will be wanted. Readiness to combine any entity is such an option, and Section 4 prices it.

2.3. Signals, and the legal boundary

Spence (1973) showed that an action conveys information only when its cost differs across the types taking it; an action every type takes at the same cost separates nobody. Akerlof (1970) showed that a counterparty who cannot verify quality discounts what is offered. Both appear in the model: the first as the condition under which a preparation programme leaks the intention, the second as the discount a buyer, or the credit a joint-venture partner, applies to an entangled business whose separation it cannot assess.

Confidentiality before announcement is legitimate and ordinary, but three bodies of law set its limits, and this paper takes them as given. For a listed group, the EU Market Abuse Regulation (Regulation (EU) 596/2014) defines inside information, requires its public disclosure as soon as possible, allows disclosure to be delayed only where the issuer has a legitimate interest, the delay is not likely to mislead the public and confidentiality can be ensured, and requires the issuer to keep a list of every person with access to it, the *insider list*. A merger, an acquisition or a joint venture under exploration can be inside information, and the deal team is the insider list. Under the Dutch Works Councils Act (Wet op de ondernemingsraden), Article 25 gives the works council the right to advise on an intended decision to transfer control of the enterprise or a part of it, which the contribution of a division to a jointly-controlled venture is as much as a sale, and requires that the advice be sought at a moment when it can still have substantial influence on the decision. The General Data Protection Regulation (Regulation (EU) 2016/679) requires each controller to keep records of its processing activities, and the NIS2 Directive (Directive (EU) 2022/2555) and, for financial entities, the Digital Operational Resilience Act (Regulation (EU) 2022/2554) require risk management, incident reporting and registers of ICT dependencies per entity. Confidentiality, in this paper, is about timing and signal within those duties: the preparation produces per-entity records because the regulations require them, keeps the deal within the insider list, and brings the works council in at the moment the Act names. What it does not do is let the organisation, the suppliers or the market infer the intention from the IT programme before the group speaks.

3. Entanglement

3.1. Definition

We say that an entity is *entangled* with the rest of its group to the degree that it depends on things outside itself without a defined interface to them. The qualification matters and is the paper's central distinction. A division that authenticates its staff through the group's identity platform depends on the group; whether it is entangled depends on whether its accounts are marked as its own, whether the service it receives is described in a catalogue with a unit price, and whether the platform can hand its tenant to another operator. If so, the dependence is *defined*: on the day the division is sold or contributed, the counterparty receives a service it can see, pay for and replace. If not, the dependence is *undefined*: nobody can say which accounts, which cost or which steps, and the carve-out begins with discovery. Sharing is not entanglement. Undefined sharing is.

For entity i and dimension d, let $e_{id} \in [0,1]$ be the share of the entity's dependencies in that dimension that are undefined in this sense, weighted by their operational importance. The entanglement index is

$$E_i = \sum_d w_d \, e_{id}, \qquad \sum_d w_d = 1, \tag{1}$$

Equation 1. The entanglement index: the weighted share of an entity's dependencies that are undefined, over six dimensions.

with weights w_d reflecting how much each dimension matters to the separation of that kind of business. $E_i = 0$ is a module: the entity can be lifted out, and plugged in elsewhere, with every dependence either internal or contracted. $E_i = 1$ is a business whose every dependence must be discovered before it can be cut or combined. Appendix A says how e_{id} is scored.

3.2. Six dimensions

Table 1 gives the six dimensions, what entangles in each, and what a defined interface looks like. The dimensions are chosen because they are the headings of an IT due-diligence questionnaire and of a TSA schedule, which is to say that they are the dimensions along which a counterparty will price entanglement and a parent will have to provide services.

3.3. Sharing is not entanglement

The distinction has a consequence that runs against intuition. A division that runs everything itself has a low index, but the group pays for the duplication and, as Section 4 shows, loses whatever vital capability sits there when it leaves. A division that depends on the group through defined interfaces also has a low index, and the group pays once. Globalisation done without design rules raises entanglement, by pooling everything into one undifferentiated estate; done with them, it lowers it, by converting undefined dependences into defined ones. The rules are the entity attribute, the catalogue with unit prices, the per-entity accounts and the affiliate-transition clause, and they cost little to adopt when a platform is built and a great deal to retrofit.

Table 1. The six dimensions of entanglement: what makes a dependence undefined, and what defines it. A defined dependence is still a dependence; it is one that can be handed over, priced and replaced, whether the entity is sold, contributed or integrated.

Dimension	Undefined (entangled)	Defined (separable and combinable)
Data	Entity's records mixed with others' in shared databases without an entity key; no map of where its personal data lives; master data owned by nobody	Entity attribute on every record and system of record; per-entity data map and record of processing; master-data ownership stated
Identity and access	One directory, one forest, accounts and groups without entity marking; privileged access held by whoever set it up	Global identity platform with the legal entity as an attribute on every account, group and application assignment; privileged access per entity
Infrastructure	Shared data centres, networks and cloud accounts with no per-entity segmentation or cost; addressing and naming group-wide	Landing zone with per-entity accounts or subscriptions, segmented networks, tagged resources and unit cost per entity
Applications and integrations	Shared ERP or CRM instance with shared master data and custom code nobody documents; point-to-point integrations across entities	One template with the entity as an organisational unit; integration layer with per-entity routing; interface register
Contracts and licences	Regional or group contracts in the parent's name with no assignment or affiliate-transition right; licences pooled; *change-of-control clauses* (which let a supplier terminate or renegotiate when ownership changes) unknown	Global contracts with affiliate-transition and transfer clauses; licence register per entity with assignability recorded; change-of-control terms listed
People and competence	Group-wide competence centres housed, by history, on one entity's payroll; key-person knowledge undocumented	Vital shared capabilities on the shared-services entity's payroll; entity-specific capabilities on the entity's; competence register

4. A minimal model

4.1. Separation cost and TSA duration

Let entity i enter a deal: carved out of its parent to be contributed to a joint venture, or sold. The parent pays a *separation cost* S_i to cut it loose, the one-off cost of the carve-out project, and the receiving entity, a buyer or JVCo, takes services under a TSA for a *duration* D_i until it can run them itself. We take both to rise with entanglement, and to rise convexly:

$$S_i(E) = s_0 + s_1 E^{\beta}, \qquad D_i(E) = D_0 + d_1 E^{\beta}, \qquad \beta > 1. \tag{2}$$

Equation 2. Separation cost and TSA duration rise convexly with entanglement.

The constants s_0 and D_0 are the floor: even a module needs a project to lift it and a few months of transition. The convexity $\beta > 1$ is the claim that an entangled business is more than proportionally harder to separate than a half-entangled one, and we give three reasons. First, dependences interact: identity must be separated before data can be, contracts before infrastructure, and an undefined dependence in one dimension blocks a defined one in another, so that the critical path lengthens faster than the number of items on it. Second, discovery compounds: an undefined dependence is found during separation rather than before it, each one found reveals others, and the rate of discovery rises with the density of what is undefined. Third, the people who could resolve an undefined dependence are the same people in every dimension, so that the work queues. Yetton et al. (2022) found that IT carve-out complexity, rather than size, predicts divestor performance, which is what convexity in the index predicts.

4.2. The deal friction

The parent's loss from entanglement is larger than the separation cost, and we call the whole the *deal friction*. It has four parts. The separation cost S_i. The *counterparty discount*: a buyer that cannot see what the separation will cost it, or a joint-venture partner that cannot value an entangled contribution, prices the uncertainty in the Akerlof manner, as a buyer's price cut or as a lower credited value that shrinks the parent's equity share in JVCo, which we write as a share θ of the entity's value V_i scaled by the same convex term. The *TSA friction*: a TSA is usually priced at cost, so the parent runs services for another company, a buyer or JVCo, at a margin that does not cover the management attention, the liability and the delay in restructuring its own cost base that the TSA imposes, at a cost c_T per month of duration. And the *stranded cost* σ_i: the share of the group's shared cost that was allocated to the entity and does not fall when it leaves, the largest of what bankers call the *dis-synergies* of a deal, treated in Section 4.6. Then

$$L_i(E) = s_0 + \big(s_1 + \theta V_i + c_T d_1\big) E^{\beta} + c_T D_0 + \sigma_i. \tag{3}$$

Equation 3. The deal friction: separation cost, counterparty discount, TSA friction and stranded cost.

Everything in the bracket is convex in entanglement, and the bracket is large: for a division worth a billion euro, a discount or lost-credit share of a few per cent is tens of millions, which is the order of the carve-out project itself. Figure 1 draws the friction for the constructed group of Section 11.

4.3. Symmetric readiness as a real option

Now suppose the group does not know which entity a deal will involve, or how. The board has a prior p_i that entity i will enter a deal within the horizon, with $\sum_i p_i \le 1$ because there may be none. A preparation programme lowers entan-

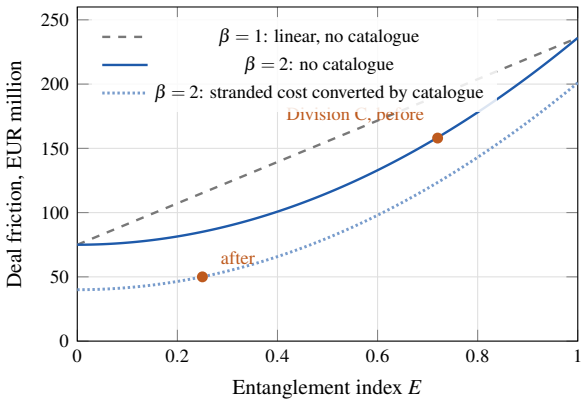

Figure 1. The deal friction as a function of the entity's entanglement index, for the constructed Group M of Section 11: $s_0 + c_T D_0 = 28$, the bracket of equation (3) is 161, and the stranded cost is 47 without a catalogue and 12 with one. The linear case understates the friction at high entanglement; the programme moves Division C from 0.72 to 0.25 and converts most of the stranded cost into a TSA line. Values illustrative.

glement from E_i^0 to E_i^1 for the entities it covers, at a cost K for the shared platforms and c_i for including entity i in them. The expected friction avoided is

$$\Delta = \sum_{i \in \mathscr{P}} p_i \big[L_i(E_i^0) - L_i(E_i^1) \big] - K - \sum_{i \in \mathscr{P}} c_i, \qquad (4)$$

Equation 4. Expected friction avoided by a preparation programme, net of its cost: the value of the real option.

where $\mathscr{P}$ is the set of entities prepared. The programme is a real option in the sense of Myers (1977) and Dixit and Pindyck (1994): it creates the right to put any prepared entity into a deal, to contribute it to a venture or to sell it, at the low exercise cost $L_i(E_i^1)$ rather than the high one, without obliging the group to do anything, and its value rises with the uncertainty about which entity, which deal and when, because the option covers all of them.

Two properties of equation (4) drive the paper. First, the inclusion rule: entity i belongs in $\mathscr{P}$ whenever $p_i [L_i(E_i^0) - L_i(E_i^1)] > c_i$. Once the shared platforms exist, c_i is the marginal cost of one more tenant, one more set of accounts, one more schedule in the catalogue, and it is small against a friction avoided that runs to tens of millions for any entity of size; so the rule includes every entity at any prior that is not negligible, and the optimal programme covers the whole group. Second, the friction is convex in E, so that for a given effort, lowering every entity's index a lot is worth more than lowering one entity's to zero and leaving the rest. The economics of readiness under uncertainty prefer a symmetric programme before secrecy is mentioned; confidentiality then comes free, as the next subsection shows.

4.4. Leakage: when a programme signals

Let an observer, a manager in a division, a supplier's account team, an analyst, hold the prior p_i and watch the programme. The observer sees a *footprint* F_i for each entity: the projects, the staff moves, the contract changes, the data segregation, the hiring freezes, the visitors, that touch entity i and are visible from outside the deal team. The observer updates in the Bayesian way:

$$q_i = \frac{p_i \Pr(F \mid i \text{ is in a deal})}{\sum_j p_j \Pr(F \mid j \text{ is in a deal})}. \qquad (5)$$

Equation 5. An observer's posterior belief about which entity is in a deal, given the programme's footprint.

We say the programme *leaks* to the degree that q differs from p: the observer has learnt something about which entity. Now if the programme is symmetric, so that the footprint vector F is the same whichever entity the board intends to move, the likelihood $\Pr(F \mid i)$ is the same for every i, and equation (5) returns $q_i = p_i$. The posterior equals the prior. The observer has watched the whole programme and learnt nothing about *which*. This is Spence's condition read backwards: an action every entity receives at the same cost cannot separate them.

The same argument answers *whether*. If every step has a rationale that would justify it at $p = 0$, if a group that intended no deal would run the programme anyway, then observing it does not raise the observer's belief that a deal is intended at all. The two conditions together, symmetry across entities and a standalone rationale for each step, are what we mean by *confidentiality by design*, and they are conditions on the programme's design, not on who is told. An asymmetric programme, one that segregates one division's data first or freezes one division's hiring, leaks by construction however few people know why; a symmetric one does not leak however many see it. Appendix B lists what leaks and what does not.

4.5. The capability-location safeguard

The friction so far is the parent's loss on the entity combined. There is a second loss, borne by the entities that remain, which the index does not capture and which is often the larger. Let $\mathscr{C}^*$ be the set of IT capabilities vital to the group as a whole: the people who know the ERP template, the manufacturing-execution competence, the integration platform, the security operations centre, the licence agreements that cover the group. Each capability sits somewhere: on an entity's payroll, under an entity's contract, in an entity's systems. Let $M_{ck} = 1$ if capability c sits in location k, where k ranges over the entities and the shared core. When entity j leaves, sold or contributed to a venture, the group that remains, which we call the *remainer*, loses

$$R_j = \{ c \in \mathscr{C}^* : M_{cj} = 1 \text{ and } M_{c,\text{core}} = 0 \}, \qquad (6)$$

Equation 6. Vital capabilities the remainer loses when entity j leaves: those sitting in j and not in the shared core.

every vital capability that sits in j and nowhere shared. The counterparty, meanwhile, receives everything with $M_{cj} = 1$, which is what it has bought or what the venture now jointly controls.

The asymmetry is the point. Capabilities specific to the departing entity should go with it; the buyer is paying for them, or the venture is built on them. Capabilities vital to the group that happen to sit in that entity, by the accident of which division implemented the system first, walk out with it, and the remainer discovers at Day 1 that it has handed over its own competence centre: a *wrong-pockets* problem

in deal language, an asset that ended up on the wrong side of the perimeter, and Williamson's hold-up on an asset the group depended on but did not, in the relevant sense, own. The hold-up bites harder in a joint venture than in a sale, because the capability does not simply leave: it passes into an entity the group jointly controls with a partner that may be a competitor, where it strengthens the partner's hand in exactly the way Inkpen and Beamish (1997) describe. The *safeguard* is to make $M_{c,\text{core}} = 1$ for every $c \in \mathscr{C}^*$ before the deal: to move every vital shared capability onto the shared-services entity's payroll, contracts and systems. Under the globalisation rationale this is the creation of global centres of excellence, which well-run groups do with no deal in view. The safeguard is itself symmetric, because it moves capabilities out of every division, and so it leaks nothing. Section 11 shows the failure mode it prevents.

4.6. Stranded cost and the TSA trap

A group's shared IT cost, C_{sh}, is allocated to its entities by a key, revenue, headcount or usage; entity j carries a share a_j. When j leaves, the allocation leaves with it but the cost does not. Only the variable fraction v of the shared cost, licences per seat, consumption, falls at once; the rest, data centres, platforms, teams, contracts with minimum terms, stays. The *stranded cost* is

$$\sigma_j = (1 - v)\, a_j C_{\text{sh}} \times H_j, \tag{7}$$

Equation 7. Stranded cost: the fixed share of the shared cost the departing entity carried, times the time it takes to remove it.

the fixed share of what j carried, times the time H_j it takes the remainer to remove it. Without a plan, H_j is years: nobody knows which of the shared cost was j's, the teams that served j are also serving the others, and the contracts run to term. This is the stranded-cost cliff: at Day 1 the remainer's unit cost of IT jumps because the same cost is spread over less business, and it stays up until the cost base is rebuilt.

The TSA is usually offered as the remedy and usually becomes the trap. Under a TSA the parent keeps serving j and bills the buyer or JVCo, so that the stranded cost is covered for the TSA's duration. But a TSA priced at cost, with no exit schedule and no unit prices, holds the remainer's cost base in place, because the people and systems serving it cannot be released; gives the receiving entity no reason to leave, because the service is cheap and the migration is not; and extends the parent's liability and attention into a company it no longer wholly owns, which in a joint venture is a company it co-owns with the partner it is serving. The TSA converts the cliff into a plateau. The same holds in the other direction for a *reverse TSA*, under which the departed business keeps serving the remainer, which is what happens when a capability walked out with it.

The globalised catalogue converts it into a ramp. If every shared service is already in a catalogue with a unit price per entity, for every entity, then the TSA is written by selecting lines from the catalogue, the receiving entity sees what it is paying for and what it would cost to replace, the exit schedule is per line, and the remainer's ramp-down plan is the mirror image of the receiver's migration plan. Stranded cost becomes a TSA revenue line that declines on a schedule

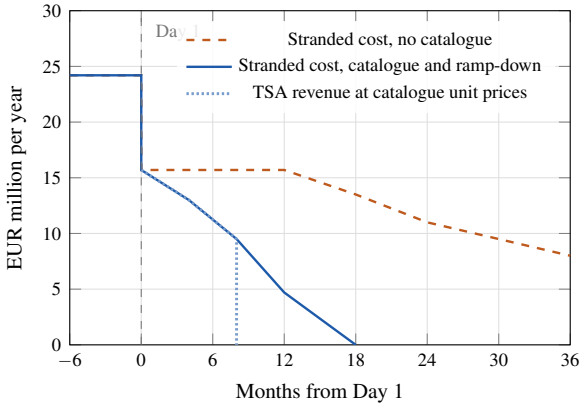

Figure 2. The stranded-cost cliff for Group M's Division C (Section 11). The division carried EUR 24.2 million of shared cost; 35 per cent falls at Day 1 and EUR 15.7 million is stranded. Without a catalogue it decays over years as contracts expire and people leave. With a catalogue the TSA covers it at unit prices for eight months while the remainer removes it on a plan, and it reaches zero at month 18. Values illustrative.

the remainer can plan its cost base against. Figure 2 draws the two cases.

5. Globalisation as readiness and safeguard

5.1. The programme

The programme that lowers every entity's index, relocates vital capabilities and builds the catalogue is not a deal programme. It is the globalisation of shared IT, which groups that have grown by acquisition undertake, sooner or later, because three regional IT departments and four divisional ones cost more than one, secure less and answer no enterprise question without a project. Ross, Weill and Robertson (2006) described the standardised core on which such groups execute; the literature reviewed in Section 2 found that the cross-business IT integration capability built by such programmes is what makes later deals cheap, whichever side of the deal the group is on. Table 2 lists the programme's components, the rationale each has on its own, and what each does to the index and to capability location when it is built with the design rules of Section 3.

5.2. Why it reveals no deal

Each row of Table 2 has a business case that closes at $p = 0$. A group that will never combine anything still saves money and risk by running one directory, one landing zone and one contract per supplier, and the savings are the kind that finance functions approve. The programme therefore passes the test of Section 4: observing it does not raise anyone's belief that a deal is intended. And because every component is applied to every entity, the second test is passed too: nothing in it says which. The confidentiality is not a story told about the programme; it is the programme's own economics.

What makes the difference between a globalisation that reveals no deal and one that does is the third column. Without the design rules, globalisation pools: one directory with no entity attribute, one cloud account for everyone, one

Table 2. The globalisation programme: each component's standalone rationale (the readiness that reveals no deal), the design rule that makes it lower rather than raise entanglement, and its effect on capability location (the safeguard). Every component is applied to every entity.

Component	Standalone rationale	Design rule that defines the dependence	Effect on capability location
One identity platform	Security posture, single sign-on, joiner-mover-leaver control, audit; NIS2 and DORA access-control duties	Legal entity as a mandatory attribute on every account, group, device and application assignment; privileged roles scoped per entity; tenant hand-over procedure documented	Identity engineering moves to the core; divisions keep application-role ownership
One network and cloud landing zone	Cost, security segmentation, exit from regional data centres, consistent operations	One account or subscription per legal entity under a group organisation; tagged resources; segmented addressing; unit cost per entity in the bill	Platform and network engineering to the core
One ERP template or integration layer	Process standardisation, master-data quality, closing speed, one reporting chain	Entity as organisational unit in the template; master data governed with entity ownership; integration layer with per-entity routing and an interface register	Template competence centre to the core; divisions keep process ownership
One service desk and operations	Service quality, cost, follow-the-sun support	Every ticket, asset and configuration item carries the entity; service catalogue with unit price per line per entity	Operations to the core
Global contracts and licence register	Volume pricing, fewer suppliers, one vendor management; licence compliance	Affiliate-transition and transfer clauses in every global contract; assignability and change-of-control terms recorded per contract; licences allocated per entity in the register	Vendor management to the core
Global centres of excellence	Scarce skills pooled, career paths, resilience against key-person loss	Vital shared capabilities on the shared-services entity's payroll with a competence register; entity-specific skills stay in the entity	The safeguard itself
Per-entity compliance registers	GDPR records of processing per controller; NIS2 and DORA registers per entity; ISO/IEC 27001 and SOC 2 scope statements	Registers kept per legal entity as the regulations require, not per region or group	Compliance function to the core; entity evidence stays attributable

ERP instance with shared master data, one contract in the parent's name. Such a programme raises entanglement for every entity, and a group that runs it and later combines a division will find the division harder to cut out, or to integrate, than before. With the rules, the same platforms carry the entity on everything, priced and transferable per entity. The rules cost almost nothing at design time, because an attribute, a tag, an account structure and a contract clause are decisions rather than work, and they cost a retrofit programme later. We at Nashua would state the practical rule as: *never build a shared platform without the entity on it*.

5.3. Why it is a safeguard

The fourth column is the safeguard. A group that has grown by acquisition carries its vital capabilities wherever history left them: the ERP competence centre in the division that implemented first, the manufacturing-systems team in the country with the largest plant, the security operations in the region that had the incident. The globalisation programme moves them to the shared-services entity under the rationale of pooling scarce skills, and the move is complete years before any deal. When a division is then sold or contributed, equation (6) returns an empty set: nothing vital to the group sits only in the division. The counterparty gets the division's own engineers, product-data systems and plant-floor integrations, which is what it paid for or what the venture is built on; the remainer keeps the core.

The safeguard has a cost the group should see. Moving a competence centre out of a division takes its best people from the business that grew them, and the division will object. The objection is legitimate when the capability is specific to the division, and the design rule is to move only what the group as a whole depends on; a group that moves everything has built a central IT monarchy, and a group that moves nothing will hand its own competence into a buyer or a venture. The competence register of Appendix A is where the line is drawn, capability by capability.

6. Compliance as a standing due-diligence dossier

A *data room* is the collection of documents a party makes available to a counterparty for *due diligence*, the examination of what it is buying, contributing or combining with. For the IT part, the questionnaire is predictable: what systems does the business run, on what, under which contracts, with which people, holding which data, certified to what,

with which incidents and which debt. A party that assembles the answers after announcement does so in a scramble, visibly, with consultants, and the scramble is itself a signal to everyone it touches. A party that already holds the answers, per entity, as the by-product of ordinary compliance, fills the room in days from registers it maintained anyway. We call that standing body of evidence the *due-diligence dossier*: not the room, which opens only when a deal is announced, but what is kept ready to fill it. In a joint venture the dossier works both ways, because each parent runs due diligence on the other's contribution and credits it on what it can verify. Table 3 lists, by dimension, the registers the programme maintains and the due-diligence questions each answers.

Three features make the dossier standing rather than occasional. First, the registers are per legal entity because the regulations are: GDPR records of processing belong to each controller, NIS2 and DORA obligations attach to each entity in scope, ISO/IEC 27001 certificates and SOC 2 reports have a stated scope, and a group that keeps them per entity is doing what it is told rather than preparing a deal. Second, the registers are maintained by the functions that own them, the data-protection officer, the security function, vendor management, the service-management team, as part of their ordinary cycle, so that the room is never more than one cycle out of date. Third, the room has an index: a single schedule, per entity, that points to every register, so that the deal team can rehearse the room without anyone outside it noticing. The rehearsal has a cover too: regulators and certifiers audit, and audit readiness is a compliance programme.

The standing dossier also changes the counterparty's price. The discount term $\theta V_i E^{\beta}$ in equation (3) is the counterparty's charge for what it cannot see. A buyer, or a venture partner, that receives, within days of signing a confidentiality agreement, the application inventory, the contract register with assignability and change-of-control terms, the identity export for the entity, the data map, the certificates, the software bills of materials and the incident and debt registers, can estimate its own separation or integration cost and need not price the uncertainty. The dossier does not lower E; it lowers θ, and in a competitive process the difference between a counterparty that can see and one that cannot is the difference between a bid and a discount, or between a fair contribution credit and a haircut. The dossier also feeds the apparatus of a deal process: the *vendor due diligence*, a VDD, the report a party commissions on its own business so that every counterparty works from the same facts; the *Q&A rounds* in which counterparties put written questions; the *red-flag report* a counterparty's advisers write on what they found; and the *clean team*, a ring-fenced group of a competitor-counterparty's advisers who alone may see competitively sensitive data such as customer records and prices, which matters most in a joint venture with a competitor. A party with the registers answers each from the same source; a party without them answers each from a different scramble.

Two registers are not yet universal: the *software bill of materials*, the list of components inside each application the entity runs or has built, which counterparties burnt by inherited vulnerabilities now ask for; and the technical-debt register, the list of what the entity's IT knows it should have fixed, which parties are reluctant to keep and counterparties trust most when they find it, because a party that documents its debt is not hiding it.

7. Confidentiality and its legal boundary

7.1. What confidentiality by design is

Confidentiality, in this paper, is a property of the programme's design: symmetry across entities and a standalone rationale for every step, so that neither the organisation, nor its suppliers, nor the market can infer the intention from what they observe. It is not a property of how many people know, though need-to-know is kept as well, within a deal team that is small, named and, for a listed group, on the insider list. The distinction matters because secrecy by restriction fails in proportion to the number of people who must act on the secret, and a carve-out or an integration requires hundreds. Confidentiality by symmetry lets them act without knowing, because what they are doing is a globalisation programme, which is what they have been told and what it is.

7.2. What is not done before announcement

The discipline of confidentiality is mostly a discipline of not doing. Four classes of action leak by construction and are deferred to the moment the group speaks.

Entity-specific sequencing. Segregating one entity's data, moving one entity to its own accounts first, migrating one entity's contracts ahead of the others. The programme sequences by platform, by region or by wave, never by the entity the board has in mind, even where that costs a few months.

Vendor change-of-control conversations. Asking a supplier whether a licence can be assigned, or what the change-of-control clause means, about one entity, tells the supplier's account team everything. The programme instead puts affiliate-transition and transfer clauses into every global contract as standard terms, at renewal, for every entity, and records the terms of every existing contract in the register as part of licence compliance. The conversation with the supplier about the specific entity happens after announcement, from a register that already says what the clause allows.

Entity-specific insider lists. A listed group's insider list names those with access to the inside information; it is the deal team. Extending it to the division's own management, its IT, its HR, before announcement, widens the circle and, because the people added know why, leaks by the fact of their knowing. The programme is designed so that the people who do the work need not be on the list.

Deal-shaped artefacts. A draft TSA, a *separation blueprint* or a *contribution agreement* with a named partner's or buyer's systems in it, a *cutover plan* for the weekend before Day 1, a Day-1 organisation chart for the venture. These are written after announcement, from the catalogue and the

Table 3. The due-diligence dossier: registers maintained per entity as compliance hygiene, and the due-diligence questions they answer for a buyer or a joint-venture partner.

Dimension	Registers maintained per entity	Due-diligence questions answered
Data	Data map; GDPR records of processing per controller; data-sharing and transfer agreements; retention schedules; master-data ownership	What personal and business data does the entity hold, where, on what basis, shared with whom, and who owns the master data it depends on?
Identity and access	Identity export filtered on the entity attribute; privileged-access register; joiner-mover-leaver evidence; access reviews	Which accounts, groups and privileges belong to the entity, and can they be lifted out as a tenant?
Infrastructure	Asset inventory tagged by entity; landing-zone account structure; network segments; cloud bills per entity; business-continuity and recovery evidence	What does the entity run on, what does it cost, and can it be transferred or rebuilt?
Applications and integrations	Application inventory with ownership, hosting and lifecycle; interface register with per-entity routing; software bills of materials for built applications; technical-debt register	What applications does the entity depend on, which are shared, which are its own, what is inside them and what is known to be wrong?
Contracts and licences	Contract register with term, assignability, affiliate-transition and change-of-control clauses; licence register allocated per entity; supplier risk assessments; TSA catalogue with unit prices	Which contracts transfer, which must be renegotiated, what licences the entity holds, and what the TSA will cost?
People and competence	Competence register; capability location; key-person list; organisation chart of IT serving the entity; open positions	Who comes with the entity, who stays, and what does the entity lose or the remainer lose?
Cross-cutting	ISO/IEC 27001 certificate and statement of applicability with entity scope; SOC 2 reports where services are provided; NIS2 and DORA registers and incident reports per entity; security assessments; incident register	What is the entity's security and resilience posture, with evidence, and what has gone wrong?

registers, in days, which is why the catalogue and registers exist.

7.3. The boundary

Confidentiality is about timing and signal; it is never a means of avoiding a duty. Three duties fix the boundary and the programme is designed to meet each.

For a listed group, inside information must be disclosed as soon as possible unless the conditions for delay are met and documented: a legitimate interest, no likelihood of misleading the public, and confidentiality ensured. An exploration of a deal that is not yet decided can meet those conditions, and the deal team's insider list is what "confidentiality ensured" means in practice. The moment the conditions fail, because the information has leaked or because a decision has been taken that the market must know, the group discloses. Nothing in the programme changes that moment; what the programme changes is that the IT preparation is not the thing that causes the conditions to fail.

Under the Dutch Works Councils Act, the works council must be asked to advise on an intended decision to transfer control of the enterprise or part of it, which the contribution of a division to a jointly-controlled venture is, at a time when its advice can still have substantial influence, with the reasons, the expected consequences for staff and the measures envisaged. The globalisation programme is itself subject to the council's rights where it changes the organisation, and is presented to the council as what it is; what the council is not told, before the board has an intended decision, is an intention the board does not yet have. The programme does not shorten the council's time; if anything it lengthens it, because a separable business gives the council a clearer picture of consequences than an entangled one.

Disclosure duties to a counterparty, under the warranties a party gives in a sale or contribution agreement, are met rather than strained by the standing dossier, because a party that has kept the registers knows what it is warranting.

8. The readiness timeline

Table 4 sets out what is done when, from three years before announcement to Day 1, with the standalone rationale of each step and what is deliberately not done. The horizon of three years is set by the slowest components: a global contract portfolio renews over two to three years, an identity platform takes a year to build and a year to migrate, and a competence centre moved too fast loses the people it was moved to keep. A group that starts later gets less of the friction back, roughly in proportion to how far down the index it has come, which is why we at Nashua would advise any group that might ever combine anything to start now.

The timeline has one feature practitioners resist: the programme is finished for every entity, including those the board has no intention of combining. Equation (4) answers the cost objection, and the feature is also what makes the timeline safe to show inside the group, which is a test worth applying to every step: if a step could not be shown to every division's management without revealing the intention, it is not symmetric and belongs after T_0.

Table 4. The readiness timeline. T_0 is the announcement that talks are under way; Day 1 is the first day of the combined or departed entity. Every step before T_0 is applied to every entity and has the rationale shown; the last column is what is not done because it would signal.

When	Done, for every entity	Standalone rationale	Not done before T_0
$T-36$	Board approves the globalisation programme; deal team of a handful formed; for a listed group, insider list opened when the exploration becomes inside information	Cost, security and operating-model review of a group grown by acquisition	No entity named in any programme document
$T-30$	Design rules fixed: entity on everything, catalogue with unit prices, per-entity accounts, standard affiliate-transition clause; capability register drawn up	Architecture principles	No ranking of entities by separability
$T-24$	Global centres of excellence; vital shared capabilities moved to the shared-services entity; identity platform and landing zone built	Pooling scarce skills; security; exit from regional data centres	No move of entity-specific skills; no hiring freeze anywhere
$T-18$	Global contracts renewed with the standard clauses as they fall due; contract and licence registers per entity with assignability and change-of-control terms	Volume pricing; licence compliance	No supplier asked about any entity's transfer
$T-12$	Identity migration complete; per-entity data maps and records of processing; ERP template rolled by wave; catalogue priced per entity	Regulatory compliance; master-data quality	No entity-specific segregation; waves set by region or platform
$T-9$	Audit-readiness rehearsal: dossier index populated per entity from the registers; certifications renewed with per-entity scope	Certification and audit cycle	No deal advisers on site; no draft TSA
$T-6$ to $T-1$	Programme completes; stranded-cost analysis per entity in the catalogue's cost model; deal team rehearses the room for every entity	Cost transparency	No counterparty contact beyond the deal team; works council not yet asked: no intended decision
T_0 to Day 1	Announcement; insider list extended as the law requires; teaser and information memorandum to counterparties under non-disclosure agreements; data room opened within days; non-binding offers or heads of terms, then due diligence from the registers and binding terms; TSA from the catalogue with an exit date per line; separation sequenced by dimension; works council asked to advise on the intended decision; supplier assignments from the contract register; signing, conditions precedent including merger clearance, closing	,	,
Day 1	Entity on its own accounts and tenant, or landed as a module on the counterparty's core; TSA lines at unit prices; remainer's ramp-down live; set of equation (6) empty	,	,

9. The partnering side

The model so far has priced the carve-out: cutting an entity loose from its parent. Partnering is that carve-out followed by a join, and the three forms it takes, a joint venture, an acquisition and a merger, differ in what the join demands. This section works through each.

9.1. Joint ventures

A joint venture is the two-sided case, and the one the carve-out literature treats least. Each parent contributes a part of itself into JVCo, a new company the two control together, and JVCo must run from Day 1. Industrial groups form such ventures to reach a scale neither party has alone and to reconfigure a footprint neither could move alone, as when two sub-scale product lines combine and their volume assembly consolidates in Central and Eastern Europe to escape the labour costs of Western European plants; the IT readiness this paper describes is what lets the contributed business be cut loose from its parent and re-sited inside the venture. The contributing parent faces the full deal friction of equation (3): it carves the division out at cost S_i, a TSA runs from the parent to JVCo while JVCo stands up its own IT, stranded cost stays in the parent, and the partner credits the contribution on what it can verify, so that entanglement the partner cannot assess becomes a lower credited value and a smaller equity share. A joint venture adds two demands a sale does not. First, the contribution must not only separate cleanly from the parent but combine cleanly with the partner's: two landing zones, two identity platforms and two ERP templates have to interoperate, or be chosen between, on Day 1, and defined interfaces on both sides are what make that a project of weeks rather than years. Second, JVCo is jointly governed, and joint governance runs on numbers both parents trust; the per-entity catalogue is JVCo's

opening cost model, the basis of the TSA each parent gives it, and the evidence that neither parent is cross-subsidising the other through the venture. A group that has globalised with the design rules contributes a priced, tenant-shaped, register-backed business; a group that has not contributes a dispute.

The capability safeguard of Section 4 is sharper here than anywhere. When a division is sold, a vital capability that walks out with it is lost to a third party; when a division is contributed, it walks into an entity the group co-owns with a partner that may compete with it elsewhere, and, as Inkpen and Beamish (1997) showed, the partner that can learn and absorb the other's contribution gains the bargaining power to reshape or end the venture on its own terms. A group that has moved its vital shared capabilities into the core before contributing keeps them whatever the venture's fate; a group that has not has taught a competitor how its manufacturing systems work and called it a partnership.

9.2. Acquisitions

The argument mirrors for acquisitions. A group that buys a business faces the integration problem the literature of Section 2 describes, and the finding there is that the acquirer's integration capability, built before the deal, determines the value created (Tanriverdi & Uysal, 2011; Henningsson et al., 2018). The globalised core is that capability made concrete. A group with one identity platform, one landing zone, one template and one catalogue does not merge two landscapes when it buys; it lands the acquired business on its own: a tenant in the identity platform, an account in the landing zone, an organisational unit in the template, a set of lines in the catalogue, each with the entity on it from the first day. The acquired business is *bolted on*, or in the smaller case *tucked in*, attached as a module to a core that was built to take modules, and its own entanglement with its former parent is cut by the same registers, now read from the other side of the table: the integration cost the acquirer bears rises with the target's entanglement and falls with the acquirer's own readiness to receive a module.

Puranam, Singh and Zollo (2006) described the dilemma of coordination against autonomy in technology acquisitions: integrate and lose what made the target valuable, or leave it alone and lose the synergy. The globalised core with design rules offers the middle that the dilemma usually lacks, because it integrates the layers that are shared, identity, infrastructure, contracts, service, while leaving the acquired business's own applications and people where they are, as an entity on the platform rather than a landscape to be merged. Haspeslagh and Jemison's (1991) preservation and absorption become a choice per layer rather than per deal.

9.3. Mergers, and the symmetry that holds throughout

A merger of equals is the case where both landscapes are in question at once, and neither is simply the module the other receives. Here the globalised core is tested as a whole: the combined group chooses one identity platform, one landing zone and one template from two, and the group whose estate was built with the design rules contributes the core the other lands on. Only a *platform deal*, an acquisition or merger so large that the acquirer's own landscape is the one in question, puts the core itself to the test in this way. Across all three forms the symmetry argument holds in reverse: an integration programme shaped around one target, or a contribution programme shaped around one division, signals as surely as a separation programme shaped around one entity to be sold, and a core that can take or give any module reveals nothing about which it will.

10. The Dutch situation

10.1. Deals in public view

Dutch-headquartered groups have, in the last decade, carried out several of Europe's largest deals, and the public record of each illustrates a part of the argument. Philips separated its lighting business and listed it as a standalone company on Euronext Amsterdam in May 2016 (Philips Lighting, 2016): the *IPO route*, which a *dual-track process* keeps open beside a *trade sale* to a single buyer, and the one that demands the most of a carve-out, because a listed company must stand entirely alone; a division of a century became a listed company with a landscape of its own, in the open, as announced strategy. Philips later sold its domestic-appliances business, announced in March 2021 and completed that September, with a fifteen-year brand-licence agreement as part of the transaction (Philips, 2021): a dependence of the departed business on its former parent converted into a defined, priced and time-limited contract, which is what this paper means by a defined interface. AkzoNobel completed the sale of its specialty-chemicals business to The Carlyle Group and GIC on 1 October 2018 for an enterprise value of EUR 10.1 billion (AkzoNobel, 2018), and the business was relaunched as Nouryon days later (Nouryon, 2018); the separation had been announced as strategy, so that confidentiality was not the question and separation time was. Unilever's tea business, operated as ekaterra before its sale, was sold to CVC Capital Partners, completed on 1 July 2022 (Unilever, 2022), the business having been constituted as a standalone entity ahead of the sale, which is the visible route to a low index. And DSM and Firmenich combined in May 2023, a merger executed as a contribution of Firmenich's shares into the Dutch group (dsm-firmenich, 2023): a combination of a Dutch and a Swiss group in which the task was integration and the question of Section 9 had to be answered, and in which the contribution structure is the very mechanism a joint venture uses.

None of these groups has published an entanglement index, and we at Nashua do not claim to know how any of them prepared. What the record shows is the scale of the task: deals of this size take a year or more from announcement to completion, carry TSAs, and are priced or credited by counterparties who can see what they are receiving. The paper's claim is that the years before announcement are where the index is lowered, and the record is consistent with groups that had already made a business separable in substance obtaining cleaner outcomes than groups that began at announcement.

10.2. The works council's moment

The Dutch Works Councils Act gives the works council the right to advise on an intended decision to transfer control of the enterprise or a part of it, which a sale, a merger and a contribution to a joint venture each are, and requires that the advice be sought at a time when it can still have substantial influence on the decision to be taken, with a statement of the reasons, the expected consequences for staff and the measures envisaged. This fixes the point at which the intention becomes a matter for the organisation rather than the board: when there is an intended decision. Before that point the board may explore, and must keep inside information within its insider list if it is listed; at that point the council must be asked, with enough time and enough information for its advice to matter.

The timing has two consequences for the programme. First, the globalisation programme is itself subject to the council's rights where it reorganises IT, and it is presented to the council on its own merits, which it has; a programme that could not be presented to the council honestly is not symmetric and should not be run. Second, a separable business gives the council what the Act requires: a statement of consequences that can actually be made. The council of an entangled division is asked to advise on a transfer whose consequences for its staff depend on a separation nobody has yet scoped; the council of a division with a low index is told which systems, contracts and people go, which stay, and under what TSA, because the registers say so. Confidentiality by symmetry, correctly done, serves the council's right rather than circumventing it.

10.3. The Dutch holding as a case in point

Dutch-headquartered international groups are, more often than most, holdings assembled by acquisition across many countries, with regional IT departments that grew up around the acquisitions and divisional IT that came with them; they are also, more often than most, listed in Amsterdam under the Market Abuse Regulation, with works councils at the Dutch entities. This combination of circumstances makes the Dutch holding the clearest case for the argument: a group that may at any time contribute, buy, be bought or merge, that cannot say so until it must, and whose entanglement was never designed. The constructed case that follows is such a group.

11. A constructed case

The case is constructed, not observed. It is not a report on any company and its figures are illustrative and internally consistent.

11.1. Group M as it stands

Group M is a Dutch-headquartered industrial group with revenue of EUR 6.0 billion in four divisions: A, industrial pumps (35 per cent of revenue); B, valves (25 per cent); C, filtration (20 per cent); and D, services (20 per cent). It operates in about thirty countries. IT is run by three regional departments, Europe, the Americas and Asia-Pacific, that grew around the acquisitions in each region, and by divisional IT in each division. Annual IT cost is EUR 180 million, of which EUR 110 million is shared (regional and group) and EUR 70 million divisional. The board has concluded that C, filtration, is sub-scale on its own and carries its assembly in high-labour-cost Western European plants, chiefly in Germany, and intends to contribute it within two years into a 50/50 joint venture with a competitor's filtration arm: a business large enough to compete, and one whose industrial logic is to consolidate the two parties' volume assembly in Central and Eastern Europe, where labour costs a fraction of what it does in the west, while their engineering and systems combine. A bolt-on acquisition in services for D is a possibility later; nobody outside a deal team of seven knows. The deal team's own priors, over three years, that an entity will enter a deal are $p_C = 0.6$, $p_D = 0.25$, $p_A = 0.1$ and $p_B = 0.05$.

The landscape is the usual one. One directory forest covers the group, with accounts and groups that carry no entity. The regional data centres host everyone, with networks addressed group-wide and a cost allocation by headcount. A and B share one ERP instance; C and D share another, with D's field-service data in C's company codes because D was once a part of C. Contracts are regional, in the regional entity's name, with assignability unknown and change-of-control clauses unread. The group's manufacturing-execution competence centre, fourteen people in Germany, sits on C's payroll because C's plants were the first to be automated; A's and B's plants depend on it. The assembly those systems run can move east to lower its labour cost; the knowledge of the systems cannot, and is exactly what the safeguard must keep in the core as the footprint shifts. The ERP competence centre sits in B for the same reason.

Table 5. Group M: entanglement index by dimension and division, before and after the globalisation programme, equal weights. Every division moves; none moves alone. Values illustrative.

	Before				After			
Dimension	A	B	C	D	A	B	C	D
Data	0.6	0.7	0.7	0.8	0.3	0.3	0.3	0.3
Identity	0.9	0.9	0.9	0.9	0.2	0.2	0.2	0.2
Infrastructure	0.8	0.8	0.8	0.8	0.2	0.2	0.2	0.2
Applications	0.5	0.7	0.6	0.7	0.4	0.5	0.4	0.5
Contracts	0.7	0.8	0.8	0.6	0.2	0.2	0.2	0.2
People	0.6	0.7	0.5	0.3	0.2	0.2	0.2	0.2
Index E	0.68	0.77	0.72	0.68	0.25	0.27	0.25	0.27

11.2. Before and after

Table 5 scores the index for all four divisions before the programme and after it, with equal weights. The scores are what the deal team's rehearsal of Appendix A would produce. Before the programme, every division is between 0.68 and 0.77: entangled in identity and infrastructure almost completely, in contracts heavily, and in data and applications according to which instance it shares with whom. After the programme, every division is between 0.25 and 0.27, and the residual is in applications, because the ERP template is still one instance and migrating a company code out of it is a project, and in data, because master data shared with a sister division must be split.

Table 6. Group M: the deal friction of contributing Division C before and after the programme, EUR million except duration. The programme's cost is shared by four divisions and largely carried by its own savings.

	Before	After
Entanglement index E_C	0.72	0.25
Separation cost S_C	60	29
Counterparty discount $\theta V_C E^{\beta}$	39	5
TSA duration D_C (months)	25	8
TSA friction $c_T D_C$	12	4
Stranded cost σ_C	47	12
Deal friction L_C	158	50
Friction as share of V_C	10.5%	3.3%

The programme is the one of Table 2, run from $T-36$ on the timeline of Table 4, under its own business case: EUR 48 million over three years, of which EUR 40 million for the shared platforms and EUR 2 million per division for inclusion, against savings of EUR 12 million a year from $T-12$ as the regional data centres close and the contract portfolio consolidates. The programme would have been approved with no deal in view, and at Group M it was presented, to the executive committee, the works council and the divisions, as exactly that.

11.3. The capability that would have walked

At $T-24$ the group creates global centres of excellence under the programme and moves the manufacturing-execution competence centre from C's payroll to the shared-services entity, and the ERP competence centre from B's. C keeps the six engineers who maintain its own filtration plants' integrations, because those are specific to C; the eight who serve A's and B's plants move. Had the move not been made, equation (6) at Day 1 would return the manufacturing-execution competence for A and B: fourteen people, now on the payroll of a joint venture half-owned by a competitor, serving the remainer's plants under a TSA with that venture, until the remainer rebuilt a capability it had owned for a decade, and with the competitor-partner learning the group's manufacturing systems from the inside meanwhile. Because the move was made two years earlier, under a rationale the divisions accepted, the set is empty. The ERP competence centre moving from B is the same safeguard for a division nobody intends to combine, and it is what makes the move symmetric: B's management saw a pooling of skills, not a preparation.

11.4. Economics of the contribution of C

Table 6 gives the friction for C with the model's parameters: $\beta = 2$, $s_0 = 25$, $s_1 = 68$, $\theta = 0.05$ on a value V_C of EUR 1.5 billion, $D_0 = 6$ months, $d_1 = 36$ months, $c_T = 0.5$ per month, and a stranded cost from C's 22 per cent share of the EUR 110 million shared cost with a variable fraction of 35 per cent, removed over three years without a catalogue and over eighteen months with one, half of it covered by TSA revenue meanwhile. Here the discount term is the credit C's contribution loses in the partner's valuation, which comes straight off Group M's equity share in JVCo. The friction on C falls by EUR 108 million. Across the four divisions, with the deal team's priors and the corresponding reductions for A, B and D (about 130, 110 and 95 respectively, from their own values and shares), the expected friction avoided by equation (4) is about EUR 107 million against a programme cost of EUR 48 million, before the programme's own savings; with the savings, the programme pays for itself by T_0 and the option comes free. That is confidentiality in numbers: the programme's rationale at $p = 0$ is positive, and the option on top of it is worth twice its cost.

11.5. The confidential timeline at Group M

The deal team is formed at $T-36$; because the group is listed and the exploration is inside information, the seven are on the insider list from the moment the board minutes record the intention to explore, and nobody else is. The programme is announced at $T-34$ with its business case and run by region and by platform, in waves the regional CIOs sequence on operational grounds: Europe first because its data-centre lease ends soonest, then the Americas, then Asia-Pacific. C is in the European wave with A and B. No document names C; no supplier is asked about C; no hiring is frozen; the works council sees a globalisation programme, advises on it, and is told the truth about it. At $T-9$ the compliance function runs an audit-readiness exercise in which the dossier index is populated for every division, and the deal team rehearses C's room within the exercise's scope. At $T-3$ the board authorises the process and engages advisers, who draft the teaser and the information memorandum under the insider list and the codename. At T_0 the group announces that it is in talks to combine C with a partner's filtration business in a joint venture.

11.6. From T_0 to Day 1

The partner and its advisers receive the data room eleven days after signing non-disclosure agreements, through a clean team for C's customer and pricing data because the partner is a competitor: for C, the application inventory, the contract register with assignability and change-of-control terms for thirty-four global and nineteen local contracts, the identity export on C's entity attribute, the data map and records of processing for C's eleven controllers, the ISO/IEC 27001 certificate with C in scope, software bills of materials for C's four built applications, and the incident and technical-debt registers. The TSA from the parent to JVCo is drawn from the catalogue: eleven lines, identity, network, hosting, ERP hosting of C's company codes, service desk and six application services, each at the unit price every division already pays, each with an exit date. The works council of the Dutch entity is asked to advise on the intended decision to transfer control of C to the venture, with the consequences for staff set out from the registers. Three Q&A rounds are answered from the registers, the partner's red-flag report finds nothing the technical-debt register had not already said, and the contribution is valued on the dossier rather than on a discount. Thirty-one global contracts are assigned to JVCo under their affiliate-transition clauses; three are renegotiated. Signing is at $T_0 + 4$ months, closing at $T_0 + 7$ after merger clearance, and the signing-to-closing period is spent on the cutover plan rather than on

discovery.

At Day 1, C runs inside JVCo on its own identity tenant and cloud subscriptions, both transferred; its company codes stay on Group M's ERP template for eight months under the TSA and are then migrated to JVCo's chosen platform; its contracts are the venture's. The remainer keeps the identity platform, the landing zone, the template, the global contracts and both competence centres, and loses nothing it depends on, nothing that passes to a competitor through the venture. The stranded EUR 15.7 million a year is covered by TSA revenue for eight months and removed on the ramp-down of Figure 2, reaching zero at month 18 as the last European data-centre hall closes and the shared teams are resized. Nothing about Day 1 was improvised, and nothing before T_0 was visible.

12. Implications

Put the entity on everything. Every shared platform, account, record, contract line and catalogue line carries the legal entity it serves. This is the cheapest design rule in IT and the most expensive to retrofit, and it is the difference between globalisation that defines and globalisation that pools.

Build the catalogue before it is needed. A service catalogue with a unit price per line per entity is the TSA, the ramp-down plan, the stranded-cost analysis and a joint venture's opening cost model, written in advance for every entity. It is also ordinary cost transparency, which is its cover.

Move what is vital, leave what is specific. Draw the capability register, and move to the shared core every capability the group as a whole depends on, from every division, under the rationale of pooling. Leave in each division what is specific to it. The register is where the argument with the divisions is settled, and where a group decides what it will never hand into a venture.

Sequence by platform, never by entity. Waves set by region, platform or lease expiry are symmetric; waves set by which entity the board has in mind are not. Where the symmetric sequence costs months, pay them.

Keep the registers as the law reads them. Records of processing per controller, security and resilience registers per entity, certificates with entity scope. The regulations ask for it, and the dossier is the by-product.

Test every step for symmetry. If a step could not be shown to every division's management without revealing the intention, it is not symmetric and belongs after T_0.

Price the friction. A board deciding whether to prepare should see equation (4): the friction avoided per entity, the priors, and the programme's cost net of its own savings. In our experience the option is worth a multiple of its cost for any group with a division it might one day move.

13. Objections and replies

"The programme is itself the signal." A group with a division the market already considers a deal candidate may fear that any programme will be read as preparation. The reply is the test of Section 4: a programme that a group would run at $p = 0$ raises nobody's belief that a deal is intended, and the market's belief about the candidate division is what it was before. What would raise it is asymmetry: a programme that touches that division first or differently. The symmetric programme does not lower the market's prior; it declines to move it.

"Counterparties want a standalone business, not one that shares." Counterparties want a business whose dependences are defined. A financial buyer with no platform of its own wants a TSA with unit prices and exit dates, which the catalogue gives; a strategic buyer or a venture partner with a platform wants to land the business on it as a module, which the entity attribute gives. Neither wants a business that runs its own duplicate of everything at a cost it will have to cut, and neither wants a business whose dependences must be discovered. Sharing is not the problem; undefined sharing is.

"The works council will see through it." The council is told the truth about the programme, which is a globalisation programme, and is asked to advise on a transfer of control when there is an intended one, as the Act requires. Nothing is hidden from the council that the council has a right to know at the time it has the right to know it. What the programme withholds, before the board has an intention, is an intention the board does not have.

"Preparing entities we will not move is waste." The inclusion rule of equation (4) answers this: the marginal cost of including an entity on platforms that exist is small, the friction avoided is large at any plausible prior, and the programme's own savings carry most of the cost. A group that prepares only the entity it means to move has saved a little and told everyone.

"Entanglement is where the synergy is." The synergy of a group is in sharing: one platform, one contract, one competence centre. None of that is lost by defining the dependence; the entity attribute, the unit price and the transition clause take nothing from the sharing. What is lost is the ability of a shared estate to hide what belongs to whom, and that was never a synergy.

"This is secrecy, and secrecy is suspect." Confidentiality before a decision is lawful, ordinary and, for a listed group, regulated in detail: inside information is disclosed as soon as possible unless the conditions for delay are met, and the insider list says who knows. The programme operates inside those rules and makes them easier to keep, because it needs fewer people to know. It is the opposite of trading on information; it is preparing in a way that gives nobody information to trade on.

14. Limitations and research agenda

This is theoretical research. Five limitations follow, each with a test.

First, the convexity β is asserted, not measured. It is the model's decisive parameter: with $\beta = 1$ the friction is linear and the case for lowering the index everywhere weakens to a case for lowering it where the prior is highest. Yetton et al. (2022) found that carve-out complexity predicts divestor performance; a study that scored the index of carved-out entities retrospectively, from the TSA schedules and separation budgets that parents keep, and regressed separation cost and TSA duration on it, would estimate β. Our prediction is that it is well above one.

Second, the index is a construct, and its weights are judgement. Appendix A says how it would be scored; whether the scores predict anything is the empirical question. The test is whether the index, scored before announcement, predicts the TSA's duration and the discount a counterparty applies, across a sample of carve-outs, contributions and integrations.

Third, the leakage model treats observers as Bayesian with a prior over entities. Real observers also react to rumour, to the group's history and to the sector's deal flow. The testable version is an event study: whether globalisation programmes announced by groups that later moved a business were followed by speculation about that entity more often than about others, and whether asymmetric programmes were followed by it more often than symmetric ones.

Fourth, the legal reading is Dutch and European, and the Dutch works council's advice right has no exact equivalent in many jurisdictions. The model's conditions, symmetry and standalone rationale, do not depend on the law; where the law applies, the programme has to be presentable to the council and kept within the insider list, and we have argued it is.

Fifth, the case is constructed. A field study following one group through the timeline, scoring the index at intervals and recording what each step cost and what it signalled, would be worth more than the case and is the natural next step.

Beyond these, the two-sided joint-venture case invites a model of its own: the contribution credit as a bargaining outcome between two parents with their own indices, in which each parent's readiness is both a cost it saves and a bargaining position it holds, and the stability results of Inkpen and Beamish (1997) are restated in terms of what each parent's entanglement lets it withhold.

15. Conclusion

A group that may partner a division, in a joint venture, an acquisition or a merger, or separate one, by sale, spin-off or listing, cannot say so until it must, and the preparation that would make the move clean takes years. The way out is not to prepare in secret, which fails as soon as enough people must act, but to prepare in a way that reveals nothing: the same programme for every entity, each step with a rationale of its own. Under uncertainty about which deal will come, that symmetric programme is also the economically correct one, because the friction of a deal, the separation cost, the counterparty's discount, the cost of transitional services and the cost stranded with the group that remains, rises convexly with entanglement, and lowering every entity's entanglement is worth more than lowering one's to zero. The programme that does it is the globalisation of shared IT, built with the design rules that turn undefined dependence into defined dependence: the entity on everything, a catalogue with unit prices, per-entity accounts and transition clauses in every contract.

We at Nashua have argued that this programme is several things at once. It is readiness, because an entity cleanly separable from its parent is also one a partner or acquirer can cleanly receive. It is confidentiality, because its business case closes with no deal in view and its footprint is the same for every entity. It is the safeguard, because it moves every vital shared capability out of the divisions into a core the remainer keeps, so that no division walks out with the group's competence, into a buyer or, worse, into a venture a competitor co-owns. And it is the standing due-diligence dossier, because the registers it keeps per entity, as the regulations require, are what a partner or acquirer's due diligence asks for and can be opened in days. The legal boundary is kept, not evaded: the insider list stays small because the programme needs few to know, the works council is asked at the moment the Act names and is given a clearer picture than an entangled division could provide, and disclosure happens when it must. What the programme buys is an option: to contribute any entity to a venture, to sell one, or to bolt one on, at a known and low exercise cost, without having told anyone which. The convexity on which the argument rests is the quantity most worth measuring, and the deals that Dutch and other European groups have carried out in public view are where it could be measured. We would rather it were measured than believed.

Appendix A. Measuring the index and readiness

Scoring e_{id}. For each entity and dimension, inventory the dependences on things outside the entity: systems, platforms, contracts, data stores, teams. Classify each as defined if it meets the criteria of Table 1 for its dimension, with the entity marked, the service priced and the hand-over documented, and as undefined otherwise. Weight each dependence by operational importance on a three-point scale (vital, significant, minor) and compute e_{id} as the importance-weighted share that is undefined. Score from the registers where they exist and from interviews where they do not; the interviews are themselves evidence of what is undefined.

Weights w_d. Set by the kind of business: a manufacturing division weights applications and people heavily, a sales division data and contracts, a services division identity and infrastructure. Where there is no reason to differ, equal weights; the case uses them.

Readiness. We define readiness as $1 - E_i$, and symmetric readiness as the condition that the index of every entity lies within a narrow band, which the group can report as the

range of E_i across entities. A wide range is a leak waiting to be read.

The capability register. List every IT capability the group as a whole depends on, the location where it sits (payroll, contract, system), and whether it also sits in the shared core. The set R_j of equation (6) for each entity j is read off the register, and the safeguard is complete when every R_j is empty.

Stranded cost. From the cost model of the catalogue: for each shared service, the fixed and variable fractions and the allocation per entity. σ_j follows from equation (7) with H_j from the ramp-down plan, or three years where none exists.

Convexity β. From a sample of completed carve-outs: score the index retrospectively, take separation cost and TSA duration from the project records, and fit $\log(S - s_0)$ on $\log E$; the slope is β. The same fit on the counterparty's discount, where a competitive process reveals it, gives θ.

Leakage. Record, for each entity, the observable footprint of the programme: projects, staff moves, contract changes, sequencing. Symmetry is the condition that the footprints are the same up to the entity's size. Where a footprint differs, record why, and whether the reason is one that could be stated to the entity's management without revealing an intention.

Appendix B. Confidentiality signal checklist

What leaks, and what does not, read from the leakage condition of Section 4: an observable differs across entities, or has no rationale without a deal.

Leaks. A project, wave or exception scoped to one entity; data segregation sequenced by entity; a supplier asked about assignment or change of control for one entity; an insider list extended into one entity's management; a hiring freeze, bonus change or retention plan in one entity; a legal-entity restructuring that isolates one entity without a tax or regulatory reason; advisers or bankers at one entity's sites; a data room opened before announcement; a draft TSA, contribution agreement or Day-1 plan with a named counterparty; a codenamed programme whose scope excludes some entities; a vendor due-diligence report commissioned for one entity; a clean team set up before announcement; a catalogue priced for one entity only.

Does not leak. A platform programme applied to every entity in waves set by region, platform or lease; the entity attribute on every account, record and resource; records of processing per controller, because the law requires them; a catalogue with unit prices for every entity, because it is cost transparency; affiliate-transition clauses in every global contract as standard terms; competence centres moved to global centres of excellence from every division; certifications with per-entity scope; an audit-readiness exercise covering every entity; a stranded-cost analysis presented as the catalogue's cost model.

References

[1] AICPA & CIMA. *System and Organization Controls: SOC suite of services*. https://www.aicpa-cima.com/resources/landing/system-and-organization-controls-soc-suite-of-services

[2] Akerlof, G. A. (1970). The market for "lemons": Quality uncertainty and the market mechanism. *Quarterly Journal of Economics*, 84(3), 488–500. https://doi.org/10.2307/1879431

[3] AkzoNobel (2018, 1 October). *AkzoNobel closes sale of Specialty Chemicals to The Carlyle Group and GIC*. Media release. https://www.akzonobel.com/en/media/media-releases/akzonobel-closes-sale-of-specialty-chemicals-to-the-carlyle-grou

[4] Baldwin, C. Y., & Clark, K. B. (2000). *Design rules: The power of modularity*. MIT Press. https://doi.org/10.7551/mitpress/2366.001.0001

[5] Brauer, M. (2006). What have we acquired and what should we acquire in divestiture research? A review and research agenda. *Journal of Management*, 32(6), 751–785. https://doi.org/10.1177/0149206306292879

[6] Capron, L. (1999). The long-term performance of horizontal acquisitions. *Strategic Management Journal*, 20(11), 987–1018. https://doi.org/10.1002/(SICI)1097-0266(199911)20:11<987::AID-SMJ61>3.0.CO;2-B

[7] Capron, L., Dussauge, P., & Mitchell, W. (1998). Resource redeployment following horizontal acquisitions in Europe and North America, 1988–1992. *Strategic Management Journal*, 19(7), 631–661. https://doi.org/10.1002/(SICI)1097-0266(199807)19:7<631::AID-SMJ963>3.0.CO;2-9

[8] Capron, L., Mitchell, W., & Swaminathan, A. (2001). Asset divestiture following horizontal acquisitions: A dynamic view. *Strategic Management Journal*, 22(9), 817–844. https://doi.org/10.1002/smj.175

[9] Dixit, A. K., & Pindyck, R. S. (1994). *Investment under uncertainty*. Princeton University Press. https://doi.org/10.1515/9781400830176

[10] Dyer, J. H., & Singh, H. (1998). The relational view: Cooperative strategy and sources of interorganizational competitive advantage. *Academy of Management Review*, 23(4), 660–679. https://doi.org/10.5465/amr.1998.1255632

[11] dsm-firmenich (2023, 9 May). *Completion of contribution of Firmenich shares*. Press release. https://our-company.dsm-firmenich.com/en/our-company/news/press-releases/2023/completion-of-contribution-of-firmenich-shares.html

[12] European Parliament and Council (2014). Regulation (EU) No 596/2014 on market abuse (Market Abuse Regulation). *Official Journal of the European Union*, L 173, 12 June 2014. https://eur-lex.europa.eu/eli/reg/2014/596/oj

[13] European Parliament and Council (2016). Regulation (EU) 2016/679 on the protection of natural persons with regard to the processing of personal data (General Data Protection Regulation). *Official Journal of the European Union*, L 119, 4 May 2016. https://eur-lex.europa.eu/eli/reg/2016/679/oj

[14] European Parliament and Council (2022). Regulation (EU) 2022/2554 on digital operational resilience for the financial sector (DORA). *Official Journal of the European Union*, L

333, 27 December 2022. https://eur-lex.europa.eu/eli/reg/2022/2554/oj

[15] European Parliament and Council (2022). Directive (EU) 2022/2555 on measures for a high common level of cybersecurity across the Union (NIS 2). *Official Journal of the European Union*, L 333, 27 December 2022. https://eur-lex.europa.eu/eli/dir/2022/2555/oj

[16] Fähling, J., Leimeister, J. M., Yetton, P., & Krcmar, H. (2013). Managing an IT carve-out at a multi-national enterprise. *Journal of Information Technology Teaching Cases*, 3(2), 106–110. https://doi.org/10.1057/jittc.2013.10

[17] Feldman, E. R., & McGrath, P. J. (2016). Divestitures. *Journal of Organization Design*, 5(1), 2. https://doi.org/10.1186/s41469-016-0002-x

[18] Graebner, M. E., Heimeriks, K. H., Huy, Q. N., & Vaara, E. (2017). The process of postmerger integration: A review and agenda for future research. *Academy of Management Annals*, 11(1), 1–32. https://doi.org/10.5465/annals.2014.0078

[19] Haspeslagh, P. C., & Jemison, D. B. (1991). *Managing acquisitions: Creating value through corporate renewal*. Free Press. ISBN 978-0-02-914165-6.

[20] Hennart, J.-F. (1988). A transaction costs theory of equity joint ventures. *Strategic Management Journal*, 9(4), 361–374. https://doi.org/10.1002/smj.4250090406

[21] Henningsson, S., & Carlsson, S. (2011). The DySIIM model for managing IS integration in mergers and acquisitions. *Information Systems Journal*, 21(5), 441–476. https://doi.org/10.1111/j.1365-2575.2011.00374.x

[22] Henningsson, S., & Kettinger, W. J. (2016). Understanding information systems integration deficiencies in mergers and acquisitions: A configurational perspective. *Journal of Management Information Systems*, 33(4), 942–977. https://doi.org/10.1080/07421222.2016.1267516

[23] Henningsson, S., Yetton, P. W., & Wynne, P. J. (2018). A review of information system integration in mergers and acquisitions. *Journal of Information Technology*, 33(4), 255–303. https://doi.org/10.1057/s41265-017-0051-9

[24] Inkpen, A. C., & Beamish, P. W. (1997). Knowledge, bargaining power, and the instability of international joint ventures. *Academy of Management Review*, 22(1), 177–202. https://doi.org/10.5465/amr.1997.9707180262

[25] ISO/IEC (2022). *ISO/IEC 27001:2022. Information security, cybersecurity and privacy protection, Information security management systems, Requirements*. International Organization for Standardization. https://www.iso.org/standard/27001

[26] Kogut, B. (1988). Joint ventures: Theoretical and empirical perspectives. *Strategic Management Journal*, 9(4), 319–332. https://doi.org/10.1002/smj.4250090403

[27] Mehta, M., & Hirschheim, R. (2007). Strategic alignment in mergers and acquisitions: Theorizing IS integration decision making. *Journal of the Association for Information Systems*, 8(3), 143–174. https://doi.org/10.17705/1jais.00118

[28] Moschieri, C., & Mair, J. (2008). Research on corporate divestitures: A synthesis. *Journal of Management & Organization*, 14(4), 399–422. https://doi.org/10.5172/jmo.837.14.4.399

[29] Myers, S. C. (1977). Determinants of corporate borrowing. *Journal of Financial Economics*, 5(2), 147–175. https://doi.org/10.1016/0304-405X(77)90015-0

[30] Nouryon (2018, 9 October). *AkzoNobel Specialty Chemicals is now Nouryon*. News release. https://www.nouryon.com/news-and-events/news-overview/2018/AkzoNobel-Specialty-Chemicals-is-now-Nouryon

[31] Parnas, D. L. (1972). On the criteria to be used in decomposing systems into modules. *Communications of the ACM*, 15(12), 1053–1058. https://doi.org/10.1145/361598.361623

[32] Philips (2021, 2 September). *Philips completes sale of Domestic Appliances business to global investment firm Hillhouse Investment*. Press release, Koninklijke Philips N.V. https://www.philips.com/a-w/about/news/archive/standard/news/press/2021/20210902-philips-completes-sale-of-domestic-appliances-business-to-global-investment-firm-hillhouse-investment.html

[33] Philips Lighting (2016, 27 May). *Philips Lighting joins Euronext Amsterdam in this year's largest IPO in Europe to date*. Press information. https://www.signify.com/b-dam/b2b-li/en_AA/investor/ipo/Philips_Lighting_Listing_Press_Release.pdf

[34] Puranam, P., Singh, H., & Zollo, M. (2006). Organizing for innovation: Managing the coordination-autonomy dilemma in technology acquisitions. *Academy of Management Journal*, 49(2), 263–280. https://doi.org/10.5465/amj.2006.20786062

[35] Robbins, S. S., & Stylianou, A. C. (1999). Post-merger systems integration: The impact on IS capabilities. *Information & Management*, 36(4), 205–212. https://doi.org/10.1016/S0378-7206(99)00018-X

[36] Ross, J. W., Weill, P., & Robertson, D. C. (2006). *Enterprise architecture as strategy: Creating a foundation for business execution*. Harvard Business School Press. ISBN 978-1-59139-839-4.

[37] Sanchez, R., & Mahoney, J. T. (1996). Modularity, flexibility, and knowledge management in product and organization design. *Strategic Management Journal*, 17(S2), 63–76. https://doi.org/10.1002/smj.4250171107

[38] Simon, H. A. (1962). The architecture of complexity. *Proceedings of the American Philosophical Society*, 106(6), 467–482.

[39] Spence, M. (1973). Job market signaling. *Quarterly Journal of Economics*, 87(3), 355–374. https://doi.org/10.2307/1882010

[40] Tanriverdi, H., & Uysal, V. B. (2011). Cross-business information technology integration and acquirer value creation in corporate mergers and acquisitions. *Information Systems Research*, 22(4), 703–720. https://doi.org/10.1287/isre.1090.0250

[41] Unilever (2022, 1 July). *Unilever completes the sale of its Tea business, ekaterra*. Press release. https://www.unilever.com/news/press-and-media/press-releases/2022/unilever-announces-completion-of-the-sale-of-its-tea-business-ekaterra/

[42] Wet op de ondernemingsraden (Works Councils Act). Wet van 28 januari 1971, BWBR0002747, artikel 25. https://wetten.overheid.nl/BWBR0002747/

[43] Wijnhoven, F., Spil, T., Stegwee, R., & Fa, R. T. A. (2006). Post-merger IT integration strategies: An IT alignment perspective. *Journal of Strategic Information Systems*, 15(1), 5–28. https://doi.org/10.1016/j.jsis.2005.07.002

[44] Williamson, O. E. (1979). Transaction-cost economics: The governance of contractual relations. *Journal of Law and Economics*, 22(2), 233–261. https://doi.org/10.1086/466942

[45] Yetton, P., Henningsson, S., Böhm, M., Leimeister, J. M., & Krcmar, H. (2022). How IT carve-out project complexity influences divestor performance in M&As. *European Journal of Information Systems*, 32(6), 962–988. https://doi.org/10.1080/0960085X.2022.2085201