

Loss at the handover: where the boundary of a full mandate is drawn

换手处的损耗：全托管的边界画在哪里

庚辛研究院 / Glacier Institute ^{*†}

From The Amundsen Institute

Glacier Institute (庚辛研究院) | Working Paper No. GI-WP-2026-P12-EN | Date: 2026-09-22 | Language: English, with Chinese title and abstract

Cite as:

Glacier Institute (2026). Loss at the handover: where the boundary of a full mandate is drawn. Glacier Institute Working Paper GI-WP-2026-P12-EN. <https://glacier.mba/research/GI-WP-2026-P12-EN.html>

庚辛研究院 (2026). 《换手处的损耗：全托管的边界画在哪里》. 庚辛研究院通讯论文 GI-WP-2026-P12. <https://glacier.mba/research/GI-WP-2026-P12.html>

Key Takeaways

- 1 **"Never handed off" is not a promise. It is a boundary line, and you have to count the steps on it first.**
- 2 **Founders are not asking what you do. They are asking why this particular piece should be yours.**
- 3 **The site carries two lists of six. One is process, one is relationship, and handing them off does not cost the same thing.**
- 4 **Not handing off does not save you the effort of passing documents. It saves you the effort of rebuilding context.**

Abstract

When a financing advisor says it is "responsible end to end, with six dimensions never handed off", that sounds like a promise. It is in fact a boundary question: how far does the mandate go before it stops? This paper treats it as a question about the boundary of the firm rather than as a service slogan.

The source text is public. A boutique investment bank states its "full mandate" in five separate places on its website. All five are reproduced verbatim here, with line numbers in the public plain text. Normalising the items each of them

enumerates yields something that no single sentence says: the site carries two different lists of six. One is a ladder of relationship levels — trust, information, narrative, tempo, relationships, accountability. The other is a sequence of process steps — deal information flow, core value narrative, matching of target capital, deal tempo, coordination of key relationships, closing. The two lists overlap on only four items. The longest public list has nine items, three more than six.

The count is not a word game. It shows that what is "never handed off" is at least two kinds of thing: some of it is process, some of it is relationship. Process can be handed over and the cost of handing it over can be estimated. A relationship cannot: change the person and it starts again. The two require different boundary tests, and while they sit in one list no boundary can be drawn at all.

The mechanism section draws on six bodies of work: transaction cost economics and the boundary of the firm (asset specificity, uncertainty, frequency); incomplete contracts and residual rights of control; multitask agency theory, in which tasks of unequal measurability distort incentives when bundled; the empirical literature on vertical integration; evidence on losses at the handoff — information lost at shift handover, relational coordination, the stickiness of knowledge; and the mirroring hypothesis, under which the structure of what is delivered comes to resemble the structure of the organisation delivering it.

*** Disclosure:** Glacier Institute is the research arm of Glacier Capital, and this paper is issued under the name of Glacier Institute. In the course of its business, Glacier Capital acts as financial adviser to a number of technology companies and invests its own capital in some of them; such relationships may overlap with the industries discussed here. This paper does not concern any specific mandate and uses no non-public information; all company, industry and data references are drawn from public sources and cited individually. The authors received no third-party compensation for this paper.

† Disclaimer: This is a methodological working paper and represents the authors' analysis at the time of writing only. It does not constitute investment advice, nor an offer or solicitation of an offer for any security, fund interest or other instrument, nor a commitment or forecast regarding the valuation, financing outcome or investment return of any company. It has not been peer reviewed and may be revised in later versions.

The paper offers one tool: an eleven-row handoff boundary audit, each row written as criterion, what argues against outsourcing, what argues for it, and its basis. The eight normalised items are then scored against it. The central claim is stated as a falsifiable proposition together with the evidence that would refute it. Five boundary cases close the paper, including one that must be stated: the price of end-to-end is that attribution becomes hard, so not handing off requires an outside assessor.

No client company is named, no description that could identify one is used, and none of the firm's own financial figures appear.

Keywords: boundary of the firm; transaction costs; vertical integration; residual control rights; handoff loss; mirroring hypothesis; financing advisory

JEL Classification: D23 (组织行为与交易成本), L14 (交易关系与网络), L84 (专业与商业服务), G24 (投资银行与创业投资)

摘要 (Chinese abstract)

一家融资顾问机构说自己「端到端负责、六个维度不换手」，这句话听起来是承诺，其实是一个边界问题：管到哪一步为止。本文把它当成企业边界问题来处理。底本是公开网页：一家精品投行在官网上有五处关于「全托管」的公开表述，本文逐字抄录并标出行号；把五处的枚举项逐项归一之后，能数出一件不在任何一句话里的事——官网有两张各含六项的清单，一张是关系层级，一张是流程环节，两张只在四项上重合，而最长的那张公开清单有九项。这说明「不换手」的对象至少有两类：工序可以交接，交接成本可以估；关系不能交接，换人就重来。机制部分用六组文献：交易成本经济学与企业边界、不完全契约与剩余控制权、多任务代理、纵向一体化的经验研究、交接损耗的实证、镜像假说。工具是一张十一行的换手边界自查表，并用它扫了那八个归一项。主张写成可以被推翻的命题，文末给出五类边界。本文不出现任何客户公司的名字，不出现这家机构自身的任何经营数字。

Keywords (Chinese): 企业边界；交易成本；纵向一体化；剩余控制权；交接损耗；镜像假说；融资顾问

1. The Problem: "How Far Does This Go?"

This is the question founders ask most often at a first meeting, and the hardest one to answer well.

Answer narrowly and you sound like an introducer. Answer broadly and the next question arrives immediately: then what is left for me to do? If you do everything, who is accountable when something goes wrong?

Peers and investors are asking the other half of the same thing. When a firm says it works "end to end" and "never

hands off", does that sentence have a boundary anyone can test, or is it just a pleasant phrase?

This paper answers both.

The first question cannot be answered with a service list. A service list says what we do. It does not say why this item belongs to us and that one does not — and the second is what the founder is actually asking. He is deciding whether handing something over is worth it, not learning your scope of business.

The second cannot be answered with "we handle all of it" either. Economics has had a framework for the boundary of the firm for almost ninety years, and the answer it gives is never "more is better". It is "stop at the point where the two costs are equal". Applied to financing advisory, this paper argues that the boundary belongs at **the point where the cost of the interface exceeds the cost of coordinating internally**. One step past that, not handing off starts to lose money.

The route: Section 2 lays out the public source text and counts two lists out of its own enumerations. Section 3 gives six bodies of mechanism literature. Section 4 returns to what "end to end" says in its own field, and to its acknowledged price. Section 5 gives the tool and applies it to the eight items. Section 6 states the claim as a proposition that can be refuted. Sections 7 and 8 draw the boundaries.

How this differs from the other papers in the series, stated here so no reader takes it for repetition.

GI-WP-2026-P8, *One-Way Doors*, is about whether an action can be undone and who presses the button: that is the **action layer**. This paper is about whether a piece of work should sit inside one system at all: the **boundary layer**. In one transaction, the boundary decides who does it; the one-way door decides who presses it.

GI-WP-2026-P11, *Borrowed Theories*, §4.3, gave a three-paragraph verdict on the "end-to-end → full mandate" pair. This paper expands that subsection into a standalone boundary analysis and adds the two things it did not have: an item-by-item normalisation of the public lists, and an audit anyone can run.

Founders are not asking what you do. They are asking why this particular piece should be yours.

2. The Source Text: Five Public Statements of "Never Handed Off"

2.1 How the source text was taken

The source is a public web page and anyone can reproduce it. Method: `curl` the page

`https://glacier.mba/institute.html`; strip HTML comments, then `<script>`, then `<style>`; replace every remaining tag with a newline; unescape entities; drop blank lines. The result is 10,079 lines of plain text, retrieved 2026-09-22. All line numbers below refer to that plain text — the same source text and the same line numbering used by GI-WP-2026-P11.

2.2 First place: Chapter 18, pair 3 (L3880–L3885)

Chapter 18 maps the concept of "end to end" onto the firm's "full mandate":

「从输入到输出，一个模型端到端负责；分成多段的流水线，误差会在段与段之间累积。」

From input to output, one model is responsible end to end. In a pipeline cut into stages, error accumulates between the stages. — theory sentence, L3881 [1]

「从恢复事实到交割收口，同一个系统端到端负责，六个维度不换手——误差没有地方累积。」

From restoring the facts to closing the deal, one and the same system is responsible end to end. Six dimensions are never handed off — error has nowhere to accumulate. — phenomenon sentence, L3883 [1]

The pair carries a source line: "Bojarski et al. (2016), End to End Learning for Self-Driving Cars, arXiv:1604.07316, Figure 4" (L3885) [1][18].

There is a number in that sentence: six. The sentence does not say which six. To find out, you have to go elsewhere on the site.

2.3 The other four places

Second place (L674–L678), the full-mandate collaboration model.

「该模式涵盖六大核心维度（信任为先）：信任、信息、叙事、节奏、关系、担当。」

The model covers six core dimensions (trust first): trust, information, narrative, tempo, relationships, accountability. — L675, repeated verbatim at L678 [1]

「庚辛资本的『全托管协作』意味着：企业保留所有重大事项的最终裁决权，同时将交易信息流、核心价值叙事、目标资方匹配、交易推进节奏、关键关系协调及交

割落地等全套流程，交由庚辛资本统筹主导。」

"Full-mandate collaboration" means the company retains final say over all material matters, while the whole process — deal information flow, core value narrative, matching of target capital, deal tempo, coordination of key relationships and closing — is led and coordinated by the firm. — L676 [1]

Note that these two sentences sit on the same screen and do not enumerate the same thing. L675 enumerates six *dimensions*. L676 enumerates six *process steps*.

Third place (L3220–L3230), home page, section 09.

「企业决策，庚辛统筹。企业家只做三件事：定战略、找对人、深入一线。全托管的本质是同侧——你的利益，就是我的利益。」

The company decides; we coordinate. The founder does three things: set the strategy, find the right people, stay close to the front line. A full mandate means being on the same side — your interest is our interest. — L3220 [1]

「全托管六级：信任是地基，往上依次是信息、叙事、节奏、关系、担当；企业保留最终决策权。」

Six levels of full mandate: trust is the foundation, then information, narrative, tempo, relationships, accountability. The company retains the final decision. — L3230 [1]

Fourth place (L4524–L4525), frequently asked questions.

「全托管是庚辛的深度交易协作：企业保留重大事项最终决策权，交易信息、核心叙事、推进节奏、关键关系与交割安排交由庚辛统筹。」

A full mandate is deep transactional collaboration: the company keeps final say on material matters, while deal information, the core narrative, tempo, key relationships and closing arrangements are coordinated by the firm. — L4525 [1]

This sentence enumerates **five** items, not six.

Fifth place (L4523), the preceding question in the same set.

「不是。投资人引荐只是工作的一部分，核心服务还包括事实恢复、资本路线设计、投资人结构、路演与投资人沟通、复杂交易设计、产业协同、尽调支持、条款协调和交割管理。」

表 1 / Exhibit 1 The enumerated items of the five public statements of "full mandate", normalised. The four columns come from L675/L678/L3230 (six core dimensions), L676 (whole process), L4525 (FAQ) and L4523 (core service). Rows are ordered by the line number of first appearance. Normalisation rule in the writing note §6.1. The table is counted here; the website contains no such comparison.

Normalised item	Six core dimensions	Whole process	FAQ	Core service
Trust	trust	—	—	—
Information	information	deal information flow	deal information	—
Narrative	narrative	core value narrative	core narrative	capital route design
Tempo	tempo	deal tempo	tempo	time cadence (see L4521)
Relationships	relationships	coordination of key relationships	key relationships	industrial coordination
Accountability	accountability	—	—	term negotiation
Capital matching	—	matching of target capital	—	investor structure
Closing	—	closing	closing arrangements	closing management
Restoring the facts	—	—	—	restoring the facts
Complex deal design	—	—	—	complex deal design
Diligence support	—	—	—	diligence support

表 2 / Exhibit 2 Counts from Table 1 (continued from Table 1). The denominator is the eleven normalised items of Table 1. The bands are mutually exclusive and sum to 11, with no double counting. The count is made here.

Band	Items	Which ones
Appears in three or more of the four columns	4	information, narrative, tempo, relationships
Appears only under "six core dimensions"	2	trust, accountability
Appears in a process list but not among the dimensions	2	capital matching, closing
Appears only in the longest list	3	restoring the facts, complex deal design, diligence support

No. Introducing investors is only part of the work. The core service also includes restoring the facts, designing the capital route, investor structure, roadshow and investor communication, complex deal design, industrial coordination, diligence support, term negotiation and closing management. — L4523 [1]

This sentence enumerates **nine** items.

2.4 Normalising the five enumerations

The table below is counted here from the five public sentences above. The normalisation rule is stated in the writing note, § 6.1: items are merged only on their surface wording where they plainly refer to the same thing (「交易信息流」, 「交易信息」, 「信息」 all become "information"); no semantic inference is made; a blank means the item does not appear in that place.

Three things in Tables 1 and 2 are worth stopping on.

First, there are two lists of six. One is a ladder of relationship levels: trust, information, narrative, tempo, relationships, accountability. The other is a sequence of process steps: deal information flow, core value narrative, matching of target capital, deal tempo, coordination of key relationships, closing. The Chapter 18 sentence uses the word 「维度」, *dimensions*, which literally points to the first. But a reader will easily take it for the second, because "handed off" is a verb that only makes sense about process.

Second, the two lists overlap on only four items: information, narrative, tempo, relationships. What these four have in common is that each is both a step and a relation. Information has to be gathered (a step), and someone has to trust you enough to give it to you (a relation).

Third, what is unique to each list is of a completely different nature. Unique to the dimensions are trust and accountability — these are not

steps but properties; they cannot be handed over, only built or lost. Unique to the process list are capital matching and closing — these are ordinary steps, with a start, an end and a deliverable.

This is where the paper starts. The object of "never handed off" mixes two kinds of thing. Handing off a step produces a loss, and the loss can be estimated and compared. Handing off a relationship is not a loss but a reset to zero. If both are explained by the same reason — that error does not accumulate at the seams — then that reason is right about one of them and unnecessary for the other. A relationship never had a seam-error problem. It has a different one: **change the person and it starts again.**

The site carries two lists of six. One is process, one is relationship, and handing them off does not cost the same thing.

3. Mechanism: A Boundary Is Compared, Not Declared

3.1 The boundary sits where two costs are equal

Coase [2] asked why firms exist at all if markets allocate resources. His answer was that using the price mechanism has a cost — finding a counterparty, negotiating, contracting, monitoring performance all cost money.

Firms exist because for some things, organising by command once is cheaper than buying by price once. The boundary test that follows is a single sentence: **bring work inside until the cost of organising one more piece internally equals the cost of buying it on the market.**

Coase did not say what determines that difference. Williamson [3] supplied three workable dimensions:

- **Asset specificity.** How much of the investment made for this transaction survives a change of counterparty. The less survives, the higher the specificity.
- **Uncertainty.** Whether what may happen later can be written down at signing. The more that cannot, the more discretion is needed afterwards.
- **Frequency.** How often transactions of this kind occur. High frequency is what amortises a dedicated internal arrangement.

The more of the three run high, the stronger the case for internalising; all three low and market exchange is cheaper. Williamson later restated this as a comparison of discrete structures [4]: market, hybrid (long-term contract, alliance) and hierarchy are three governance forms each with its own range of fit, not a ranking of better and worse. **This matters here: a full mandate is a governance structure, not a service tier.**

3.2 Residual control: who decides what the contract left out

Grossman and Hart [5] and Hart and Moore [6] come in from another direction. Their starting point is that no contract is complete; something is always left out. So what really has to be allocated is not "who does what" but **who decides in the situations nobody wrote down** — the residual right of control. Whoever holds it has the incentive to make specific investments up front, because he knows he will not be held up afterwards.

This has a direct consequence here. One sentence recurs almost unchanged in three of the five public statements:

「企业保留所有重大事项的最终裁决权」 *the company retains final say over all material matters* — L676 [1]

「企业保留最终决策权」 *the company retains the final decision* — L3230 [1]

「企业保留重大事项最终决策权」 *the company keeps final say on material matters* — L4525 [1]

In the Grossman–Hart frame this is not a courtesy. **It places the residual right of control firmly on the client's side.** The consequence can be derived: the advisor therefore does not hold ex post discretion, so whatever specific investment the advisor makes has to be protected some other way — by repetition (frequency), by reputation, or by the part that is written down. That is where row 11 of the audit in Section 5 comes from.

3.3 Multitask: bundling the measurable with the unmeasurable goes wrong

Holmström and Milgrom [7] study an agent doing several things at once. The conclusion is counter-intuitive but robust: when some tasks have easily measured output and others do not, strengthening the incentive on the measurable ones **draws effort away** from the rest. The right response is often not to tune the incentive but to **redraw who owns which task** — put tasks of very unequal measurability with different people, or give none of them strong incentives.

Applied here: the eleven items of Table 1 differ enormously in measurability. Closing has a clear completion test. So does restoring the facts — do the materials reconcile or not. Trust and accountability have almost no measurable current output. Putting them in one mandate and then measuring only the measurable few will systematically pull effort away from the rest. **This yields a boundary test: if a step's output cannot be signed off on its own, do not carve it out and give it to someone else — once carved out, nobody can prove it was ever done well.**

3.4 Empirical work: when integration actually happens

There are decades of evidence behind the theory. Monteverde and Teece [8] studied automotive components: the more specialised the engineering knowledge required and the costlier it is to switch supplier, the more likely the part is made in-house. Masten [9] studied make-or-buy decisions in aerospace and found the same direction — specificity and complexity raise the probability of internalisation. Lafontaine and Slade [10] reviewed thirty years of empirical work on vertical integration and report that this literature is highly consistent in direction: **the higher the asset specificity and the uncertainty, the more likely integration**, in line with the transaction cost prediction.

The use here is simple. The boundary test is not invented by this paper, nor by the firm whose text it reads. It has been tested repeatedly in other industries. There is no reason financing advisory runs on different physics.

3.5 How real the loss at the handoff is

"Error accumulates at the seams" sounds like a metaphor. It is not.

Arora and colleagues [12] carried out a critical incident analysis of communication failure when patients are handed between shifts in a hospital. The most common failure type they recorded was **omission** — information the outgoing shift held but did not pass on, and information passed on but not acted on. The value of the study is that it turns "handoff loss" from a saying into something that can be counted.

Gittell [11] gives evidence from the other side: across service settings requiring several occupational groups to work together, **relational coordination** — shared goals, shared knowledge and mutual respect, together with frequent, timely, problem-solving communication —

is associated with performance, and matters more where input uncertainty is high. In other words, the loss at the seam is not removed by more documentation. It is removed by relationship — which is exactly the thing that cannot be handed off.

Szulanski [13] adds a third finding: even inside one company, moving a practice from one place to another is hard. He calls it **internal stickiness**, and finds that the main impediment is not weak motivation but the recipient's lack of absorptive preparation and an arduous relationship between the two sides. Argote and Ingram [14] generalise the point: knowledge resides in the combination of people, tools and tasks, and moving only one of the three does not move the knowledge.

Together these give a restatement of what "never handed off" actually saves: **not the cost of passing documents, but the cost of rebuilding context.**

3.6 Mirroring: what you deliver grows into the shape of your organisation

Conway [15] wrote in 1968 a sentence that has been quoted ever since: the structure of a system copies the communication structure of the organisation that designs it. MacCormack, Baldwin and Rusnak [16] tested this directly on software, comparing products of similar function built by tightly coupled organisations and by loosely coupled distributed ones, and found the former significantly less modular — **organisational structure and product structure really are mirrored**. Colfer and Baldwin [17] aggregated more than a hundred studies and confirm that the hypothesis holds in most settings, while identifying the conditions for exceptions: where an organisation can establish high-bandwidth communication across a boundary, mirroring can be broken.

Two uses here. First, a deliverable produced in segments by several organisations will most likely grow the corresponding seams — that is structural, not a failure of execution. Second, and more important: **mirroring can be broken, at the cost of building a high-bandwidth channel at the boundary**. So the question is never "should we hand off" but "is the channel at that handoff worth building". If it is, hand off. If it cannot be afforded, do not.

Not handing off does not save you the effort of passing documents. It saves you the effort of rebuilding context.

表 3 / Exhibit 3 The handoff boundary audit. Eleven criteria to pass before deciding whether a step belongs inside one system. The basis is in the last column. The table is this paper's own; the website contains no equivalent.

No.	Criterion	Argues against handing off	Argues for handing off	Basis
1	Asset specificity	What this step produces is near worthless with another counterparty	The output is equally useful to anyone	[3][8][9]
2	Uncertainty	What will come up cannot be written down in advance	Cases are enumerable and fit on a deliverables list	[3][4][10]
3	Frequency	This kind of thing happens many times a year	Rarely; a dedicated arrangement cannot be amortised	[3]
4	Standalone sign-off	The output cannot be judged complete on its own	There is a clear completion test a third party can apply	[7]
5	Coupling to neighbours	The handoff carries a large, hard-to-state context	The handoff list fits on one page	[13][14]
6	Recipient's preparation	The recipient has no prior knowledge; it has to be taught	The recipient already knows this class of work	[13]
7	Verifiability of the handoff	The recipient cannot check on the spot whether what arrived is right	What is passed over can be independently reconciled	[12]
8	Attribution need	A failure must be locatable to a specific step	Redoing the whole thing is not expensive	[18]
9	Relational content	The value lives between one person and another	The value lives in documents and process	[11][14]
10	Structural mirroring	A handoff would leave a matching seam in the deliverable	A high-bandwidth channel can be built at the boundary	[15][16][17]
11	Residual control	Situations not in the contract must be decided by the client on the spot	Unspecified matters can be handled by existing rules	[5][6]

4. What End-to-End Says in Its Own Field

4.1 Where it wins

The source the website cites [18] trains a neural network to map camera pixels directly to steering angle, dropping the hand-designed intermediate stages of lane detection and path planning. Its argument is that each stage of a segmented pipeline optimises its own objective, the objectives across stages are not aligned, and each stage being locally optimal does not make the whole optimal — while error accumulates at the seams.

The argument needs two conditions: **enough end-to-end training signal, and error genuinely arising at the seams rather than inside individual stages.**

The second usually holds in cross-organisation work; the four empirical findings of §3.5 are the evidence. The first is weak in a transaction: the end-to-end learning signal is the closing outcome, which is sparse and slow.

4.2 Where it loses

End-to-end carries a price acknowledged in its own field: **it is hard to interpret and hard to localise.** The advantage of a segmented pipeline is precisely that a failure can be attributed to a stage — lane detection is wrong, so go and fix lane detection. When an end-to-end system fails, you retrain; saying which layer learned the wrong thing is hard.

That price has an exact organisational counterpart. A firm that does a piece of work from beginning to end by itself loses the cross-check that a handoff produces for free. Handoffs cost something, but a handoff is also a free review — whoever picks up the next leg starts by asking how this came about. Remove the handoff and you have to buy that review back.

So the first conclusion of this paper is a pairing requirement: end-to-end must be paired with an outside assessor. This is not a cautious footnote; it is the condition under which end-to-end holds. It is the same matter, seen from the other side, as

表 4 / Exhibit 4 The eight normalised items of Table 1 scored against the eleven criteria of Table 3 (the three that appear only in the longest list are discussed in the text). "Keep" = most criteria argue against handing off; "May" = most argue for; "Split" = the process part and the relational part of the same item point opposite ways. The scoring is this paper's judgement, based on Sections 3 and 4; the website contains no such evaluation.

Normalised item	Verdict	Deciding criteria	Reason in one line
Trust	Keep	9, 11	Not a step; change the person and it restarts
Information	Keep	5, 6, 7	The context to be passed far exceeds the documents
Narrative	Keep	1, 5, 10	The most specific item; near worthless with another counterparty
Tempo	Keep	2, 8	Cannot be written down in advance; failures must be locatable
Relationships	Keep	9, 1	The value lives between people
Accountability	Keep	4, 9, 11	Cannot be signed off on its own; only ownership holds it
Capital matching	Split	3, 4, 6	The list can be divided; the judgement cannot
Closing	Split	4, 7, 11	The process can be outsourced; the discretion must be present

GI-WP-2026-P13 in this series — that paper is about what the assessor has to produce.

4.3 Back to the sentence on the site

Return to the phenomenon sentence: from restoring the facts to closing the deal, one and the same system is responsible end to end, six dimensions are never handed off, error has nowhere to accumulate.

By the count of §2.4, the scope of that sentence has to be stated in parts:

- For the **process** items (capital matching, closing, restoring the facts, complex deal design, diligence support), the end-to-end argument holds, on the second condition of § 4.1 — error arises mainly at the seams. The evidence of § 3.5 supports it.
- For the **relationship** items (trust, accountability), the end-to-end argument is **unnecessary**. The reason they are not handed off is not error accumulation but that a relationship cannot be transferred. Supporting a correct practice with the wrong reason leaves the practice undefended when it is questioned.
- For the **amphibious** four (information, narrative, tempo, relationships), both reasons hold and reinforce each other.

记忆点 · KEY POINT

The price of end-to-end is that you cannot tell which stage failed. So not handing off means buying one review back.

5. The Tool: A Handoff Boundary Audit

5.1 Eleven criteria

The six bodies of literature in Section 3, compressed into one table, for whoever has to decide tomorrow whether a given step should be handed over.

How to use it: go down the rows. The more that land in the middle column, the more that step belongs inside one system. The more that land in the right column, the more it should be handed over — **handing it over is not a failure, it is a saving**. The table gives no total, because different steps do not carry equal weight. What it gives is a list of questions that all have to be answered.

Row 11 comes last because it is the hardest. The sentence repeated three times on the site — the company retains the final decision — is not a statement of posture in this table but a **structural constraint**: it permanently excludes a class of matters from "can be handed over", not because the advisor is incapable, but because the residual right of control has already been allocated to the other side.

5.2 Running the audit over the eight items

"Split" deserves a sentence. Capital matching and closing are the two most process-like items in Table 1, and for that reason the two most easily misjudged. Their process parts — assembling a list, preparing documents, chasing signatures — put criteria 3, 4 and 7 squarely in the right column: frequent, separately verifiable, independently reconcilable. But each wraps a core that cannot be outsourced: for the first, the judgement of *why this one fits*; for the second, the discretion of *whether to concede this term*. **Keep the whole item and you keep the savings you could have had. Hand over the whole item and you hand over the core.** The right move is to cut once more inside the item.

Where exactly to cut, this paper cannot say in general; it depends on the arrangement. What it can say is the plane of the cut: **between what can be signed off separately and what can only be decided in the room.**

记忆点 · KEY POINT

The boundary is not drawn between steps. It is drawn between what can be signed off and what can only be decided in the room.

6. A Proposition That Can Be Refuted

Main proposition. If handing off a step requires transferring a large amount of context that the recipient cannot independently verify — that is, criteria 5, 6 and 7 all fall in the left column — then giving that step to another organisation produces an observable loss, which shows up as more rework in that step and downstream of it than when it is not handed off. Conversely, if the handoff list for a step is short and its output can be independently signed off, then keeping it inside merely adds coordination cost without a corresponding reduction in error.

What evidence would refute it. Take a set of comparable transactions, pair them step by step into a handed-off group and a not-handed-off group, and count rework in that step and downstream. If the two groups do not differ, the proposition fails — the verifiability of context would then bear no relation to handoff loss, and criteria 5, 6 and 7 would not be criteria.

The data this requires is not something the Glacier Institute can supply on its own: it needs cross-organisation, step-paired process records. This paper states plainly that it cannot run the test, and does not substitute internal data for it.

Corollary one (list test, run once here). If what is "never handed off" really mixes two kinds of thing, then a firm's public statements should contain **more than one list**, and the lists should not fully coincide. Tables 1 and 2 show that normalising five public statements gives eleven items, of which only four appear in three or more columns, while the two lists of six overlap on only four.

Corollary one is not refuted on this one sample. Take another firm and find its public statements carry a single, consistently repeated list, and the corollary is refuted — the split would then reflect inconsistent wording rather than two kinds of object.

Corollary two (mirroring test, not run). If the mirroring hypothesis [15][16][17] holds for financing transactions, then a transaction completed in segments by several organisations should leave observable inconsistencies at the corresponding seams: two versions of the same fact in different documents. The test would take a set of publicly disclosed transaction documents, count how often the same fact is stated differently, and group by how many organisations handled it in segments. This paper does not run it, because public information about segmentation is insufficient. **Left explicitly open; not filled with conjecture.**

Corollary three (pairing test). §4.2 argues that end-to-end must be paired with an outside assessor. Testable form: in organisations that work end to end without an outside assessor, the time to locate a problem should be longer. Observe that they locate it just as fast, and the pairing claim fails.

记忆点 · KEY POINT

If you cannot say what would show it false, it is not a boundary. It is a sentence.

7. Boundaries and Counter-Cases

An analysis with only supporting cases does not deserve trust. Five places where this boundary test is known to fail.

Counter-case one: the price of end-to-end is hidden by its own success. This is the important one. An organisation that does not hand off may genuinely fail less often — but when it does fail, it locates the failure more slowly, and **slow location leaves no record**. An organisation can only see what it writes down, so it will systematically under-count the cost of attribution. There is one remedy: **manufacture a handoff where there is none** — have someone who was not involved review the work at a fixed point. That handoff is not there for division of labour. It is there to buy back the free check.

Counter-case two: reading "should not be handed off" as "should not be checked". Six items in Table 4 are marked Keep. Staying inside one system does not mean staying in one person's hands, still less going unchecked. Among Williamson's three governance forms [4], the price of hierarchy is precisely that internal monitoring cost rises — internalising saves market transaction cost and buys management cost in exchange. **Treating internalisation as free is the most common misuse of this framework.**

Counter-case three: criterion 9 can be used to exempt anything. Any piece of work can be described as having "a relational element", and thereby escape being split, signed off or compared. So the criterion needs a harder form: not "is there a relational element" but "**how much value is lost if someone else does it**". If you cannot state a specific quantity, you cannot claim exemption under criterion 9.

Counter-case four: mirroring has exceptions, and the exceptions are what you can win. Colfer and Baldwin [17] state explicitly that mirroring breaks where an organisation builds high-bandwidth communication at the boundary. So "a handoff always leaves a seam" is a **default, not a law**. A firm that refuses all handoffs because it believes the law is giving up the division of labour it could have had by building a channel. **Defaults can be changed, at the cost of paying for the channel.**

Counter-case five: the list here comes from one firm's public statements, and the denominator is unknown. Table 1 counts what the site writes. Work that is done but not written up is not in the denominator; work that is written up but actually performed by someone else cannot be detected here either. The eleven items are therefore a **sample of public statements**, not the firm's actual scope of work. This is the same point as §7,

counter-case three, of GI-WP-2026-P11: **what is not written down does not enter the denominator, and then the denominator is not real.**

Internalising saves transaction cost and buys management cost. It was never free.

8. Scope

One, the object. This paper analyses five statements about service boundaries in one public text from one professional service firm. Its conclusions apply to boundary questions in intermediary services characterised by knowledge work, highly specific deliverables and long feedback cycles. They do not extend to make-or-buy decisions in manufacturing, nor to services with short feedback and easily measured output.

Two, the nature of the evidence. Table 1 is a normalised count of verbatim extracts; the normalisation rule is in the writing note and the line numbers are reproducible. Table 2 is counted from Table 1. **Table 3 is a tool constructed here, and Table 4 is a judgement, not a measurement.** The eight rows of Table 4 would need internal information to be decided accurately; judged only from public text, they are conservative — leaning towards Keep or Split. Another judge would differ on individual rows. The three readings offered here depend on the overall shape (two lists, four overlapping, two unique on each side), not on any single cell.

Three, a single sample. One firm, one text. The main proposition of Section 6 needs cross-organisation, step-paired process data, which this paper cannot obtain. Corollary one is not refuted on this one sample, which is not the same as being confirmed.

Four, no causal inference. This paper does not claim that a full-mandate arrangement produces any outcome. Section 5 is about where the boundary should be drawn, not about what drawing it there yields. Reading Table 4 as a performance evaluation is a misreading.

Five, the summarised theory. The account of six literatures in Section 3 is simplified to serve a boundary test and does not substitute for the originals. There is a long-running debate inside transaction cost economics over whether asset specificity is the principal mechanism; this paper does not enter it, and takes only the three dimensions repeatedly supported by empirical work.

Six, an honest gap. What this paper does not do is corollary two, the mirroring test, which needs a set of publicly disclosed transaction documents whose segmentation can be determined. Until that work is done, the judgement that a handoff always leaves a seam is a working assumption that can be refuted, not a conclusion.

Seven, what this paper does not do. It does not evaluate any business outcome of the firm, does not constitute investment advice, and gives no information traceable to any specific transaction or company.

Table 4 is a judgement, not a measurement. What refutes it is paired process records, not better wording.

9. Conclusion

Back to the two questions of Section 1.

"How far does this go?" There is an answer, and it is not a service list. The answer is a test: **as far as the point where the cost of the interface exceeds the cost of coordinating internally.** To place a given step on one side or the other, Section 5 gives eleven rows. The three most useful are rows 5, 6 and 7 — how much unstatable context the handoff carries, whether the recipient is prepared to absorb it, and whether what arrives can be checked on the spot. All three in the left column and the step should not be handed off. All three in the right and keeping it merely adds management cost.

Does "end to end, never handed off" have a testable boundary? It does, and the boundary is not in the firm's own wording but in the distinction between two kinds of object. What the count shows: the public statements carry two lists of six that overlap on only four items; what is unique to each is either a property (trust, accountability) or a step (capital matching, closing). **Handing off a step is a loss, and losses can be compared. Handing off a relationship is a reset to zero, and that cannot be compared.** One reason cannot carry both; for one of them it is surplus.

This paper makes one recommendation, and only one: **where nothing is handed off, manufacture a handoff.** §4.2 shows that the price of end-to-end is that attribution gets harder; §7, counter-case one, shows that this price leaves no record and is therefore under-counted. The remedy is not to abandon end-to-end but to buy back the cross-check a handoff would have given for free — have someone who was not involved

review the work, at a fixed point. How to do that is the subject of the next paper in this series.

Three sentences summarise the paper:

1. "Never handed off" is not a promise but a boundary line, and you have to count the steps on it — and ask whether they are the same kind of thing.
2. What not handing off really saves is not the effort of passing documents but the effort of rebuilding context.
3. Internalising saves transaction cost and buys management cost. It was never free.

Where nothing is handed off, manufacture a handoff — and buy back the review the handoff would have given you for free.

References

Format: author (year). Title. Journal/publisher, volume(issue), pages. DOI or accessible link. Grouped by class of source: primary material is the firm's own public text; secondary material is peer-reviewed literature and public preprints. Each entry was checked by live request on 2026-09-22 using the `check()` function of `_仓储/_build/refcheck.py`; all 18 are reachable, and the title of each was compared word for word against the title returned by Crossref or the arXiv API. Details are in the writing note, § 4. No unverified entry is listed.

Primary material (the firm's public text; every quoted sentence comes from [1], unchanged)

[1] Glacier Institute (2026). Glacier Institute website: the full-mandate collaboration model (L674–L678), home page section 09, FULL-MANDATE ADVISORY (L3215–L3230), Chapter 18 pair 3, "end to end / full mandate" (L3880–L3885), and frequently asked questions (L4520–L4525).
<https://glacier.mba/institute.html> (accessed 2026-09-22)

Secondary (boundary of the firm and transaction costs)

[2] Coase, R. H. (1937). The Nature of the Firm. *Economica*, 4(16), 386–405.
<https://doi.org/10.1111/j.1468-0335.1937.tb00002.x>

[3] Williamson, O. E. (1979). Transaction-Cost Economics: The Governance of Contractual Relations. *The Journal of Law and Economics*, 22(2), 233–261.
<https://doi.org/10.1086/466942>

[4] Williamson, O. E. (1991). Comparative Economic Organization: The Analysis of Discrete Structural Alternatives. *Administrative Science Quarterly*, 36(2), 269–296. <https://doi.org/10.2307/2393356>

Secondary (incomplete contracts and incentives)

[5] Grossman, S. J., & Hart, O. D. (1986). The Costs and Benefits of Ownership: A Theory of Vertical and Lateral Integration. *Journal of Political Economy*, 94(4), 691–719. <https://doi.org/10.1086/261404>

[6] Hart, O., & Moore, J. (1990). Property Rights and the Nature of the Firm. *Journal of Political Economy*, 98(6), 1119–1158. <https://doi.org/10.1086/261729>

[7] Holmström, B., & Milgrom, P. (1991). Multitask Principal–Agent Analyses: Incentive Contracts, Asset Ownership, and Job Design. *The Journal of Law, Economics, and Organization*, 7(special issue), 24–52. https://doi.org/10.1093/jleo/7.special_issue.24

Secondary (empirical work on vertical integration)

[8] Monteverde, K., & Teece, D. J. (1982). Supplier Switching Costs and Vertical Integration in the Automobile Industry. *The Bell Journal of Economics*, 13(1), 206–213. <https://doi.org/10.2307/3003441>

[9] Masten, S. E. (1984). The Organization of Production: Evidence from the Aerospace Industry. *The Journal of Law and Economics*, 27(2), 403–417. <https://doi.org/10.1086/467071>

[10] Lafontaine, F., & Slade, M. (2007). Vertical Integration and Firm Boundaries: The Evidence. *Journal of Economic Literature*, 45(3), 629–685. <https://doi.org/10.1257/jel.45.3.629>

Secondary (losses at the handoff, and coordination)

[11] Gittell, J. H. (2002). Coordinating Mechanisms in Care Provider Groups: Relational Coordination as a Mediator and Input Uncertainty as a Moderator of Performance Effects. *Management Science*, 48(11), 1408–1426. <https://doi.org/10.1287/mnsc.48.11.1408.268>

[12] Arora, V., Johnson, J., Lovinger, D., Humphrey, H. J., & Meltzer, D. O. (2005). Communication Failures in Patient Sign-out and Suggestions for Improvement: A Critical Incident Analysis. *Quality and Safety in Health Care*, 14(6), 401–407. <https://doi.org/10.1136/qshc.2005.015107>

[13] Szulanski, G. (1996). Exploring Internal Stickiness: Impediments to the Transfer of Best Practice within the Firm. *Strategic Management Journal*, 17(S2), 27–43. <https://doi.org/10.1002/smj.4250171105>

[14] Argote, L., & Ingram, P. (2000). Knowledge Transfer: A Basis for Competitive Advantage in Firms. *Organizational Behavior and Human Decision Processes*, 82(1), 150–169. <https://doi.org/10.1006/obhd.2000.2893>

Secondary (mirroring: organisational structure and delivered structure)

[15] Conway, M. E. (1968). How Do Committees Invent? *Datamation*, 14(5), 28–31. https://www.melconway.com/Home/Committees_Paper.html (accessed 2026-09-22)

[16] MacCormack, A., Baldwin, C., & Rusnak, J. (2012). Exploring the Duality between Product and Organizational Architectures: A Test of the "Mirroring" Hypothesis. *Research Policy*, 41(8), 1309–1324. <https://doi.org/10.1016/j.respol.2012.04.011>

[17] Colfer, L. J., & Baldwin, C. Y. (2016). The Mirroring Hypothesis: Theory, Evidence, and Exceptions. *Industrial and Corporate Change*, 25(5), 709–738. <https://doi.org/10.1093/icc/dtwo27>

Theory-side source (the one the website itself cites)

[18] Bojarski, M., Del Testa, D., Dworakowski, D., Firner, B., Flepp, B., Goyal, P., Jackel, L. D., Monfort, M., Muller, U., Zhang, J., Zhang, X., Zhao, J., & Zieba, K. (2016). End to End Learning for Self-Driving Cars. <https://arxiv.org/abs/1604.07316> (accessed 2026-09-22)

Verification note: [1] HTTP 200; [15] 200; [18] arXiv page 200, arXiv API returns the title "End to End Learning for Self-Driving Cars", matching this list. [4][8][14][16] resolve to 200 directly; [2][3][5][6][7][9][10][11][12][13][17] return 403 to scripts (anti-scraping, not absence), and were checked against Crossref metadata for title, volume, issue and pages, all matching. **One correction is on record:** Masten 1984 was first written with DOI 10.1086/467070, which resolves but returns a different article's title in Crossref; the correct DOI, found by title search, is 10.1086/467071, and has been substituted.

Cross-references within this series (no DOI yet; cited by number)

- GI-WP-2026-P8, *One-way doors: irreversible actions in the primary market, and who presses them*
- GI-WP-2026-P11, *Borrowed theories: how frontier-technology concepts explain an investment bank's daily work*
- GI-WP-2026-P13, *Soft labels: why a review should record the options not taken*

Sources of the quoted Glacier Institute text

All quotations come from [1]

<https://glacier.mba/institute.html> (retrieved 2026-09-22; plain text obtained by stripping comments, scripts, styles and all tags — 10,079 lines). Line by line:

- 「该模式涵盖六大核心维度（信任为先）：信任、信息、叙事、节奏、关系、担当。」 — L675 (repeated at L678)
- 「庚辛资本的『全托管协作』意味着：企业保留所有重大事项的最终裁决权……交由庚辛资本统筹主导。」 — L676
- 「企业决策，庚辛统筹。企业家只做三件事：定战略、找对人、深入一线。全托管的本质是同侧——你的利益，就是我的利益。」 — L3220
- 「全托管六级：信任是地基，往上依次是信息、叙事、节奏、关系、担当；企业保留最终决策权。」 — L3230
- Chapter 18, pair 3, theory sentence 「从输入到输出，一个模型端到端负责；分成多段的流水线，误差会在段与段之间累积。」 — L3881
- Chapter 18, pair 3, phenomenon sentence 「从恢复事实到交割收口，同一个系统端到端负责，六个维度不换手——误差没有地方累积。」 — L3883
- Chapter 18, pair 3, source line "Bojarski et al. (2016), End to End Learning for Self-Driving Cars, arXiv:1604.07316, Figure 4" — L3885
- 「不是。投资人引荐只是工作的一部分，核心服务还包括事实恢复、资本路线设计、投资人结构、路演与投资人沟通、复杂交易设计、产业协同、尽调支持、条款协调和交割管理。」 — L4523
- 「全托管是庚辛的深度交易协作：企业保留重大事项最终决策权，交易信息、核心叙事、推进节奏、关键关系与交割安排交由庚辛统筹。」 — L4525
- 「orchestrator 是在企业关键资本窗口中，统一事实、资本表达、投资人角色、交易结构、产业关系、时间节奏和交割闭环的组织者。」 — L4521 (cited for the "time cadence" cell of Table 1)

One note on punctuation: the inner quotation marks in L676 are 「」 in the original; quoted inside an outer 「」 here, they are rendered 『』 per Chinese convention. Only the shape of the quotation marks changed; not one character did.

English renderings in italics are translations supplied for readers of this edition. The Chinese is the original and governs.

Redaction note: none of the nine quoted sentences names any client company, so no "a certain company" rewriting was required; rewrites = 0. Nothing in this paper's own prose describes any feature traceable to a specific company. See the writing note, §5.

Disclosure: Glacier Institute is the research arm of Glacier Capital, and this paper is issued under the name of Glacier Institute. In the course of its business, Glacier Capital acts as financial adviser to a number of technology companies and invests its own capital in some of them; such relationships may overlap with the industries discussed here. This paper does not concern any specific mandate and uses no non-public information; all company, industry and data references are drawn from public sources and cited individually. The authors received no third-party compensation for this paper.

Disclaimer: This is a methodological working paper and represents the authors' analysis at the time of writing only. It does not constitute investment advice, nor an offer or solicitation of an offer for any security, fund interest or other instrument, nor a commitment or forecast regarding the valuation, financing outcome or investment return of any company. It has not been peer reviewed and may be revised in later versions.