GO-TO-MARKET & PRICING

The integration savings come from what you put in

Integration plans are lists of things to remove. In the one model that measures all four moves, the largest cost lever is the one that adds.

Sinan Isoglu, MBA (Quantic)

Commercial growth leader, lecturer and doctoral researcher

PUBLISHED	UPDATED	READING TIME	WORDS
18 August 2026	20 August 2026	10 min	2,127

https://isoglu.com/insights/journal/integration-savings-come-from-what-you-put-in/

CONTENTS

MANAGEMENT SUMMARY

Post-merger integration plans are written as rationalisation lists, and the measured record shows that rationalisation falls overwhelmingly on the acquired company. In the one structural model that estimates all four available moves against the same outcome, cutting is the weakest of them. Divesting the acquirer's own assets raises cost savings; divesting the target's assets does not, and it carries a significant negative path into both revenue-based synergies. The largest coefficient on cost savings belongs to no cut at all: it is the redeployment of the acquirer's own resources into the acquired business. The direction of that transfer is not symmetric, and an integration plan almost never states which way anything is moving.

Keywords: Post-merger integration · Cost synergies · Asset divestiture · Resource redeployment · Commercial due diligence

The plan arrives about a fortnight after closing, and it is a list. Systems to consolidate, offices to combine, overlapping product lines to rationalise, functions to merge. It is a competent document, it is organised by workstream, and almost every line on it describes something the acquired company has.

That is not an accusation. It is what the deal model asked for. Cost synergies were underwritten before the wire cleared, and the overlap that justified them sits, by definition, in the business that was just bought.

The question worth asking is narrower than whether the plan is any good. It is whether removal is the thing that produces the savings.

What happens to acquired operational capabilities during post-merger integration?

The extent of it is not in dispute, and it has been measured on administrative data rather than on recollection. Vojislav Maksimovic, Gordon Phillips and N. R. Prabhala examined 1,483 mergers using plant-level records held at the US Census Bureau, following what acquirers did to the physical operations they had bought. Within three years of the acquisition, acquirers of full firms had sold 27% of the target's plants and closed a further 19%.

Read as disjoint, which is our reading and not the paper's arithmetic, that is close to half the acquired productive base in other hands or gone, inside three years.

One inference follows and I want to mark it as an inference rather than let it pass as evidence. From "divestiture lands on the acquired company" I am reading "the plan was written about the acquired company." Those are not the same claim. Divestiture is an outcome, a plan is a document, and I have not read a corpus of integration plans. What would settle it is a coded sample of real integration charters, which does not exist publicly and which I have not built. Until someone does, treat the step as a reasonable reading of the outcome data and not as a measured fact.

How does empirical M&A modeling value the four integration moves?

Most of what is known about integration comes from studies that measure one action at a time. There is an exception, and it is old enough that its sample has aged into a period piece while its structure has not.

Laurence Capron surveyed the managers of 253 horizontal acquisitions initiated by European and US manufacturers between 1988 and 1992, and estimated a structural model in which four post-acquisition actions compete to explain the same outcomes. The four are worth naming precisely, because the whole argument lives in the distinction between them:

1. Divesting the acquirer's own assets
2. Divesting the target's assets
3. Redeploying the acquirer's resources into the target
4. Redeploying the target's resources into the acquirer

In the practitioner sources I searched, nothing separates these four. Integration is discussed as depth and speed, as light-touch versus full absorption, as day-one readiness. The four moves above are treated as one activity called integration, and they are not one activity.

Why does pure cost-cutting deliver the lowest synergy return?

Take the two cuts first, since they are what the plan is made of.

Divesting the acquirer's own assets raises cost savings, at 0.19, and it is the most strongly estimated path in that part of the model. Divesting the target's assets does not: the coefficient is −0.09, and it is negative. Capron tested whether those two are genuinely different rather than comparing their significance stars, which is the correct and frequently skipped step: she compared the full model against a nested one constraining the two coefficients to be equal, and the difference holds at roughly the conventional threshold.

Two things should be said about that immediately, because the temptation is to make it carry more than it can.

The evidence that cutting your own assets works is strong. The evidence that cutting theirs actively hurts is weak: one coefficient, small, at a permissive threshold, from a single-informant survey. The honest statement is that the first is well supported, the second is suggestive, and the gap between them is real but not commanding.

The more interesting number is elsewhere in the same figure. Divestiture also runs a negative path of −0.11 into both revenue-based synergies, market coverage and innovation capability, and those paths are estimated considerably more precisely than the cost effect everyone is chasing. Removal buys a contested amount of cost and reliably costs something on the revenue side.

Why does capability transfer outperform headcount reduction in acquisitions?

Now the part that changes the plan.

In the same model, on the same outcome, the largest coefficient on cost savings belongs to none of the removals. It belongs to redeploying the acquirer's own resources into the acquired business, at 0.28.

Larger than cutting your own assets. Larger, obviously, than cutting theirs. And precisely estimated.

This is not a rhetorical inversion. It is the same dependent variable, in the same estimation, on the same 253 deals, which is the only basis on which these numbers may be compared at all. Read it plainly: the largest-coefficient route to the cost synergy in this model was to move capability in, not to take capacity out. What that capability costs to move is a question the model does not answer and this piece will not pretend to.

Capron says as much in her own discussion. Where standalone divestiture of the target fails to produce the expected savings, changes in how the acquired business operates, driven by resources moved into it, do reduce costs.

Why is integration direction an unstated strategic choice?

The fourth move behaves differently from the third, and the asymmetry is the practical finding.

Moving the acquirer's resources into the target helps on every route: cost savings at 0.28, market coverage at 0.17, innovation capability at 0.16.

Moving the target's resources into the acquirer helps the revenue routes more strongly, at 0.28 and 0.29, and then carries a direct negative path of −0.27 onto acquisition performance itself.

I am not going to build a prescription on that last number. It is a single-star estimate and it deserves a question rather than a rule. But the question is a good one, and no integration plan I have seen asks it: when a capability moves between the two organisations, which way is it moving, and has anyone decided that on purpose?

Capability, of course, is not a line item that sits still while you move it: much of it is attached to people who can leave, which is a cost of the transfer the model does not price.

The plan says "consolidate the CRM." It does not say whose CRM practice is being absorbed into whose. It says "harmonise the pricing function." It does not say which side's pricing capability survives contact. Those are the same decision, and the model says the two directions do not behave the same way.

Why has the cost-reduction synergy thesis always been empirically weak?

There is a final number that reframes the rest, and it sits in the same estimation.

Cost savings reaches acquisition performance at 0.22. Market coverage reaches acquisition performance at 0.69, more than three times the weight, and better estimated. Innovation capability sits at 0.22 as well, estimated more loosely still.

So the plan is optimised for the weaker of the two routes that the model estimates at conventional confidence, and the stronger one by a factor of three is the one integration is known to put at risk.

This site has made an adjacent argument before, from a different study, and it is worth reading alongside this one: integration depth buys cost savings by spending market position. Furthermore, to prevent post-close accountability from eroding, promised synergies need a ledger that audits actual headcount, timeline, and asset transfers over time. That piece and this one use different papers and different constructs, and they converge on the same uncomfortable place. The cost column is the one the plan can measure, and it is not the one that carries the deal.

What do post-merger synergy datasets not establish?

Three limits, stated here rather than in a footnote.

It does not say that cutting destroys value. Capron, with Will Mitchell and Anand Swaminathan, later showed on the same 253 deals that divestiture is best read as a logical consequence of reconfiguring two businesses into one, and not as evidence of failure. Divestiture that follows a reconfiguration decision is a different act from divestiture that substitutes for one.

It does not transfer coefficients between studies. The 27% and 19% come from US census plant records covering 1981 to 2000. The path coefficients come from a survey of 253 European and US manufacturing deals from 1988 to 1992. They describe the same behaviour from two directions and they may not be arithmetically combined.

The sample is old and narrow, and one feature of it works in your favour. Manufacturing, horizontal deals, a single informant on the acquirer's side. It is also 70% cross-border, which is closer to the situation of most readers of this journal than a domestic-only sample would be, and closer than the figures that circulate about European deal failure. The numbers belong to their period. The structural point, that the four moves are different moves, does not obviously expire.

I also could not find the acquirer-versus-target split made in the practitioner sources I searched, in English or in German, nor in the turnaround and restructuring vocabulary, where the segmentation is by profit contribution per division, facility or product line rather than by which side of the deal an asset arrived from. That is a statement about what I searched, not a claim about everything that exists.

Which question should executives bring to the post-merger integration review?

Not a scorecard, and deliberately not a threshold: the evidence here cannot support a number, and inventing one would be worse than saying nothing.

Take the rationalisation lines in the current plan. For each one, ask a question with a right answer:

Was the acquirer-side equivalent considered and rejected, or was it never raised?

“Considered and rejected” is a good answer. It means someone looked at both balance sheets and made a decision, which is what the strongest coefficient in Capron’s model rewards. “Never raised” is a finding, and it is a finding about the plan rather than about the deal.

There is one honest complication. In many organisations the integration team is staffed only to look at the acquired company, so “never raised” will sometimes be a fact about the org chart rather than about anyone’s judgement. That is worth knowing too, and it is cheaper to discover in a planning meeting than in the year-two review.

Then ask the second question, which the plan almost certainly does not answer: for every capability being combined, which direction is it moving, and did anyone choose?

Is the counter-argument available? Certainly. The acquirer’s operations were already optimised, the target’s overlap is the entire reason the deal cleared its hurdle rate, and there is nothing left on your own side to cut. That may be true in a given deal. The reason I do not think it is true in general is that acquirer-side capability turns out to be a variable rather than a constant: Maksimovic and colleagues found that acquirers who are demonstrably skilled at running their peripheral divisions retain more of what they buy. Skill differs between acquirers. Where something differs, there is something to look at.

The plan will always be easier to write as a list of subtractions. Subtractions have owners, dates and a number attached. The measured record says the largest cost lever in the set was an addition, and it was pointed at the company that had just been bought.

Where are the methodological boundaries of synergy realization?

Boundary. The four moves are a decision decomposition, not a validated savings formula. Test the direction, cost and revenue effect of each move in the actual integration plan.

Evidence base. The analytical frame also draws on these additional sources: Capron (1999); Capron et al. 2001; Maksimovic et al. 2011. The links identify the exact works; they support the mechanisms and boundary conditions discussed here, not every claim in isolation.

Capron, L. (1999). The long-term performance of horizontal acquisitions. Strategic Management Journal, 20(11), 987–1018.

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

Maksimovic, V., Phillips, G., & Prabhala, N. R. (2011). Post-merger restructuring and the boundaries of the firm. Journal of Financial Economics, 102(2), 317–343. https://doi.org/10.1016/j.jfineco.2011.05.013

END OF ESSAY

HOW TO CITE THIS ESSAY

Isoglu, S. (2026, August 18). The integration savings come from what you put in. Sinan Isoglu. https://isoglu.com/insights/journal/integration-savings-come-from-what-you-put-in/

KEEP READING

Go-to-market & pricing

The integration plan spends what the deal bought.

https://isoglu.com/insights/journal/integration-spends-what-the-deal-bought/

Go-to-market & pricing

What is price adjustment cost? Changing a price changes more than a number

https://isoglu.com/insights/journal/what-is-price-adjustment-cost/

Go-to-market & pricing

What is price fairness? A reference and process problem

https://isoglu.com/insights/journal/what-is-price-fairness/

Sinan Isoglu, MBA (Quantic)
Commercial growth leader, lecturer and doctoral researcher

Online version: https://isoglu.com/insights/journal/integration-savings-come-from-what-you-put-in/