



GROWTH THAT COMPOUNDS

A customer P&L needs a cost boundary

Revenue is not customer profitability. Trace order, account, market, and shared-capacity costs to the object and decision they can support.

Sinan Isoglu, MBA (Quantic)

Commercial growth leader, lecturer and doctoral researcher

PUBLISHED

25 August 2026

READING TIME

6 min

WORDS

1,266

<https://isoglu.com/insights/journal/a-customer-p-and-l-needs-a-cost-boundary/>

Management summary	2
Why must a customer P&L start with the cost object rather than the dashboard?	3
Why is support ticket volume an incomplete measure of account cost?	3
How does cost hierarchy modeling expose hidden account servicing burdens?	3
How do activity-based costing methods alter customer profitability rankings?	4
Why do customer-facing coordination costs often exceed direct product labor?	5
How should leadership structure cost allocation decisions across accounts?	5
Where are the empirical boundaries of customer profitability accounting?	5
References	6

MANAGEMENT SUMMARY

A customer can produce attractive revenue and still consume service capacity that the gross-margin view does not show. Helgesen's customer-profitability working paper offers a useful boundary: keep revenue at the order object, assign costs to the level where they are incurred, and aggregate only after the object is visible. Its 564 orders, 176 customers, and 36 markets are a Norwegian fish-export setting, not a software benchmark. The paper also shows why the method matters: full, variable, and activity-based costing produce different estimates. Add the cost of customer-facing work carefully, and keep shared capacity shared when the driver is not defensible.

Keywords: Customer profitability · Cost to serve · Customer economics · Growth measurement

Revenue is the easiest customer number to see. It is also the easiest number to mistake for an explanation.

An account can grow while consuming implementation time, support capacity, infrastructure, credit work, custom requests, or relationship effort that the gross-margin view does not show. The question is not whether those costs matter. It is which customer object they belong to, what driver makes the assignment defensible, and what decision the resulting P&L is meant to support?

Why must a customer P&L start with the cost object rather than the dashboard?

Helgesen's working paper studies four Norwegian exporters of klipfish and frozen fish or fillets and their customers. It combines customer accounts and profitability analysis with market-survey measures. The accounting sample covers 564 orders, 176 customers, and 36 geographical markets, with about 180 million NOK in order revenue.

That setting matters. It is an industrial export context, not a software benchmark or a universal cost-to-serve distribution. The value of the paper is the accounting boundary it makes visible. Revenue is assigned at the order level. Costs are assigned to the level where they are incurred. Results are then aggregated from orders to customers and markets.

This gives a first diagnostic: what is the object of the decision? An order P&L can ask whether a transaction covered its direct and order-related costs. A customer P&L can ask what the set of orders and account work produced. A market P&L can ask whether a route or geography is attractive. Those are different questions, even when they share a revenue field.

Why is support ticket volume an incomplete measure of account cost?

Direct work is often easy to trace. A support interaction has an account, an implementation project has a customer, and a custom request has a requesting team. That traceability is useful, but it does not automatically make the whole capacity cost account-specific.

The support engineer's time may be directly assigned. The training that made the support system usable, the manager's capacity, the tooling, and the unallocated queue may remain shared. If the P&L assigns all of those costs to the account because one ticket opened the work, the account result becomes an allocation convention presented as a customer fact.

The reverse error is possible too. If recurring customer work is always buried in a shared platform line, the P&L can make a high-touch account look like a low-touch account. The practical boundary is therefore not "direct costs good, shared costs bad." It is: which cost object was observed, where was the cost incurred, and what driver links shared work to this decision?

How does cost hierarchy modeling expose hidden account servicing burdens?

Table 1 The customer P&L boundary

Level	Revenue or work object	Cost classes to inspect	Defensible driver	What the result can say	Failure mode
Order	Transaction, invoice, delivery, or implementation unit	Product, freight, transaction, rush, order handling	Order lines, events, hours, or documented transaction driver	Whether this order covered the chosen order boundary	Order margin is called customer profitability
Customer	Set of orders and account-specific work	Support, account management, customer credit, custom service	Account events, time, cases, or a tested service driver	Customer result under a stated method	Shared capacity is assigned by habit
Market	Customers and orders in a route, segment, or geography	Market development, channel, regulatory, and route costs	Market activity or route-level driver	Market or channel result under a stated scope	Country is used as a proxy for every cost
Business unit	Shared infrastructure and common capacity	Platform, leadership, tooling, unused capacity	A declared capacity or allocation rule	Unit economics under the chosen reporting purpose	The allocation is mistaken for causal cost

Author's framework grounded in Helgesen (2000) and Zbaracki, Ritson, Maklan, and Dean (2004). Helgesen's working paper is a Norwegian export setting; Zbaracki et al. is one price-adjustment firm. Neither supplies a universal service-cost allocation rule.

How do activity-based costing methods alter customer profitability rankings?

Helgesen explicitly notes that full costing, variable costing, and activity-based costing produce different customer-profitability estimates. The paper uses an activity-based hierarchy, but that does not turn the output into the one true customer profit. It makes the method legible.

That distinction matters when a manager changes service policy. A variable view may be useful for a short-run capacity decision. A fuller view may be useful for a long-run account portfolio decision. An activity-based view may expose the work objects and drivers that a broad overhead rate hides. The same account can look different under each purpose without one calculation being dishonest.

The minimum record is therefore not just a number. It includes the reporting purpose, the cost method, the period, the included revenue, the included work, the driver, and the costs deliberately left shared.

Why do customer-facing coordination costs often exceed direct product labor?

Zbaracki and colleagues studied the costs of price adjustment in one firm. Their reported comparison put managerial costs at more than six times menu costs and customer costs at more than twenty times menu costs. They also reported a total equal to 1.22% of revenue and 20.03% of net margin in that firm and year.

Those ratios are not a service benchmark. They are a warning about the object. A price change can create internal management work and customer-facing work, and a ledger that records only the internal effort will miss part of the economic event. The same logic applies to implementation, support, escalation, credit, and custom service. The cost category is a hypothesis until its object and driver are visible.

How should leadership structure cost allocation decisions across accounts?

Suppose a team wants to decide which accounts deserve a different service level. The customer P&L may need account-specific support and implementation effort, a defensible share of recurring service work, and a separate view of shared platform capacity. Suppose the team instead wants to decide whether to add a new support region. The market or business-unit object may matter more than the account result.

The correct boundary can change with the decision. That is not inconsistency. It is a reason to write the purpose into the P&L rather than silently changing the denominator.

Before a service policy changes, record:

- the customer, order, market, or business-unit object;
- the revenue and time window included;
- the directly traceable work;
- the shared work and its driver, if any;
- the costs left unallocated and why; and
- the outcome that would show the policy helped.

Where are the empirical boundaries of customer profitability accounting?

Helgesen's context cannot establish how a software company should allocate cloud infrastructure or support. Zbaracki's one-firm price-adjustment ratios cannot establish a general cost-to-serve curve. Neither source tells a manager which customer to fire, which service to remove, or which account is strategically valuable.

They support a narrower conclusion: customer profitability is a constructed result whose boundary, cost objects, drivers, method, and decision purpose must be visible. Revenue can start the analysis. It cannot finish it.

The boundary belongs beside customer lifetime value as a forecast and the cost of goods sold question, because both keep the economic object visible before a number is treated as a decision.

REFERENCES

- Helgesen, Ø. (2000). Are loyal customers profitable? Customer satisfaction, customer loyalty and customer profitability at the individual level. SNF Working Paper 07/2000. https://snf.no/media/m0pdbai2/a07_00.pdf
- Zbaracki, M. J., Ritson, M., Levy, D., Dutta, S., & Bergen, M. (2004). Managerial and customer costs of price adjustment: Direct evidence from industrial markets. *The Review of Economics and Statistics*, 86(2), 514–533. <https://doi.org/10.1162/003465304323031085>

END OF ESSAY

HOW TO CITE THIS ESSAY

Isoglu, S. (2026, August 25). A customer P&L needs a cost boundary. Sinan Isoglu.
<https://isoglu.com/insights/journal/a-customer-p-and-l-needs-a-cost-boundary/>

KEEP READING

Growth that compounds

Long customer relationships can be low-profit

<https://isoglu.com/insights/journal/long-customer-relationships-can-be-low-profit/>

Growth that compounds

The 20-55 rule: customer prioritization misclassifies the portfolio

<https://isoglu.com/insights/journal/the-20-55-rule-customer-prioritization-misclassifies-the-portfolio/>

Growth that compounds

One customer model cannot predict every outcome

<https://isoglu.com/insights/journal/one-customer-model-cannot-predict-every-outcome/>



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