

What is sales forecasting?

Sales forecasting estimates what will sell, for a named unit and period, before the actual is known. It is not the target, and it needs a baseline.

Author	Published	Reading time	Words
Sinan Isoglu, MBA (Quantic) Commercial growth leader, lecturer and doctoral researcher	18 October 2026	12 min	2,670

Management summary Sales forecasting is estimating what will sell, of a named product or group, over a named period, before the actual is known, so that a decision can be taken in time. This concept guide separates the forecast from the target, the plan, the pipeline total and the forecast category; shows how forecasts are made in practice, with judgment often layered on a system forecast; reports what two field studies found about accuracy and bias, and what an incentive analysis says about salespeople's forecasts; and gives a one-page record that keeps the estimate, its baseline and its adjustments apart so a miss can be read.

Contents	§ 1	What is sales forecasting?	2
	§ 2	How are sales forecasts made?	2
	Figure 1	How often planners adjusted the system forecast	3
	§ 3	What is a sales forecast not?	4
	Figure 2	Five numbers, five questions	4
	§ 4	How accurate are company sales forecasts?	6
	§ 5	Why are sales forecasts biased?	7
	§ 6	How do you write a sales forecast down?	8
	Figure 3	A sales forecast record	9
	§ 7	Where does the evidence on sales forecasting stop?	10
		References	11

What is sales forecasting?

§ 1 of 7

Sales forecasting is estimating what will sell, of a named product or group, over a named period, before the actual is known, so that a decision can be taken in time. Lawrence et al. (2000)² name the decisions plainly: “Sales forecasting plays a pivotal role in decision making in relation to manufacturing scheduling and inventory replenishment.” Mantrala and Raman (1990)³ add why a miss costs money in either direction: firms need accurate sales forecasts “to set plans that reflect growth opportunities as well as avoid costs such as those due to stock-outs and excess inventories.”

² Lawrence, M., O'Connor, M., & Edmundson, B. (2000). A field study of sales forecasting accuracy and processes. *European Journal of Operational Research*, 122, 151-160.

³ Mantrala, M. K., & Raman, K. (1990). Analysis of a sales force incentive plan for accurate sales forecasting and performance. *International Journal of Research in Marketing*, 7, 189-202.

A sales forecast is a dated estimate that is later judged against what sold. In this guide that makes four things part of the definition: the object (what is forecast, in which unit and at which grain, such as product, customer segment, region or account), the horizon (which period closes the forecast), the information cutoff (what was known when it was made) and the decision it feeds. Leave one out and a later miss cannot be read: it may be a forecasting error, or a change in what was being forecast.

How are sales forecasts made?


§ 2 of 7

With judgment, often layered on a system forecast. In the mid-1990s surveys that Lawrence et al.² (2000) cite, management judgment was the preferred method: despite the wide availability of software, “most surveys show that management judgement is the preferred method for forecasting”. Fildes et al.¹ (2009) describe the usual approach in supply-chain companies as a system forecast that planners then adjust: “The most common approach to forecasting demand in these companies involves the use of a computerized forecasting system to produce initial forecasts and the subsequent judgmental adjustment of these forecasts by the company’s demand planners”. All four companies they studied worked that way, with a forecasting meeting agreeing the final number, and how often planners stepped in varied widely (Figure 1).

¹ Fildes, R., Goodwin, P., Lawrence, M., & Nikolopoulos, K. (2009). Effective forecasting and judgmental adjustments: An empirical evaluation and strategies for improvement in supply-chain planning. *International Journal of Forecasting*, 25(1), 3–23.

Figure 1 How often planners adjusted the system forecast

Planners at three manufacturers adjusted between 63% and 91% of system forecasts; at a retailer forecasting weekly, 10% to 14%.



The same practice, judgment on top of a system forecast, ran at very different rates. A record that keeps only the final number cannot show which kind of process produced it.

Source Fildes, Goodwin, Lawrence and Nikolopoulos (2009), Table 1, p. 8: 68,984 complete triples of system forecast, final forecast and actual from three manufacturers forecasting monthly and one retailer forecasting weekly for two product groups, 2003 to 2005; 21% of forecasts adjusted overall. Demand forecasts at SKU level.

The judgment carries information a statistical model does not see. In the forecasting meetings Lawrence et al. (2000)² attended, “New marketing initiatives, promotion plans, actions of competitors, industry developments, manufacturing problems” and other contextual information dominated the discussion, and “time series data received relatively little attention.” Salespeople hold a second kind of information: Mantrala and Raman (1990)³ start from the premise that forecast accuracy can improve with information “from salespeople who possess better knowledge than central planners about their territories’ sales prospects.”

Three routes follow, and one process can combine them: a statistical baseline built from sales history, a judgmental forecast or adjustment by managers using contextual information, and a bottom-up forecast submitted by salespeople. What matters for the record is that each contribution stays visible.

What is a sales forecast not?

§ 3 of 7

Four numbers sit next to the forecast and can be mistaken for it: the target, the plan or budget, the pipeline total and the forecast category (Figure 2). Each answers a different question.

Figure 2 Five numbers, five questions

Only the forecast answers what will sell. The other four answer what to demand, what to commit, what is open and which declared state a deal is in.

Number	The question it answers	Judged against
Sales forecast	What will sell, for this object and period, given what is known now?	The actual sales for that object and period
Sales target or quota	What should be asked of a team or a person?	Attainment
Plan or budget	What will the firm commit resources to?	Budget variance
Pipeline total	What open opportunity value exists now?	Nothing on its own; conversion still has to be estimated
Forecast category	Which declared state is a deal in, under which entry rule?	Its entry rule, and later the outcome

Each of the four can feed a forecast. None of them is one.

Source Author's comparison. Mantrala and Raman (1990) and Lawrence, O'Connor and Edmundson (2000) both describe settings in which the forecast and the target were one number (Mantrala and Raman 1990, p. 190; Lawrence, O'Connor and Edmundson 2000, p. 158); the table separates them by the question each number answers. The pipeline and category rows follow this site's pages on pipeline coverage and forecast categories.

The target asks what should be demanded, not what will happen, and mixing the two can bend the estimate. In the analysis of Mantrala and Raman (1990)³, a salesperson paid a fixed bonus for meeting a quota they set themselves has an incentive to understate it; they note that “salespeople may understate sales estimates when their compensation depends on fulfillment of quotas based on their sales forecasts.” Lawrence et al. (2000)² offer the same mechanism as one possible explanation of the bias they measured, not as something they observed: “If managers are rewarded for exceeding the sales target, there may be pressure for the forecasts to be less than the actual”.

The plan commits resources, and the meeting that produces it serves more than accuracy. In the companies Lawrence et al. (2000)² studied, forecasting and budgeting ran through a monthly meeting of all interested parties. In the meetings the authors attended, “The major objective appeared to be gaining commitment to the target final forecast from all parties”: sales, marketing, manufacturing, purchasing and finance. The phrase makes one number both the target and the forecast. “Communication also appeared to be an important objective”. Commitment to a number is useful. It is not the same thing as an estimate, which is why this guide keeps the two apart.

Fildes et al. (2009)¹ found a related confusion in the field: at the retailer they studied, “managers were confusing forecasts of demand with decisions on the levels of inventory required to meet customer service levels”, for example raising a forecast of 200 units to 250 to cut the risk of a stock-out.

A pipeline total is inventory: open opportunity value that still has to convert. A forecast category is a declared state assigned under a rule, with an owner, later compared with what happened. Both can feed a forecast; neither is one.

How accurate are company sales forecasts?

§ 4 of 7

It depends on the company: in the two field studies cited here, judgmental forecasts beat a naive forecast in some firms and not in others. Lawrence et al. (2000)² compared one-month-ahead forecasts by 13 large Australian manufacturers of branded consumer goods and durables, in a paper that reached the journal in January 1998, with a naive forecast that “is the most recent actual value”, not even seasonally adjusted. Their database held around 24,000 forecasts and 4,500 actual sales values for products each company named as important for forecasting accuracy; the one-month-ahead comparison covers about 4,600 forecast and actual pairs. For 4 of the 13 companies the judgmental forecasts were significantly more accurate; for 2 the naive forecast was significantly better; for 7 the difference was not significant. The authors conclude that “the company judgemental forecasts are not generally more accurate than naive forecasts even though the naive forecasts do not incorporate any contextual information”, while finding “no strong evidence that, in general, the companies should abandon judgement and embrace a naive forecast.”

They attribute most of the error to two sources, “serial correlation in the errors” and “bias in the forecasts”. In their regression test the four companies that beat the naive forecast showed neither, although two of the four were significantly biased in the sign test reported below.

Fildes et al. (2009)¹ made a similar comparison at four UK-based supply-chain companies in 2003 to 2005. On the median error the authors focus on, the three manufacturers’ final forecasts, taken together, were more accurate than a naive forecast; at the retailer, “both the system and final forecasts were less accurate than the naïve forecasts”.

Adjusting a system forecast does not settle the question either. Fildes et al. (2009)¹, with more than 60,000 forecasts and outcomes, found that “In three of the companies, on average, judgmental adjustments increased accuracy.” At the retailer, adjustments did not significantly improve accuracy at all. Direction mattered most: upward adjustments were much less likely to improve accuracy than downward ones and were more often made in the wrong direction, “suggesting a general bias towards optimism”. Size mattered inside that. The smallest quarter of upward adjustments at the manufacturers lowered accuracy, and at the retailer upward adjustments lowered it whatever their size, while downward adjustments paid off at every company except for the retailer’s largest. The authors expected that stopping small adjustments would “marginally enhance accuracy”. In their

hold-out test it gave little improvement for upward adjustments and cut accuracy for downward ones, and they concluded it was “not likely to lead to improvements in accuracy, whatever the information direction”. The case they keep for it is the planners’ time. How such errors are measured and summarised is the subject of forecast accuracy can hide offsetting errors.

Why are sales forecasts biased?

§ 5 of 7

Partly because the organisation prefers them to be, on the authors’ reading. Every company told Lawrence et al. (2000)² it wanted accurate forecasts; from closer questioning in the opening interviews, the authors judged that six would rather under-forecast, five would rather over-forecast and two preferred no bias in either direction. In a sign test, eight of the 13 companies showed significant bias, seven of them towards over-forecasting, and across all 4,568 forecasts 56% were above the actual. The bias ran against the elicited preference in two companies and appeared in one that preferred no bias; on that basis the authors conclude that the results “suggest that the bias in the company forecasts reflects company preference”.

The pressure can run both ways. Rewards for beating a target may push forecasts down, and worry about lost sales may push them up: “Conversely, if product managers are concerned that there be no lost sales, there may be a tendency to over-forecast.” In this study the measured bias ran mostly upward. The executives in the study did not call it bias. Asked directly, they told the authors “we try to be spot on”, and spoke of “hedging a bet” or “playing safe”; only close questioning about errors above or below the actual produced answers the authors read as a policy of bias.

Can the estimate be kept apart from the bias? A reviewer of the paper suggested exactly that: “constructing the most accurate forecasts that they can and then explicitly adjusting these forecasts to account for organisational objectives”. The authors called the idea “a rational way to proceed” but doubted it could be done, because forecasting and “biasing are tightly interwoven”. The record below is a version of it, so it is offered as this guide’s design, untested in the studies cited here. Two features may make it more workable than asking about bias after the fact: the preferred direction is set before the period, from what each kind of miss costs, as in the Gonik plan below; and each adjustment carries a recorded reason, which Fildes et al. (2009)¹ suggest requiring in a standard format. The first still needs someone to state a preferred error openly, which Lawrence et al.² doubted executives would be willing or able to do.

Incentive design is the other response, and it works on the fused number rather than separating it. Mantrala and Raman (1990)³ analyse the Gonik plan, which rewards a salesperson for “self-selecting a demanding forecast that becomes his quota” and for the sales then achieved against it. They treat quota, target and forecast as one number, which they say is often the case in practice. Management sets the plan’s parameters so that salespeople “report sales estimates with the probabilities of fulfillment desired by management”, chosen from what it costs to miss low and to miss high. That is a preferred direction of error written into pay, the same idea as the last field of the record below. Their analysis is a one-period model whose practical directions assume a risk-neutral salesperson, not a field test, and the case for such a plan rests on what forecast errors cost: it is the firm’s “loss function associated with prediction errors rather than individual sales representatives’ risk attitudes that is a dominant concern underlying the use of the Gonik plan”.

Read together: judgment brings information a model lacks, and the field evidence shows it also brings bias. Lawrence et al.² read some of it as chosen; Fildes et al.¹, whose forecasters called accuracy the principal objective of the process, could not tell incentive from cognitive bias. Neither study shows a method that wins everywhere. What a record can do is keep the estimate, the baseline it is compared with and the preferred direction of error visible, so a miss can be read.

How do you write a sales forecast down?

§ 6 of 7

Write one row per forecast object and period with five fields, in this order: object and grain, period and cutoff, baseline and final, adjustments, and preferred error and use (Figure 3). The fields are also the data a forecast needs: without the baseline, the cutoff and the actual, no later review can say whether the process added anything.

Figure 3 A sales forecast record

Write the object, the period and cutoff, the baseline and the final number, every adjustment with its reason, and the preferred direction of error in separate fields.

Object and grain	Period and cutoff	Baseline and final	Adjustments	Preferred error and use
What is forecast, in which unit and at which grain: product, segment, region?	Which period closes the forecast, and on which date was information frozen?	What did the system or the naive forecast say, and what is the final number?	Who changed the baseline, in which direction, by how much, and on what information?	Which error would the business rather make, and which decision does this forecast feed?

A final number with no baseline cannot show whether judgment helped. A forecast with no stated preferred error cannot show whether its bias followed a stated policy.

One row per forecast object and period. If the baseline field is empty, the review has nothing to compare the final number with.

Source Author's decision worksheet. The baseline field follows the naive benchmark in Lawrence, O'Connor and Edmundson (2000); the adjustment field follows the system-plus-judgment process in Fildes, Goodwin, Lawrence and Nikolopoulos (2009) and their call for better baselines and recorded adjustment reasons (p. 18). The fields are reader inputs; no company data is supplied.

- Object and grain.** Name what is forecast, in which unit and at which grain: product, customer segment, region or account. A total can look accurate while its parts miss in opposite directions.
- Period and cutoff.** Name the period that closes the forecast and the date on which information was frozen, so the forecast is judged on what was known then.
- Baseline and final.** Keep a baseline next to the final number: a statistical forecast, or at least the naive forecast of the last actual, the benchmark Lawrence et al.² used.
- Adjustments.** When someone changes the baseline, record who, in which direction, by how much and why. A recorded change of this kind is a forecast override.

- 5 **Preferred error and use.** Say which way the business would rather be wrong, and why, and which decision the forecast feeds, before the period starts. Then a bias in the results can be compared with a policy stated in advance.

Where does the evidence on sales forecasting stop?

§ 7 of 7

The field evidence comes from 13 large Australian manufacturers of branded consumer goods and durables, in a paper that reached the journal in January 1998 (Lawrence et al., 2000)², and from three manufacturers and a retailer in the UK forecasting SKU demand one step ahead between 2003 and 2005, whose planners, the authors note, had minimal training in forecasting (Fildes et al., 2009)¹. The incentive analysis is a one-period model whose practical directions assume a risk-neutral salesperson (Mantrala & Raman, 1990)³. None of it studies B2B pipeline revenue, SaaS bookings or services, so applying it there is reasonable as a structure for thinking and an assumption, not a finding.

The naive benchmark is a choice, not a law. Lawrence et al. (2000)² adopt it as a “conservative benchmark”, and changing it changes the count: against a naive forecast built on two-month-old data, the company forecasts were significantly better in five companies, significantly worse in three and not significantly different in five. No study cited here measures what better forecasting is worth in money.

 End of essay

References

- 1 Fildes, R., Goodwin, P., Lawrence, M., & Nikolopoulos, K. (2009). Effective forecasting and judgmental adjustments: An empirical evaluation and strategies for improvement in supply-chain planning. *International Journal of Forecasting*, 25(1), 3-23.
<https://doi.org/10.1016/j.ijforecast.2008.11.010>
- 2 Lawrence, M., O'Connor, M., & Edmundson, B. (2000). A field study of sales forecasting accuracy and processes. *European Journal of Operational Research*, 122, 151-160.
[https://doi.org/10.1016/S0377-2217\(99\)00085-5](https://doi.org/10.1016/S0377-2217(99)00085-5)
- 3 Mantrala, M. K., & Raman, K. (1990). Analysis of a sales force incentive plan for accurate sales forecasting and performance. *International Journal of Research in Marketing*, 7, 189-202.
[https://doi.org/10.1016/0167-8116\(90\)90021-E](https://doi.org/10.1016/0167-8116(90)90021-E)

How to cite this essay

Isoglu, S. (2026, October 18). What is sales forecasting? Sinan Isoglu.
<https://isoglu.com/insights/journal/what-is-sales-forecasting/>
Online version: <https://isoglu.com/insights/journal/what-is-sales-forecasting/>

Sinan Isoglu, MBA (Quantic)

Commercial growth leader, lecturer and doctoral researcher

© 2026 Sinan Isoglu · isoglu.com · All rights reserved