

Optimal Order Quantity: Cost-effective ordering with the batch size formula

The optimal order quantity refers to the order or batch size at which the sum of procurement, ordering and inventory holding costs over the period under review is minimised. It is one of the central variables in operational procurement and inventory management and is classically determined using the Andler batch size formula or the Harris formula. Anyone who maintains and regularly updates it reduces capital tied up and warehousing costs without jeopardising the ability to deliver.

Key facts

- The optimal order quantity minimises the total cost of ordering and inventory holding costs.
- The calculation basis is typically the Andler formula (optimal order quantity according to Andler) or its further development by Harris and Wagner/Whitin.
- The key input variables are annual requirement, ordering costs per transaction and the inventory holding cost rate.
- In practice, minimum order quantities, price tiers, discounts and transport costs are also factored into the calculation.
- The optimal order quantity is not a static value: it must be adjusted to changing demand, prices and capital costs.
- It forms the basis for planning parameters such as order interval, safety stock and reorder point.

Definition: optimal order quantity

The optimal order quantity is the quantity procured per order or production order when total procurement-side costs per unit of time are minimal. It balances two opposing cost drivers: frequent orders increase fixed ordering costs, while large order quantities increase inventory holding and capital tie-up costs.

Distinguishing it from the optimal batch size

In procurement, one speaks of the optimal order quantity; in production, of the optimal batch size; both follow the same logic and are frequently calculated using the same formulas. The difference lies in the cost drivers: with in-house production, set-up and changeover costs take the place of ordering costs. For buyers, this means: the same methodology is applicable across procurement and production, provided the fixed costs per transaction are captured correctly.

Position within requirements planning

The optimal order quantity is both an output and an input parameter of material requirements planning. Together with safety stock and reorder point, it determines how and when purchase requisitions are generated. If the value changes, this must be updated in the material master records – this is where planning parameter maintenance comes in, ensuring that system values reflect the real cost situation.

Methods and approach for determination

The standard route is via the Andler batch size formula, which weighs ordering and inventory holding costs against each other as a square root formula. It assumes constant demand, linear costs and prices independent of quantity – conditions that frequently only apply approximately in practice. The calculated optimal order quantity is therefore usually understood as a starting point, which is then supplemented with operational constraints.

The Andler formula at a glance

The formula is: optimal order quantity = square root of $(2 \times \text{annual requirement} \times \text{ordering costs})$ divided by $(\text{purchase price} \times \text{inventory holding cost rate})$. Inventory holding cost rates typically lie between 15 and 30 per cent of the value of goods per year, with capital costs, warehouse space, insurance and shrinkage factored in. Ordering costs capture not just the postage, but all transaction costs from purchase requisition through to invoice verification – in practice often a multiple of what is intuitively estimated.

Extensions for practice

Real-world procurement situations require extensions: volume discounts and price tiers shift the cost function, minimum order quantities and pack sizes set lower limits, and limited storage space sets upper limits. Financing headroom, seasonal pricing models and delivery dates also play a role. In practice, it has proven worthwhile to calculate the theoretical optimal order quantity and then adjust it to manageable rounded values, pack sizes and framework agreements.

Step order quantities with discount tiers

With price tiers, it is checked whether the next discount tier is cheaper despite higher inventory holding costs. The approach: calculate the optimal order quantity, check which tier it falls into, and compare the total costs (material plus ordering plus holding costs) with those of the next higher tier. If the calculated quantity jumps a discount threshold, it is often worthwhile to procure exactly at the level of that threshold.

Key KPIs around the optimal order quantity

Whether the chosen order quantities are actually economical only becomes apparent in the steering KPIs of requirements planning. These KPIs should be regularly compared against the calculated optimal order quantity.

Coverage and turnover frequency

Inventory coverage shows how many periods the stock will last at current consumption; turnover frequency relates consumption to average stock. If coverage rises above the target values, the chosen order quantities tend to be too large. A systematic inventory analysis identifies items with persistently excessive coverage and provides the initial signal for adjusting the optimal order quantity.

Order frequency and transaction costs

The number of orders per material and period is the direct mirror of the chosen batch size. Frequent very small orders indicate order quantities that are too low – with a rising share of transaction costs in the procurement volume. A clean supplier segmentation matrix helps here, separating A parts with intensive attention from C parts with automated, high-volume ordering routines.

Delivery capability and order reliability

The trade-off in every optimisation is delivery capability: order quantities and stock levels that are too tight increase the risk of missing parts and production stoppages. KPIs such as on-time delivery, shortage rate and the number of emergency orders show whether savings in stock were bought at the price of delivery problems. Emergency orders are a double cost signal – higher procurement costs and potential production downtime.

Risks and controls in order quantity management

The optimal order quantity is only as good as its input parameters. Outdated demand data, incorrect cost rates or ignored price tiers lead to systematically incorrect values. A robust control framework prevents errors from perpetuating themselves over years.

Data quality and master data maintenance

The most common weak point is maintenance errors in the material master: outdated prices, estimated ordering costs or never-updated cost rates. Mandatory fields for planning parameters and documented change processes provide a remedy here. Consistency of vendor master data is also important, since minimum order quantities, pack sizes and delivery terms flow from supplier master data into the ordering logic.

Process controls

A periodic review cycle – quarterly in practice for A parts, annually for B and C parts – checks whether demand, prices and inventory cost rates are still correct. Plausibility rules in the ERP system (e.g. warning when order quantities fall below the minimum order quantity) and approval levels for quantity-relevant changes complement the technical side. Procurement should assign and document responsibility for the parameters clearly.

Supplier and procurement risks

The optimal order quantity is based on the assumption of reliable suppliers. For critical sources, supplier-side risks – such as financial strength or delivery capability – should be assessed via supplier status and a Supplier Lifecycle Management (SLM) approach. With uncertain sources, an order quantity deliberately above the calculated optimum can be rational; alternative sourcing and supplier risk management complement this.

Current developments and digital support

Modern ERP and replenishment systems no longer calculate optimal order quantities statically, but dynamically from forecasts and current cost data. At the same time, sustainability requirements and supply chain transparency are changing the classic cost analysis.

Dynamic planning and AI support

AI-supported demand forecasts adjust order quantities to seasonal patterns and demand signals instead of working with annual averages. In addition, transport costs, load optimisation and minimum order quantities can be factored in algorithmically. The human buyer nonetheless remains responsible: systems provide proposals, while approval and strategic assessment remain the responsibility of procurement.

Sustainability and packaging aspects

Larger order quantities reduce packaging and transport effort per unit of quantity and can thereby support ESG goals. Targeted packaging reduction can often be achieved through optimised ordering rhythms and recurring reusable packaging. For critical materials, a material substitution (ESG) assessment is also worthwhile, since the cost logic of the optimal order quantity can also influence the choice of material.

Supply chain transparency and framework agreements

Stable framework agreements with agreed call-off quantities relieve the classic individual order and enable more flexible call-offs close to the calculated optimum. Instruments such as supply chain mapping and supplier evaluations provide the data basis for adapting order quantity strategies to the actual supplier risk. Legal aspects – such as cancellation rights or the supplier's storage obligations – must be reviewed contractually on a case-by-case basis.

Typical metrics

Inventory coverage: 60–120 days (often higher for C parts)

Benchmark for assessing whether the chosen order quantities are too large

Stock turnover: > 6 turns per year (depending on the industry)

Lower values indicate excessive order quantities

Order frequency per material: within the calculated optimal order quantity

Frequent very small orders signal quantities that are too low

Emergency order rate: < 5% of orders

High values show that order and safety stock levels are too tight

Practical example

A medium-sized mechanical engineering company with around 8,000 stock-kept items frequently ordered standard commercial goods in very small quantities, sometimes several times a week. At the same time, an inventory analysis showed that expensive A components with high coverage levels were being held in stock – precisely the wrong way round. Procurement introduced systematic maintenance of planning parameters: ordering and inventory holding costs were captured realistically, the Andler batch size was calculated for all A and B items, and the values were stored in the ERP system in volume-based tiers. Volume discount tiers were transferred into new framework agreements, C items were moved to consolidated collective orders, and certain standard parts were switched to alternative sourcing.

- Calculation of the optimal order quantity for all A and B items with realistic ordering costs
- Regular inventory analysis to identify excessive coverage levels and dead stock
- Transfer of discount tiers into framework agreements with flexible call-off quantities
- Quarterly review of planning parameters with clear responsibility within procurement

Conclusion

The optimal order quantity is a proven instrument for bringing inventory holding costs and ordering effort into a balanced relationship. Its strength lies in its simplicity; its limitation lies in the assumption of stable demand and costs. In practice, it only unfolds its effect in combination with well-maintained master data,

regular reviews and clear responsibility within procurement. Those who supplement the calculated optima with operational and supplier-side constraints achieve sustainably lower total costs without any loss of delivery capability.

Frequently asked questions

What is the optimal order quantity in simple terms?

The optimal order quantity is the quantity purchased per order at which ordering and inventory holding costs combined are at their lowest. Ordering too frequently increases transaction costs; ordering too much ties up capital in stock and incurs costs. The Andler formula provides a calculated optimum for this.

How do you calculate the optimal order quantity according to Andler?

The Andler formula is: Order quantity = square root of $(2 \times \text{annual requirement} \times \text{ordering costs})$ divided by $(\text{purchase price} \times \text{inventory holding cost rate})$. You therefore need the annual requirement, the fixed cost per order and the inventory holding cost rate. In practice, the result is then adjusted for pack sizes, minimum order quantities and price breaks.

Which costs are included in the optimal order quantity?

On the ordering side, all fixed transaction costs are counted, from the enquiry through to invoice verification, including coordination and goods receipt. On the inventory side, capital costs, warehouse space, insurance, shrinkage and obsolescence are counted, typically aggregated as a percentage of the stock value. Volume discounts and transport costs are factored in through extended models.

How often should you recalculate the optimal order quantity?

An annual review is common in practice for B and C items, and quarterly for A items. Event-driven recalculations make sense in the case of significant price changes, demand jumps or supplier changes. It is important that any changes are properly documented in the material master data.

What is the difference between the optimal order quantity and the optimal batch size?

Both concepts follow the same cost logic and are calculated using similar formulas. For the optimal order quantity, the fixed ordering costs of procurement are decisive; for the optimal batch size, it is the set-up costs of production. The difference therefore lies less in the mathematics than in the cost drivers considered.

What are the disadvantages of the classic Andler formula?

The formula assumes constant demand, constant prices and unlimited storage capacity – assumptions that rarely fully apply in practice. It also does not take into account price breaks, supply restrictions or demand uncertainty. The result should therefore always be understood as a starting point and supplemented with real-world constraints.