Before doing any math, ask: **What are the units?** A graph of dollars per day tells a different story from a graph of total dollars, even if both involve https://activecalculus.org/author/Running-Totals-Rates-of-Change money and days. Units help identify what is being measured, what the graph can tell us, and whether a straight-line connection between data points makes sense.

A **piecewise linear graph** is made of straight-line sections joined at specified points. It can be a useful way to describe or estimate changing quantities, especially when you have measurements at particular times and need to **approximate between points**. But it is still a model simplification: a convenient representation, not a promise that the real-world quantity changed in perfectly straight lines.

This distinction matters in calculus, business dashboards, shipping, inventory, and finance. A line between two observations may suggest a steady change that was never actually observed. To use a piecewise linear graph responsibly, first decide what the recorded quantity means, identify the time intervals involved, and be clear about whether you are graphing a total or a rate.

What Does "Piecewise Linear" Mean?

"Piecewise" means that a graph is described in separate parts. "Linear" means that each part is a straight line. For example, suppose a delivery dashboard records the cumulative number of packages delivered at 9 a.m., noon, and 3 p.m. Connecting the three recorded points with line segments produces a piecewise linear graph.

On each segment, the graph has a constant slope. If $Q(t)$ is the cumulative number of packages delivered by time t, then the slope between two observations is

$$ \fracQ(t_2)-Q(t_1)t_2-t_1. $$

If Q is measured in packages and t in hours, the slope's units are packages per hour. That slope is an **average rate of change** over the interval from t_1 to t_2. It is not necessarily the delivery rate at every moment within that interval.

A piecewise linear model can also be written as a function that uses a different formula on each interval. For instance, if a quantity rises from 12 units at time $t=0$ to 20 units at $t=4$, a linear model for that interval is

$$ Q(t)=12+2t,\qquad 0\le t\le 4. $$

Here, t is measured in hours and $Q(t)$ in units. The slope is 2 units per hour. If the next interval has a different slope, the model changes formula at $t=4$—hence "piecewise."

Start With What Was Recorded

Before drawing a line or calculating a slope, define the recorded quantity in one clean sentence. For example: **"The table records the cumulative number of packages delivered since the shift began at 8 a.m."** That sentence answers two important questions: What is being measured, and what is the starting time?

The starting time matters because "total" is not meaningful without a reference point. A cumulative count of packages delivered since 8 a.m. is different from the number delivered during the most recent hour. Likewise, an account balance at noon is a total balance at a particular time; it is not the amount earned per hour.

For a dashboard, it may help to write down:

- **Recorded quantity:** What does each vertical-axis value represent?

- **Start time or reference point:** Since when has the total been accumulated, or relative to what baseline is it measured?
- **Time unit:** Are horizontal-axis values in minutes, hours, days, or another unit?
- **Measurement schedule:** At what times were observations actually recorded?

Label both axes with quantities and units. “Time” and “Amount” are often too vague. “Time since shift start (hours)” and “Packages delivered (packages)” are much more useful. A well-labeled graph makes it harder to confuse a total with a rate.

Totals and Rates Are Different Questions

A common source of confusion is treating a total as if it were a rate. They are connected, but they answer different questions. A total answers, “How much has accumulated by this time?” A rate answers, “How quickly is the amount changing?”

Suppose a tank contains 40 liters at 10 a.m. and 64 liters at noon. The recorded quantity is the volume in the tank, measured in liters. The average rate of change over those two hours is

\[\frac64-40\text liters2\text hours =12\text liters per hour. \]

The tank’s volume is not 12 liters per hour; 12 liters per hour is the average rate at which the volume increased over that interval. The distinction is the same in inventory: “items in stock” is a total or level, while “items sold per day” is a rate. In finance, “account balance” is an amount of money, while “dollars earned per month” is a rate of change.

A graph of a total can be piecewise linear even when the underlying events happen in separate batches. But the line’s slope then represents an interval average, not a claim that every fraction of a package, item, or deposit occurred smoothly through time.

Intervals Matter for Average Rate of Change

An average rate belongs to a specific interval. If inventory falls from 500 items on Monday to 380 items on Friday, you can calculate an average decrease per day only after defining the time interval and counting its duration correctly. From Monday to Friday is four elapsed days—not five—if the measurements are taken at the same time on each day.

In general, for a quantity Q recorded at times t_1 and t_2, the average rate is

\[\textaverage rate of change =\fracQ(t_2)-Q(t_1)t_2-t_1. \]

The numerator has the units of the recorded quantity, and the denominator has units of time. The result therefore has units of “quantity per time.” If a shipment total changes by 90 boxes over 3 hours, the average rate is 30 boxes per hour. It does not follow that the shipment moved at exactly 30 boxes per hour throughout those three hours.

Choosing a different interval can give a different average rate. A rate averaged over one hour may hide a rush of activity followed by a pause; a rate averaged over a full day may smooth over both. So always say what the interval is when reporting an average rate. “The average was 30 boxes per hour from noon to 3 p.m.” is more informative than simply saying “The rate was 30 boxes per hour.”



When Is Linear Interpolation Okay?

Linear interpolation estimates a value between two known data points by using the straight-line segment between them. If you know a quantity at the beginning and end of an interval, interpolation can estimate a value at a time in between.

For example, suppose a dashboard shows a cumulative total of 100 items at 1 p.m. and 160 items at 3 p.m. If a rough estimate is needed for 2 p.m., linear interpolation places 2 p.m. halfway through the interval and estimates 130 items. In general, between (t_1,Q_1) and (t_2,Q_2), the interpolated value at t is

$$ Q(t)=Q_1+\fract-t_1t_2-t_1(Q_2-Q_1). $$

This estimate is reasonable when the quantity is expected to change at a fairly steady pace over the interval, the observations are close enough together for the purpose, and an estimate—not an exact reading—is acceptable. It is often useful for planning, quick comparisons, or filling in a gap in a report.

Interpolation is less suitable when there are known surges, pauses, scheduled batches, cutoffs, or other events that make the change uneven. Imagine that all 60 items were processed in a single batch at 2:50 p.m. A straight line from 1 p.m. to 3 p.m. would imply gradual processing throughout the interval. That implication would be false, even though the line connects the correct endpoint values.

Interpolation Is Not Extrapolation

Interpolation estimates *between* observed points. Extrapolation predicts beyond the observed range. Extending a line past the final recorded time assumes that the same pattern will continue. That can be a much riskier assumption, especially when schedules, capacity, demand, or other conditions may change.

A piecewise linear graph is generally most defensible within the intervals it was built to represent. For values outside those intervals, label the prediction as an extrapolation and explain the assumption behind it.

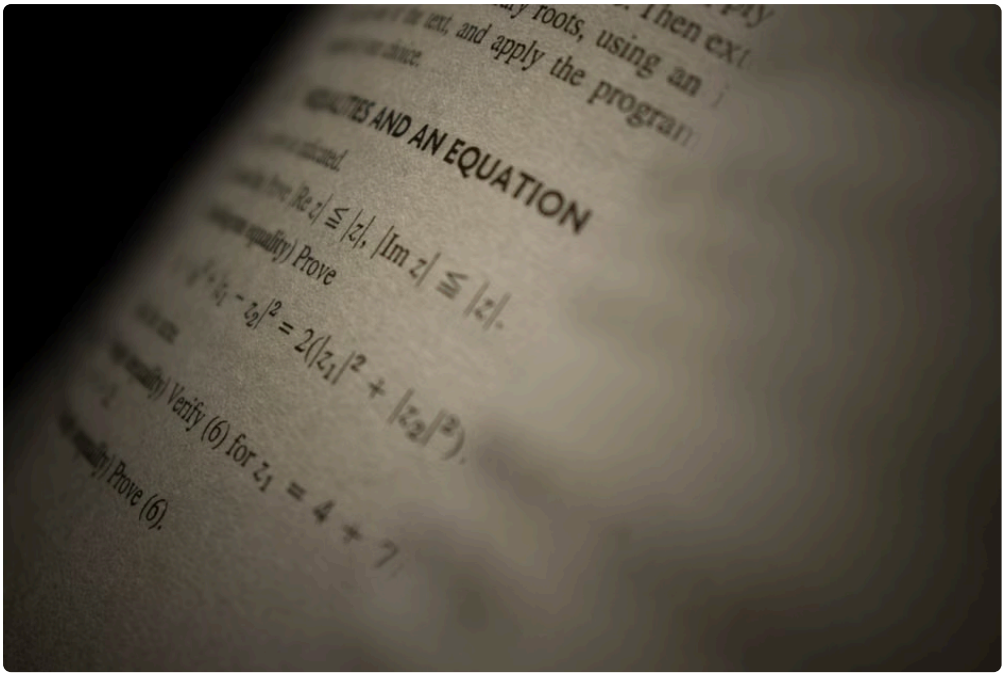
When a Straight Line Misleads

A straight-line connection is a claim about what happens between observations. It may imply continuous, steady change even when the process actually changes in jumps. When events happen in jumps, I draw a **step graph** on

purpose: horizontal stretches show that a recorded total stays the same, and vertical jumps show when it changes. The jump's exact timing should only be shown if the data support it.

Consider a warehouse count measured once each morning. If the count drops from 500 items to 420 items between Tuesday and Wednesday, the data tell us the two recorded totals. They may not tell us whether items left in one shipment, in several orders, or gradually throughout the day. A diagonal line invents a particular pattern between the measurements.

Other situations where straight-line connections need caution include:



- **Batch events:** deliveries, deposits, or production runs may occur all at once.
- **Scheduled changes:** rates may rise or fall at opening time, shift changes, or billing cutoffs.
- **Accumulated counts:** a count may remain constant for a while, then jump when an event is recorded.
- **Coarse measurements:** widely spaced observations can hide meaningful changes between timestamps.
- **Changing conditions:** a trend may not continue when staffing, demand, or capacity changes.

A useful model does not need to reproduce every detail. But it should not suggest details that contradict what is known. When the data do not show what happened between times, say that the connecting line is an approximation—or use a graph that better matches how the process works.

Piecewise Linear Graphs as Model Simplification

Real-world behavior is often more complicated than a few straight lines. A piecewise linear graph is a form of model simplification: it reduces a complicated pattern to manageable sections with constant slopes. That simplification can make a graph easier to read, calculate with, and explain.. Pretty simple.

For example, a finance model might use one rate of change for a period before a fee and another rate afterward. A shipping dashboard might use separate average rates for morning and afternoon. In each case, the change points should have a real reason—such as a policy change, a shift boundary, or a scheduled event—not just be added to make the graph look precise.

Ask whether the simplified model is accurate enough for the decision at hand. A rough estimate may be fine for a high-level planning conversation. It may not be fine for billing, auditing, safety decisions, or evaluating an

individual worker's activity. The appropriate level of detail depends on how the graph will be used.

A Practical Checklist

Before relying on a piecewise linear graph, check the following:

1. **Define the quantity.** State exactly what each vertical value records.
2. **Name the reference point.** If the quantity is cumulative, specify when accumulation begins.
3. **Write the units.** Include units on the axes and in any rate you report.
4. **Identify the interval.** A slope or average rate applies to a particular start and end time.
5. **Choose the connection deliberately.** Use straight segments when a steady-change approximation is reasonable; consider a step graph when events happen in jumps.
6. **Mark estimates as estimates.** Distinguish recorded values from interpolated or predicted ones.
7. **Check the decision.** Decide whether the model's assumptions are suitable for the purpose.

Graph feature What it tells you What it does not automatically tell you Recorded point The measured value at a specified time What happened between measurements Line segment A linear approximation between two points That change was actually steady Slope Average rate of change over an interval The exact rate at every moment Step graph A quantity that stays level between changes, with jumps The jump's exact timing unless recorded

Using MathJax and Runestone to Explore the Idea

MathJax can render the equations in this discussion clearly in a web page, including the average-rate formula and the interpolation formula. That is useful when readers need to see the difference between a total, such as \(Q(t)\), and a rate, such as \(\Delta Q/\Delta t\). A well-rendered equation still needs a plain-language explanation and units.

Runestone is an interactive publishing platform that can support lessons where learners inspect tables, adjust graph points, and compare straight-line and step-graph representations. An activity might ask learners to move a timestamp, recalculate an interval's average rate, and explain why a different interval gives a different result. The key learning goal is not just producing a number; it is interpreting what the number represents.

The Bottom Line

A piecewise linear graph connects measured points with straight-line sections. It is useful when you need a simple model or want to approximate between points, and when a roughly steady change is a reasonable assumption for each interval. Its slopes represent average rates of change over those intervals—not automatically instantaneous rates.

Start by defining the recorded quantity and its start time. Ask what the units are. Separate totals from rates, name the intervals, and remember that a diagonal line can imply smooth activity that never occurred. When the real process changes in jumps, a step graph may tell the more honest story. The best graph is not the most polished one; it is the one whose assumptions match the data and whose meaning readers can understand.