

## Revenue operations &amp; AI

# What is Little's law?

Little's law,  $L = \lambda W$ , ties three averages: items in a system, the rate they arrive and the time each spends inside. It measures and does not explain.

| Author                                                                                       | Published       | Reading time | Words |
|----------------------------------------------------------------------------------------------|-----------------|--------------|-------|
| Sinan Isoglu, MBA (Quantic)<br>Commercial growth leader, lecturer and<br>doctoral researcher | 18 October 2026 | 18 min       | 3,852 |

**Management summary** Little's law states that the average number of items in a system equals the average rate at which they arrive multiplied by the average time each spends inside:  $L = \lambda W$ . This concept guide gives the conditions under which Little proved it in 1961 (finite means, strictly stationary processes, an arrival process that is metrically transitive with a nonzero mean) and the finite-period versions in his 2011 review, which hold exactly when all three averages are taken, as he defines them, over the same observed period. It explains why the equality holds, separates what the law states from what it cannot say about causes or forecasts, and shows how a revenue team could read one pipeline stage with it: deals in the stage, deals entering per week and weeks in the stage. The sources are a mathematical proof and a review of theory and practice; the sales reading is the author's and untested.

|          |          |                                                         |    |
|----------|----------|---------------------------------------------------------|----|
| Contents | § 1      | What is Little's law?                                   | 2  |
|          | Figure 1 | One law in six vocabularies                             | 4  |
|          | § 2      | Why does counting one area twice make it true?          | 5  |
|          | § 3      | Under what conditions did Little prove it in 1961?      | 6  |
|          | § 4      | Does the law hold over one quarter of data?             | 7  |
|          | § 5      | What does the law not tell a manager?                   | 8  |
|          | § 6      | How could a revenue team read a pipeline stage with it? | 9  |
|          | Figure 2 | A Little's law check for one pipeline stage             | 11 |
|          | § 7      | Which pipeline questions are beyond these two papers?   | 12 |
|          |          | References                                              | 13 |

## What is Little's law?

§ 1 of 7

Little's<sup>21</sup> law says that in a system which items enter, stay in for a while and leave, the average number of items inside,  $L$ , equals the average rate at which they arrive,  $\lambda$  (lambda), multiplied by the average time each one spends inside,  $W$ :  $L = \lambda W$ . That is how Little (2011) states it fifty years after his proof, and his verdict on its status is short: "Basically, it is a theorem in queuing theory." His examples of items are cars at a toll booth, people in a cafeteria line, aircraft on a production line and instructions waiting to be executed inside a computer. The time inside is the whole stay in the queuing system: where service is what creates the queue, it includes the service as well as the wait.

<sup>1</sup> Little, J. D. C. (1961). A proof for the queuing formula:  $L = \lambda W$ . Operations Research, 9(3), 383-387.

<sup>2</sup> Little, J. D. C. (2011). Little's law as viewed on its 50th anniversary. Operations Research, 59(3), 536-549.

The 1961 paper defines the three quantities as means of a queuing process:  $L$  is "the mean number of units in the system",  $W$  is "the mean time spent by a unit in the system", and  $1/\lambda$  is "the mean time between the arrivals of two consecutive units", so  $\lambda$  is the arrival rate. Others had argued for the formula heuristically and proved it for particular models, one of them with Poisson arrivals at a facility serving first come, first served; Little proved it under considerably more general assumptions. He also names why anyone wants it: it is sometimes easier to find  $L$  than  $W$ , or the other way round, when solving a queuing model, and the law gives the one from the other.

The three are different kinds of average, and Little (2011)<sup>2</sup> is explicit about it.  $L$  "is a time average of the number of items in the system": the count inside, averaged over the period, a pure number of items that Little calls dimensionless.  $W$  is an average over items, the mean of their individual times, so it is measured in time units. And  $\lambda$  "is an average rate calculated by counting the" arrivals in the period and dividing by its length, so it has the "dimensions of items per unit time".

In this guide, Little's<sup>21</sup> law means the statement that, for a boundary drawn consistently around a flow of items, the average number inside equals the rate at which items enter multiplied by the average time each spends inside, with all three averages taken over the same period or, in the long-run version, over the same process. It is a relation that measurements of a flow must satisfy under the conditions set out below, not a model of what produces those measurements. Operations

management writes the same relation as throughput equals work in process divided by cycle time, and Figure 1 shows the three averages under the names that Little (2011) reports from several fields.

**Figure 1** One law in six vocabularies

Each field renames the three averages, and the relation between them stays the same. The sales pipeline row is this guide's reading.

| Setting                                     | L, number in the system                   | Lambda, rate                                                    | W, time in the system             | Where in Little (2011) |
|---------------------------------------------|-------------------------------------------|-----------------------------------------------------------------|-----------------------------------|------------------------|
| Queuing theory                              | Average number of items in the system     | Average arrival rate of items                                   | Average wait in the system        | p. 545                 |
| Factory (Hopp and Spearman)                 | Work in process (WIP)                     | Throughput (TH), an output rate                                 | Cycle time (CT)                   | pp. 545 and 547        |
| Any flow unit (Cachon and Terwiesch)        | Average inventory                         | Average flow rate                                               | Average flow time                 | p. 545                 |
| Emergency department (Harris case)          | Average patients in process (PIP)         | Average arrival rate of patients                                | Average length of stay (LOS)      | p. 546                 |
| Computer server                             | Average number of requests in process     | Average arrival rate of requests                                | Average response time per request | p. 542                 |
| Sales pipeline stage (this guide, untested) | Average number of open deals in the stage | Deals entering the stage per week, with those open at the start | Average weeks in the stage        | not in the sources     |

Operations management states the law with the output rate rather than the arrival rate. The first five rows are reported by Little (2011); the sales pipeline row is this guide's.

**Source** Little (2011), Operations Research 59(3): the queuing theory notation and the operations management notations of Hopp and Spearman and of Cachon and Terwiesch, p. 545, with Hopp and Spearman's definitions on p. 547; the emergency department terms in which Little restates the law for the Harris case, p. 546; the server terms, p. 542. The sales pipeline row is this guide's reading and is not in either source.

Little (2011)<sup>2</sup> quotes the factory definitions from Hopp and Spearman's textbook: throughput is "the average output of a production process" per unit of time, work in process is "the inventory between the start and end points" of a product routing, and cycle time is the average "time from release of a job at the beginning of the routing" until the job reaches the end. He notes that operations management states the law with the output rate "rather than the arrival rate", and that operations management applications draw their vocabulary from the "jargon of the client setting".

## Why does counting one area twice make it true? § 2 of 7

Little (2011)<sup>2</sup> proves the simplest version in a few lines. Take a system observed over a period of length  $T$  that is empty at the start and at the end, and draw the number of items inside it as a curve over the period. Call the area under the curve  $A$ , measured in item-hours or item-weeks. Dividing  $A$  by  $T$  gives  $L$ , the average number inside. Dividing  $A$  by  $N$ , the number of items that arrived, gives  $W$ , the average time per item.  $N$  divided by  $T$  is  $\lambda$ . So  $L = A / T = (N / T) (A / N) = \lambda W$ . The proof uses the same area twice, once for  $L$  and once for  $W$ , and Little traces that to a "simple physical fact": "An item in queue is also waiting." While a customer stands in line and can be counted, the customer is also accumulating minutes of waiting. In his concluding thoughts he calls this simultaneity the "glue that creates the equality".

The 1961 proof runs on the same picture with a complication. Over an arbitrary stretch of time, the area under the curve equals the sum of the times spent inside by the units that arrived in the stretch, "except for certain carry-over effects at the ends of the interval": units already inside when the stretch begins, and units still inside when it ends. Looking back in 2011, Little recalls that he could see "end effects in the picture" and needed a way to "get rid of them in the limit", and that the picture seemed to need little more than a count going up and down by one item at a time and "some guarantee of steady state" and of conservation of items.

## Under what conditions did Little prove it in 1961? § 3 of 7

Little (1961)<sup>1</sup> works with a queuing process running over all time that generates three random processes: the number of units in the system at each moment, the time each arriving unit spends in the system, and the time between one arrival and the next. Part of the definition is a bookkeeping rule: the number in the system at a moment is the number of units “whose time of arrival is before (or equal to)” that moment and whose time of departure is after or equal to it. The theorem then needs two conditions. First, the three processes are strictly stationary, each with a finite mean; in the abstract, the law holds “if the three means are finite and the corresponding stochastic processes strictly stationary”. A stationary process, as the 2011 review defines it, is one whose joint probability distribution does not change with time. Second, the arrival process is “metrically transitive with nonzero mean”: the mean time between arrivals is finite and above zero, and, in this guide’s gloss, the long-run average time between arrivals comes out at that same mean on almost every realisation of the process.

Under those conditions Little proves two results. The first is about one realisation: with probability one, the long-run averages of the number in the system, the time per unit and the time between arrivals exist, are finite and satisfy the law, so an equivalent of the law “using time averages holds with probability one for any specific realization of the queuing process”. The second is about expected values: the law is a “valid relation among phase averages”, which Little calls the principal result for applications.

What the theorem does not need is as telling. The results are “remarkably free of specific assumptions about arrival and service distributions, independence of interarrival times, number of channels, queue discipline”, and the rest. Little flags the strong assumption himself: “A requirement is made for strict stationarity”, which is “probably not the weakest requirement possible”, though he judged that “the steady state in most current queuing models would appear to be strictly stationary” and that, in cases of practical interest, the arrival process is “likely to be metrically transitive”. Later work removed the probability structure. In the sample path version that Little (2011)<sup>2</sup> reports, it is enough that the path comes from some sample space: “If limits  $\lambda$  and  $W$  exist and are finite,  $L$  exists” and is finite, and  $L = \lambda W$ . He credits Stidham’s sample path proof of 1972 as the first fully rigorous proof in the generality known today, and he describes queuing theory as “primarily concerned with relationships among” long-run averages.

Little (1961)<sup>1</sup> also leaves the boundary of the system to the user. The theorem “only requires consistency of meaning in the phrases” for the number of units in the system, the time spent in the system and arrival to the system. Call the queue alone the system, and the law holds for the queue: the mean number waiting equals the arrival rate times the mean wait. It holds for each priority class on its own, with that class’s own arrival rate. And it fails in a case Little names, models in which “not all arrivals join the system” while the rate counts them all. It holds again if the rate counts “only those arrivals that join the system”, or if “the units that do not join have a zero waiting time in the system” and are included in  $W$ .

## Does the law hold over one quarter of data?

§ 4 of 7

It does, and exactly, if the three averages are defined the way Little (2011)<sup>2</sup> defines them for a finite period. He wrote those versions because proofs over an infinite time axis leave out what practitioners face: they “operate over a finite time period”, and “all observations of practice take place in a finite” time interval, so the numbers obtained there need an interpretation. His first finite theorem, LL.1, covers a system observed from time 0 to time  $T$ , with  $T$  finite, “that is empty at 0 and  $T$ ”, such as a supermarket from opening to closing. Then “the formula  $L = \lambda W$  holds”, with  $L$  the area divided by  $T$ ,  $\lambda$  the number of arrivals divided by  $T$  and  $W$  the area divided by the number of arrivals. Three remarks go with it. “LL.1 holds under nonstationary conditions”: the supermarket’s customers arrive at a rate that may peak a little after noon and taper off at the end of the day, and nothing in the proof assumes otherwise. “LL.1 holds independent of queue discipline”, whether items are served first in, first out, last in, first out, at random or by priority. And “ $L = \lambda W$  is deterministic and holds exactly”, but only after the fact, once the path of the system is known: in Little’s<sup>1</sup> phrase, the user is in the measurement business, “not the forecasting business”. Collecting many days of such data and forecasting from them with standard statistical methods remains possible, as a separate step.

Many systems are not empty at both ends of a period: a wine cellar ordinarily is not, and neither is a sales pipeline. Little’s<sup>21</sup> second theorem, LL.2, headed “Little’s Law Over  $[0, T]$ ”, drops that restriction. Items already inside at the start are counted with the arrivals: the count includes not only arrivals “but also any items that were in the” system at time 0, and  $\lambda$  is that count divided by  $T$ .  $W$  is the area divided by the same count, so it is the average time spent inside during the period, not the average total stay; for the wine cellar it measures the

“average incremental age being added” to the wines over the period. With those definitions, “LL.2 holds exactly even if the queuing system is” never empty, and it “holds even if there is no service operation”. Little adds a warning. The algebra is right, but the items “may have a history of waiting prior” to the start and may stay after the end, so whether this version answers “the questions we want to ask” depends on the task, and a “practitioner should always check his or her model” assumptions for relevance. The finite versions also hold class by class, which Little’s heading names “Classes of items or market segments”: the law holds within each class, the class counts and arrival rates add up to the whole, and the overall  $W$  is the arrival-weighted mean of the class times.

## What does the law not tell a manager?

§ 5 of 7

Why the averages are what they are, or what to change. Little (2011)<sup>2</sup> is plain on both. The law ties three quite different measures of a system’s performance together, and “LL will not tell the managers how” to handle trade-offs between them or how to improve the measure they care about; “it lays down a necessary relation”. Of the lean manufacturing practice he reports, he writes: “Note that Little’s<sup>1</sup> Law is not in charge of changing these” quantities; the improvement work is. The activities of the organisation that runs a system are “separate from Little’s Law but are constrained by it”.

What produces the averages needs other models. In a server load test Little reports, the average queue rose roughly in line with the request rate and then climbed steeply (Little obtained that queue from the measured response times through the law itself); his interpretation is queue-related overhead, the extra work of keeping track of large queues, which “is not a consequence of” the law “but is a second-order effect requiring a new submodel”. For emergency departments, Little explains why staffing to the minimum is not enough by “basic queuing theory”, not by the law: queuing models such as the  $M/M/1$  queue predict that, “as the arrival rate approaches the service rate”, patients in process and length of stay rise “slowly at first and then rapidly”. Of the server case he makes the general point: it is characteristic of applications of the law in practice “that it is used in conjunction with other models and other” measurements.

Nor does the law say which of the three is set and which follows. If the arrival rate is fixed, the “only way to reduce inventory ( $L$ ) is to find a way to reduce” the average time per item ( $W$ ), which is arithmetic, not a mechanism. Which quantity is fixed depends on the operation: in many manufacturing settings the output rate is given and the manager

decides when to release work, so that “In effect, the arrival process is endogenous”. If something is wrong, Little expects it to show up in “one or more of the three measures”, but, in his words, “We cannot, a priori, tell the manager how to fix a problem”.

On this guide’s reading, that makes Little’s<sup>21</sup> law an identity between averages, not a causal model. Over a finite period it is arithmetic on one set of records; over the long run it is a theorem about averages that exist. If one of the three changes between two periods, at least one of the others has changed with it, but the law cannot say which change caused which, or what intervention would move any of them.

## How could a revenue team read a pipeline stage with it? § 6 of 7

Neither paper discusses sales pipelines. Little (2011)<sup>2</sup> points the law beyond the factory: he argues that managers should consider collecting and displaying the three measures for any stream of items they manage, names a pharmaceutical firm’s “new product development program” as one such stream, and recalls looking for “ways to use Little’s<sup>1</sup> Law in marketing” after a chapter he wrote with Graves. What follows is this guide’s example. The numbers are invented and the reading has not been tested on sales data.

Take one stage of a pipeline, discovery for example, as the system. Its boundary is the event that puts a deal into the stage and every event that takes one out: a deal that moves to the next stage, moves back to an earlier one, is won, is lost or is disqualified leaves the system on that date, and a deal moved back into the stage counts as a new arrival. A deal is an item. Over a 13-week quarter,  $L$  is the average number of deals open in the stage, the deal-weeks in the stage (the weeks the deals spent there, added up) divided by 13;  $\lambda$  is the number of deals that were in the stage at any point in the quarter, those open on the first day plus those that entered, divided by 13; and  $W$ , the average time in the stage, is the same deal-weeks divided by that number of deals. Suppose the stage held 120 deals on the first day, 270 entered during the quarter, and the weeks those 390 deals spent in the stage inside the quarter add up to 1,560. Then  $L = 1,560 / 13 = 120$  deals,  $\lambda = 390 / 13 = 30$  deals per week and  $W = 1,560 / 390 = 4$  weeks, and  $30 \times 4 = 120$ .

Three readings follow from the identity alone. First, any two give the third, which Little calls “perhaps the most common and also the most valuable use of LL”: records of each deal’s entry and exit dates give all three. Without the dates, the weekly counts of open deals, averaged

over the quarter, give  $L$  approximately, since a count taken once a week only approximates the area, and divided by  $\lambda$ , the deals open on the first day plus the entries, per week, they give  $W$ . Divided by the entries alone they give a different number, which is not the  $W$  of the finite-period version (on the numbers above, 5.8 weeks instead of 4). Second, the check bites only where the three numbers come from different places. Computed from one set of entry and exit dates, they agree by construction; what can disagree is an average from the dates and a figure the CRM reports on its own, such as its weekly count of open deals or its average time in stage, and a disagreement beyond the small gap that counting once a week leaves is a finding about the records or the definitions. Lovejoy, whose hospital study Little reports, put it as a check on raw data that many people gather in a decentralised way, each perhaps with definitions of their own: “Little’s<sup>21</sup> Law provides a reality check.” On a pipeline, likely causes would be entry dates rewritten when a deal is moved back, so that the dates lose a spell the weekly counts saw, or a CRM time in stage taken only from the deals that left during the quarter, which is not the  $W$  of the finite-period version because it leaves out the deals still open at the end and counts weeks spent before the quarter began. Lost deals left open do not show up as a disagreement: they sit in every count alike and make the stage look slower, which is why they return below among the reasons a stage can slow. Third, if  $\lambda$  stays at 30 deals a week and the stage’s average count of open deals grows from 120 to 150, the average time in the stage has risen from 4 weeks to 5: at the same inflow, a fuller stage is a slower stage, by arithmetic.

What the identity cannot say is why. A stage that has slowed may hold buyers who take longer to decide, sellers with too many deals to work, or lost deals that nobody has closed, and the law fits all three equally. Nor does it give a win rate, a deal value or next quarter’s numbers. Which records count as open at all is the question that [pipeline hygiene](#) answers, with the tests a record must pass before it enters an open-pipeline view. Because  $W$  depends on where a stage starts and stops, [the sales-cycle essay](#) shows how a cycle number changes when its start event, stop event and outcomes are defined differently. And a stage total in money is another object again: [pipeline coverage](#) explains why a coverage multiple needs its stage mix and probability assumptions beside it.

### Figure 2 A Little's law check for one pipeline stage

Fix the boundary and the window, count the deals and the deal-weeks, and compare the three averages with the figures the CRM reports.

| Stage                             | Window                           | Deals                         | Deal-weeks                            | Averages                                 | Records                              |
|-----------------------------------|----------------------------------|-------------------------------|---------------------------------------|------------------------------------------|--------------------------------------|
| Entry event and every exit event? | First day, last day, T in weeks? | Open on day one plus entries? | Weeks in the stage inside the window? | Deal-weeks/T, deals/T, deal-weeks/deals? | Same as the figures the CRM reports? |
|                                   |                                  |                               |                                       |                                          |                                      |
|                                   |                                  |                               |                                       |                                          |                                      |
|                                   |                                  |                               |                                       |                                          |                                      |
|                                   |                                  |                               |                                       |                                          |                                      |

The three averages agree by construction. A gap with the CRM figures points to the records or the definitions, not to the law.

One row per stage and window. Count deals open on the first day with the entries, and count only the weeks spent inside the window.

**Source** Author's decision worksheet. The definitions of the three averages over a period that need not start or end empty follow Little (2011), pp. 537 to 539; the check against the records follows the reality check reported on p. 547. Applying them to a sales pipeline stage is this guide's reading and has not been tested. The fields are reader inputs; no company data is supplied.

The worksheet keeps the three choices that decide the numbers in view: where the stage starts and ends, which period is counted, and whether deals open at the start of the period are counted with the entries. The last column is the reality check: if the averages computed from deal dates do not match the figures the CRM shows, the gap is a question about the records or the definitions, not about the law.

## Which pipeline questions are beyond these two papers? § 7 of 7

Most of the practical ones. Little (1961)<sup>1</sup> is a mathematical proof that reports no data. Little (2011)<sup>2</sup> is a review with two short proofs of its own: it proves the finite-period versions, summarises what others proved about the long-run versions, and in its practice sections reports cases, most of them from people the author “personally communicated with” (emergency department staffing, a hospital’s observation unit, a lean consulting practice), and a server benchmark. They show how the law is used, not tests of it, and as a theorem it needs none. Neither paper treats sales pipelines, stage definitions or how long deals spend in a stage.

On this guide’s reading, the finite-period identity carries over to a pipeline exactly, because it is arithmetic on the records of one period, provided the boundary and the counts are consistent. The long-run version does not carry over without more: a pipeline whose inflow grows, or swings with the fiscal quarter, may have no long-run averages to settle on, and strict stationarity is a strong assumption for it. A quarter’s  $L$ ,  $\lambda$  and  $W$  describe that quarter. Using them for the next one is a forecast with assumptions of its own, and how the three respond to more sellers, fewer stages or faster follow-up is a question neither source tests.

 End of essay

---

---

**References**

- 1 Little, J. D. C. (1961). A proof for the queuing formula:  $L = \lambda W$ . Operations Research, 9(3), 383–387. <https://doi.org/10.1287/opre.9.3.383>
- 2 Little, J. D. C. (2011). Little's law as viewed on its 50th anniversary. Operations Research, 59(3), 536–549. <https://doi.org/10.1287/opre.1110.0940>

---

**How to cite this essay**

Isoglu, S. (2026, October 18). What is Little's law? Sinan Isoglu.  
<https://isoglu.com/insights/journal/what-is-littles-law/>  
Online version: <https://isoglu.com/insights/journal/what-is-littles-law/>

---

**Sinan Isoglu, MBA (Quantic)**

Commercial growth leader, lecturer and doctoral researcher  
© 2026 Sinan Isoglu · [isoglu.com](https://isoglu.com) · All rights reserved