Setting the Asking Rent of a Vacant Long-Term Rental

Days on Market, the Value of a Lease, and Learning Demand During the Vacancy

Mathematics & Physical Sciences Series

***Abstract*—A landlord with a vacant unit sets one number, the asking rent, and lives with it for the length of the lease. Ask too much and the unit sits empty; ask too little and the shortfall repeats every month of a tenancy that lasts about 27 months. This tutorial values the decision as the present value of the vacancy, the lease and every tenancy after it. In a simulated market where qualified applicants arrive at 0.30 a week at the market rent, and 120 dollars more divides that rate by e, the best ask has a closed form. Without side effects it is $62 above the market rent. Tenants who leave sooner when they pay more, and defaults that rise with the rent, move it to $22 below. Asking $100 too much costs $632 a year; asking $100 too little costs $355. The comparables give the market rent to within about $120, and a plain regression on them puts the price sensitivity at $392 against a true $120. One landlord's history is far too thin to do better. The weekly inquiries during the vacancy are the most useful data. In a worked example, a Bayesian update from the weekly count cuts the uncertainty about the market rent from $119 to $51 in four weeks. A markdown policy built on it loses $536 per vacancy against an oracle that knows the market, where pricing at the comparables loses $1,541 and a common cut-every-two-weeks rule loses $1,181.**

***Index Terms*—Rental pricing, days on market, renewal-reward processes, optimal stopping, markdown pricing, Bayesian updating, Gamma-Poisson model, comparables, rent regulation.**

Contents

I Introduction 1
I-A How this differs from a nightly rate . . . 1
I-B The running example . . . 2
I-C The program . . . 2
I-D How to read this tutorial . . . 2

II The Value of a Vacancy 2
II-A The cycle . . . 2
II-B Present value . . . 2
II-C Why one month matters . . . 2

III Demand During a Vacancy 2
III-A The applicant rate . . . 2
III-B Days on market . . . 3
III-C Inquiries . . . 3

IV The Best Asking Rent 3
IV-A The closed form . . . 3
IV-B The full model . . . 3
IV-C A rule of thumb for price cuts . . . 3
IV-D Accepting less than the ask . . . 4

V Second-Order Effects 4
V-A Tenant quality and default . . . 4
V-B Tenure . . . 4
V-C Season and lease length . . . 4
V-D Concessions . . . 4
V-E Rent caps . . . 4

VI Estimating Demand from Comparables 5
VI-A What each source measures . . . 5
VI-B The market rent . . . 5
VI-C The price sensitivity . . . 5
VI-D How much an error in θ costs . . . 5

VII Why One Landlord's History Is Too Thin 5
VII-A Past vacancies . . . 6
VII-B Price changes within a vacancy . . . 6
VII-C A workable estimate . . . 6

VIII Learning Demand During the Vacancy 6
VIII-A The funnel . . . 6
VIII-B The update . . . 6
VIII-C A worked example . . . 6

IX When to Cut, and by How Much 7
IX-A Why a known market needs no cuts . . . 7
IX-B The markdown policy . . . 7
IX-C The cost of starting high . . . 7

X Comparing the Policies 7
X-A The setup . . . 7
X-B The results . . . 8
X-C How big is this . . . 8

XI A Routine for One Vacancy 8

XII What the Model Leaves Out 8

XIII Conclusion 9

References 9

A tutorial in the *Mathematics & Physical Sciences* series, and a companion to *Pricing a Short-Term Rental* [1], which sets a nightly rate. This volume sets the monthly rent of a unit leased for a year or more. Every number in it comes from a simulated market in `examples/ltrprice.py` beside this document, which the build runs; a test checks the claims the text makes about its output. The market is simulated so that the true demand is known and each method can be checked against it. No real unit's data is used. This is a modelling tutorial. It gives no financial or legal advice. Its account of rent law is a short summary for orientation.

Rev. `0f8a1ceb3`, built 2026-10-30 10:16 UTC.

I. Introduction

When a tenant moves out, the landlord has to pick one number: the rent to advertise. The choice is made once per vacancy, and it holds for the whole lease, often longer. If the ask is too high, the unit sits empty while the mortgage, taxes and utilities keep running. If it is too low, the unit leases in days, and the shortfall repeats every month of the tenancy.

Most landlords set the ask from comparable listings and adjust it by feel. A common rule is to start a little above the comparables and cut the price if nobody applies within two weeks. This tutorial asks what the best ask is, how to estimate the demand it depends on, and how to change the ask as a vacancy goes on.

A. How this differs from a nightly rate

A companion volume sets the nightly rate of a short-term rental [1]. There a host prices each night separately, changes the price every day, and sees hundreds of bookings a year. The long-term problem differs in four ways, and this tutorial is built around them.

- **One decision, locked in.** The ask becomes the rent of a lease of a year or more. An error is not corrected tomorrow.
- **The value is a whole lease.** A lease is worth its rent over the expected tenancy, minus the vacancy before it and the turnover costs after it.
- **Very little data.** A unit comes vacant every two or three years. One landlord's history holds a handful of vacancies.
- **A short window to learn.** During a vacancy the landlord sees how many people inquire each week. That count arrives every week while the ask can still be changed.

B. The running example

The running example is a fictional two-bedroom apartment in a mid-sized city, called *the apartment*. Its market is simulated, so the true demand is known and every method can be checked against it. The comparables put its market rent at $2,400 a month. At that rent, in an average week of the year, the apartment attracts 0.30 qualified applicants a week, so it leases in about 27 days. A tenant at the market rent stays 27 months on average. A move-out costs $3,000 in cleaning, paint, repairs and leasing fees. These numbers were chosen to look plausible. They describe no real unit.

The comparables are not exact. The apartment's true market rent may differ from $2,400 by about $120 either way, one standard deviation, because no comparable is the same unit. The landlord has to narrow that range during the vacancy.

C. The program

A program beside this document, `examples/ltrprice.py`, simulates the market, solves the landlord's problem and runs every policy. It uses only the Python standard library and fixed random seeds, so a rerun gives the same numbers. The build runs it and reads its numbers into the text, and a test checks each claim the text makes about them.

D. How to read this tutorial

Section II values a vacancy and the lease that ends it. Section III models how fast applicants arrive at each ask, and Sec. IV derives the best ask. Section V adds tenant quality, tenure, the season, concessions and rent caps. Sections VI and VII are about estimating demand before the vacancy, from comparables and from the landlord's own history. Sections VIII and IX learn demand during the vacancy and turn it into a schedule of price cuts. Section X compares the policies, Sec. XI gives a routine for one vacancy, and Sec. XII lists what the model leaves out.

II. The Value of a Vacancy

A vacancy ends with a lease, and the lease ends with the next vacancy. To compare two asks, the landlord has to value the whole cycle, from the first month's rent to the next vacancy. This section sets up that valuation.

A. The cycle

Figure 1 shows one cycle. The unit is listed at the ask p, in dollars a month. Each week it stays empty costs the holding cost $c_h = \$250$ a month (utilities and upkeep) and the rent it is not earning. The mortgage, property tax and insurance are paid whether the unit is let or empty, so they do not change which ask is best, and c_h leaves them out. When a qualified applicant signs, the lease starts the next week. A *qualified applicant* is one who passes the landlord's screening and will pay the ask.

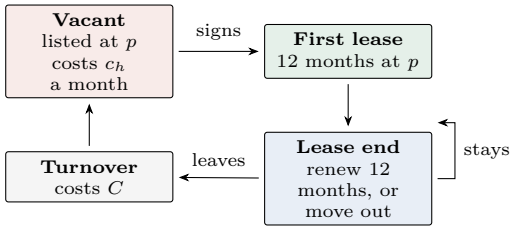


Fig. 1. One cycle of a long-term rental. The ask p decides how long the unit stays vacant, the rent of the first lease, and, through the tenant's chance of leaving, how long the tenancy lasts.

At each lease end the tenant moves out with some probability, and otherwise renews for another 12 months. For a tenant paying the market rent that probability is 45%, so a tenancy lasts 27 months on average: one lease plus a geometric number of renewals. A move-out costs $C = \$3{,}000$. The ask also changes the tenancy: a tenant who pays more than the market rent is likelier to leave at each lease end (Sec. V).

B. Present value

Money now is worth more than money later, so the program discounts every cash flow at 5% a year. All rents, costs and the discount rate are in today's dollars, that is, in real terms after inflation. The *present value* of the vacancy, V, is the discounted sum of everything that follows from it. That is the holding costs of this vacancy, the rent and turnover cost of this tenancy, and the value of every later vacancy and tenancy. Each later vacancy is priced the best way, so V answers one question: what is this vacancy worth if the landlord prices well from now on?

The present value is a large number that is hard to read. The text reports it as the *level income a year*: the constant yearly income with the same present value. For the apartment at its best ask that is $26,253 a year, after holding and turnover costs.

C. Why one month matters

One vacant month costs a month's rent, which is 8% of a year's. A tenancy lasts about 27 months. So one month of vacancy costs as much as cutting the rent by 1/27, or 3.8%, for the whole tenancy, before holding and turnover costs. That ratio drives the rest of this tutorial. A cut in the ask that saves a few weeks of vacancy can pay for itself, if the cut is small and it does shorten the vacancy.

III. Demand During a Vacancy

Demand for a vacant unit shows up as a rate: so many qualified applicants a week. The ask lowers that rate, the season moves it, and the first applicant to arrive ends the vacancy. This section writes that down and shows what it implies for days on market.

A. The applicant rate

Qualified applicants arrive at random, as a Poisson process. In calendar week w, at the ask p, their rate is

$$\lambda(p, w) = \lambda_M \, s(w) \, e^{-(p-M)/\theta} \tag{1}$$

applicants a week. The parameters are:

- M, the unit's *market rent*: the rent at which it attracts $\lambda_M = 0.30$ qualified applicants a week in an average week;

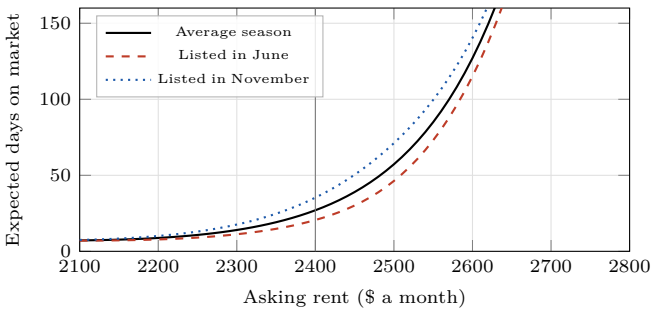


Fig. 2. Expected days from listing to signing, with the ask held fixed. The vertical line is the comparables' market rent, $2,400.

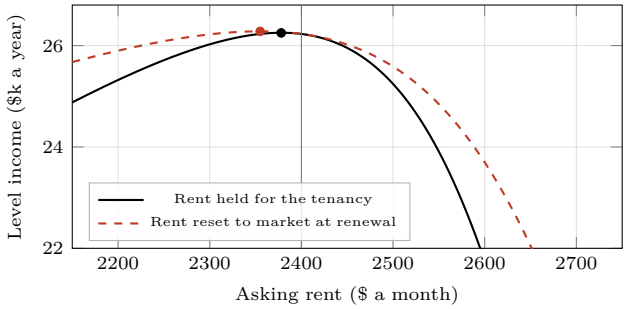


Fig. 3. Level income a year against the ask, one season all year. The dots mark the best asks; the vertical line is the market rent. The two renewal policies are explained in Sec. V.

- $s(w)$, the *season*, a factor with geometric mean 1 that peaks in mid-June and is lowest in December. For the apartment the peak is 2.0 times the trough;
- θ, the *price sensitivity*, in dollars a month: asking θ more divides the rate by $e \approx 2.7$. For the apartment $\theta = \$120$, which is 5% of the rent.

The exponential form says that each extra dollar of rent loses the same fraction of applicants. It is the long-term analogue of the willingness-to-pay model in the companion volume [1]. Real demand need not have this shape, and Sec. XII returns to that.

The chance that the unit leases in a given week is $1 - e^{-\lambda}$. A vacancy is a run of weeks without a signing, so its length depends on the ask through the rate, and through nothing else. In particular it does not depend on how long the unit has already been empty, because a Poisson process has no memory. The program counts a vacancy in whole weeks, so its expected length is $1/(1 - e^{-\lambda})$ weeks.

B. Days on market

Figure 2 shows the expected days from listing to signing against the ask. At the market rent, in an average week, the apartment leases in 27 days. Asking 5% more stretches that to 67 days; asking 5% less shortens it to 13. The curve is steep above the market rent and flat below it, because the expected wait grows exponentially with the ask.

The season moves the days on market too. At the market rent, a vacancy listed in June lasts 21 days; one listed in November lasts 35 days, and one listed in January 32. A November vacancy that drags into the new year meets the weakest weeks of the year.

C. Inquiries

Most people who contact a landlord never apply, or fail screening. The model counts *inquiries*: messages and calls about the listing. It assumes they arrive at $k = 6$ times the rate of qualified applicants, as a Poisson count independent of the applicants themselves. That is a simplification: in Fig. 6 the applicants are a subset of the people who inquire. The landlord sees the inquiries every week, which makes them the main signal for learning the rate during the vacancy (Sec. VIII). The ratio k is taken as known here. In practice it comes from the landlord's past listings or from a property manager's records.

IV. The Best Asking Rent

The best ask balances two losses: rent given up on every month of the tenancy, and rent lost while the unit sits empty. This section derives the balance in closed form, then solves the full model numerically.

A. The closed form

Take one season all year, ignore discounting, and fix the tenancy at L months whatever the rent. A cycle then lasts $L + D(p)$ months, where $D(p) = 1/\lambda(p)$ is the expected vacancy in months. Over a cycle the landlord earns pL, pays the turnover cost C and holds the empty unit for D months at c_h a month. By the renewal-reward theorem [2], the long-run income per month is the expected net income of a cycle divided by its expected length:

$$R(p) = \frac{pL - C - c_h D(p)}{L + D(p)}. \tag{2}$$

Under (1), $D'(p) = D(p)/\theta$. Setting $R'(p) = 0$ gives

$$p^* = \theta \left(1 + \frac{L}{D(p^*)}\right) - c_h + \frac{C}{L}. \tag{3}$$

The ask still appears on both sides, through D, but the right side falls as p rises, so a bisection finds the root. For the apartment, $L = 26.7$ months and the root is $2,462, which is $62 above the market rent. The unit then waits 39 days for a tenant. The program's own model counts a vacancy in whole weeks and starts the lease the week after signing. Under the same assumptions it gives $2,463, close to the closed form.

Equation (3) reads as a markup rule. The term $\theta L/D$ is large when a tenancy is long compared with the vacancy, so a long tenancy justifies a high ask. A small θ, meaning applicants who care a lot about price, pulls the ask down.

B. The full model

The full model drops three simplifications. It discounts at 5% a year. It lets a tenant paying more than the market rent leave sooner, and it charges an expected default cost that rises with the rent. Section V describes both effects. Discounting alone moves the best ask to $2,454. With all three, it is $2,378, $22 *below* the market rent, and the unit waits 23 days.

Figure 3 shows the level income a year against the ask. Asking $100 below the best costs $355 a year, and $200 below costs $1,122. Asking $100 above costs $632, and $200 above costs $3,443. Overpricing is the expensive mistake, because the vacancy grows exponentially with the ask while the rent grows only linearly.

C. A rule of thumb for price cuts

The condition $R'(p) = 0$ also gives an exchange rate between rent and days. Cutting the ask by x dollars a month loses xL over the months the rent binds. It gains the vacancy it saves,

valued at the income rate plus the holding cost, $R + c_h$ a month. The cut pays when

$$\Delta D > \frac{x\,L}{R + c_h}, \tag{4}$$

where ΔD is the vacancy it saves, in months. The rule holds the tenancy and the default risk fixed. In the full model a lower rent also lengthens the tenancy and lowers defaults, so a cut pays with somewhat fewer days saved. For the apartment $R = \$2{,}188$ a month. With the rent held for the whole tenancy, $L = 28$ months. That is a little longer than at the market rent, because the best ask is below it. A $50 cut then has to save more than 18 days of vacancy. If the landlord resets the rent to market at the first renewal, the cut binds for 12 months, and it has to save only 7 days.

D. Accepting less than the ask

An applicant sometimes offers less than the ask. The landlord should accept any offer whose lease is worth more than continuing to wait. In the one-season model the market has no memory, so waiting is worth the same V in every week, and the test is the same in week one and in week six. The test compares a lease that starts next week with an empty unit next week; it ignores a better applicant arriving later in the same week. It also assumes the lower rent stays for the whole tenancy, which lengthens the tenancy. If the rent can be raised at renewal, the lowest acceptable offer is closer to the ask. For the apartment at its best ask of $2,378, the lowest acceptable offer is $2,277, $101 below the ask. This is the *reservation rent* of search theory [7]. An offer between the two is worth taking, and turning it down to hold out for the ask costs money on average.

V. Second-Order Effects

The closed form treats the tenant as the same person at any rent and the year as one season. This section adds five effects that change the answer: tenant quality, tenure, the season, concessions and rent caps.

A. Tenant quality and default

A higher ask shrinks the pool of applicants, and it shrinks it unevenly. Applicants with comfortable incomes still apply; those who would stretch to afford the rent are likelier to fall behind. The model charges an expected default cost at signing: $400 for a tenant at the market rent, multiplied by e for every $300 of rent above it. This stands for unpaid rent, legal costs and a longer turnover, times the chance they happen. On its own it lowers the best ask by $10. The effect is small because the screening that defines a qualified applicant already removes most of the risk.

B. Tenure

Tenants compare their rent with what else they could rent. A tenant paying above the market is likelier to leave at a lease end. The model multiplies the move-out chance of 45% by e for every $400 of rent above the market. This effect lowers the best ask by $69, the largest of the corrections. An extra $50 a month buys a shorter tenancy and therefore an earlier turnover, at $3,000 plus another vacancy.

TABLE I
The best ask and the expected vacancy by listing month with a 12-month lease, and the best first lease when its length can be chosen. The gain is the level income a year that choice adds over always signing 12 months.

Listed in	Ask	Days	Lease (mo.)	Gain/yr
January	$2,380	29	18	$161
April	$2,385	21	15	$135
June	$2,385	18	12	$179
September	$2,375	24	18	$246
November	$2,370	29	18	$217

C. Season and lease length

The program solves the landlord's problem week by week through the year, with the season of (1). Table I gives the results for five listing months. With a 12-month lease, the best ask ranges from $2,370 in November to $2,385 in June, and the vacancy from 18 days in June to 29 in November. The ask hardly moves; the length of the wait does. In this model a landlord should not chase the season with the ask.

The length of the first lease is a second lever. Renewals run 12 months, so a lease signed in November ends in November every year after. Each later move-out then lands in the weakest weeks. The program lets the first lease run from 6 to 18 months. It assumes that a tenant's reasons to move build up with time, so a longer first lease ends with a higher chance of a move-out. A June vacancy keeps the 12-month lease. A November vacancy takes 18 months, which ends in May.

The gains in Table I are for this vacancy and every one after it, since each later vacancy can choose too. Most of the gain does not come from the season. With no season at all, the program still picks 18 months and gains $182 a year, because a longer first lease delays the next turnover. The season adds at most the difference to the largest gain in the table. That split rests on an assumption the model cannot check: that applicants accept an 18-month lease as readily as a 12-month one.

D. Concessions

A *concession* is a discount given as free time, such as one month free on a 12-month lease. At a face rent of $2,378, one free month makes an effective rent of $2,180, or $198 a month less. The arithmetic is the same as a lower rent; the behaviour may not be. Renters who search with a rent filter see the face rent, so a concession does not bring the listing into their results. A lower face rent does. At renewal, the tenant compares the new rent with the face rent, so a concession can make a renewal look like a smaller increase. The model has no data on how renters respond to concessions and does not try to value these effects.

E. Rent caps

The model so far holds the rent at p, in real terms, for the whole tenancy. That is the case under a strict local rent cap, or for a landlord who does not raise rents on sitting tenants. If the landlord resets the rent to the market at each renewal, the ask binds only for the first 12 months. The best ask is then $2,355, $45 below the market, and asking $100 too little costs $178 a year, against $355 (Fig. 3). A cap makes the starting rent matter more, because an error at signing can no longer be fixed at renewal.

California's Tenant Protection Act of 2019 (AB 1482) is an example of a cap [14]. For covered units, it limits rent increases

over any 12 months to 5% plus the change in the regional consumer price index, or 10%, whichever is lower (Cal. Civ. Code § 1947.12(a)). The increase is measured from the lowest rent charged in the prior 12 months. It lets the owner set the initial rent of a new tenancy when no tenant of the old one remains (§ 1947.12(b)), so the ask on a vacant unit is not capped. Its exemptions include housing issued a certificate of occupancy within the previous 15 years. They also include single-family homes and condominiums not owned by a corporation, a real estate investment trust or an LLC with a corporate member, when the tenant receives the required written notice. A duplex is exempt too if its owner has lived in the other unit, as a principal residence, since the tenancy began. The section is due to expire on 1 January 2030. Local ordinances can be stricter. Under the Costa-Hawkins Act, even in a city with rent control, the owner of most units may set the initial rent of a new tenancy (§ 1954.53). The same Act exempts single-family homes, condominiums and newer buildings from local caps (§ 1954.52) [15]. This summary is not legal advice; whether a unit is covered, and how a concession counts towards "the lowest rent charged," are questions for a lawyer.

While inflation stays below 5%, AB 1482 lets the rent rise by about 5% a year in real terms; above 5%, the 10% ceiling leaves less. The program models the first case as `cap5`: a rent below the market moves towards it by at most 5% a year, and a rent above it stays where it is. The `reset` policy lowers an above-market rent to the market at renewal, so the two differ above M. For the apartment, whose best ask is only \$45 below the market, the cap closes the gap at the first renewal. The best ask is the same as with a reset, \$2,355. The cap binds only when the starting rent is about 5% or more below the market, or when the market rises faster than the cap allows.

VI. Estimating Demand from Comparables

The model has two numbers a landlord must estimate before listing: the market rent M and the price sensitivity θ. Comparable listings, or *comps*, are the usual source for both. This section shows how well they estimate each one.

A. What each source measures

- **Listing sites** show asking rents and how long each listing has been up. A signed rent can be lower than the asking rent. A snapshot of the listings up today also over-represents overpriced units, because units priced well lease quickly and drop out of view.
- **Rentometer** [12] reports the spread of rents for units with the same bedroom count near an address. Its numbers come largely from advertised rents, so they share the asking-rent bias.
- **HUD Fair Market Rents** [11] estimate, in most areas, the 40th percentile of gross rent (rent plus utilities) paid by recent movers into standard-quality units, by bedroom count. They are published for each metropolitan area or county, and by ZIP code in areas that use Small Area FMRs. Local housing authorities set voucher payment standards from them, usually between 90% and 110% of the FMR. They start from survey data a few years old, trended forward. They give a benchmark for an area, too broad to price one unit.
- **Rent indices** such as Zillow's ZORI [13] track how a market's typical asking rent changes over time. ZORI compares rents of the same units listed at different times, so a change in which units are listed does not move it. It brings old comps up to date.
- **Signed leases** from a property manager or a local agent are the best source, because they record what tenants paid. They are also the hardest to get.

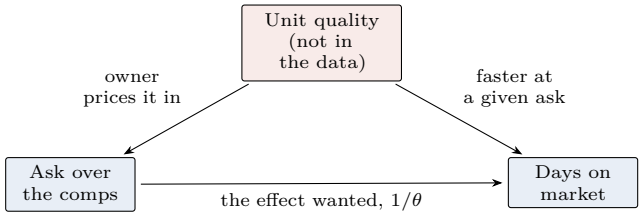


Fig. 4. Why a regression across comps overstates θ. Quality raises the ask and shortens the vacancy, so it hides part of the effect of the ask.

B. The market rent

A good set of comps puts the market rent within a few percent. The apartment's comps say \$2,400, and the truth is within about \$120 of that, one standard deviation, because every comp differs from the apartment in floor, light, finish and street. The model takes that uncertainty as a normal prior on M, centred on the comps. Narrowing it is the job of the vacancy itself (Sec. VIII).

C. The price sensitivity

The obvious way to estimate θ is to regress each comp's days on market against how far its ask sits above the comps' typical rent for its size and location. Under (1), the logarithm of the days on market rises by $1/\theta$ for each dollar of ask. The program simulates 400 comps whose true θ is \$120.

The regression puts θ at \$392, with a 95% interval from \$275 to \$682. The interval misses the truth by a factor of two or more. The cause is unit quality that the analyst cannot see, shown in Fig. 4. Owners of better units know it and ask more: in the simulation, the ask rises by 70% of the quality premium. Better units also lease faster at a given ask. So the units that ask more are partly the units that are worth more, and their days on market barely rise. In expectation only 0.25 of the true slope survives, so θ looks about four times too large. In this sample it is 3.3 times too large. This is the same confounding that the companion volume finds in a host's own booking history [1].

Richer controls help. If the analyst can measure three quarters of the quality variation, from photographs, finishes and condition, the same regression gives \$150 (95% interval \$127 to \$184). That is closer, but the interval still misses. A hedonic model [9], one that prices each feature of a unit, reduces the bias in proportion to the quality it explains.

D. How much an error in θ costs

Figure 5 sets the best ask with a wrong θ and values it under the true one. A θ 25% too low or too high costs \$6 or \$7 a year. Twice the truth costs \$118, and four times costs \$724. Near the truth, the curve is flat. The plain regression's \$392 sets an ask of \$2,467 and costs \$483 a year. An estimate within a factor of about 1.5 is good enough. The plain regression is not.

VII. Why One Landlord's History Is Too Thin

A host of a short-term rental sees hundreds of bookings a year, and can estimate demand from them with care [1]. This section counts the data that one unit gives.

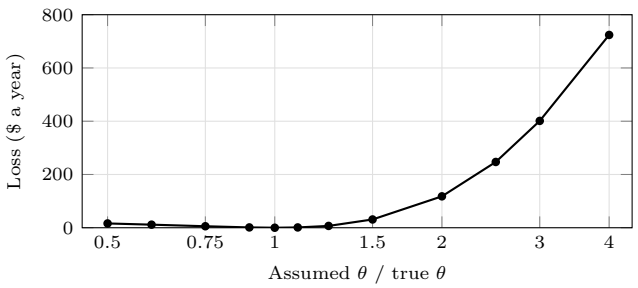


Fig. 5. Income lost by setting the ask with a wrong price sensitivity, one season all year.

A. Past vacancies

At the market rent the apartment comes vacant about once every 2.3 years. In twenty years it has about 9 vacancies. Suppose the landlord's past asks varied by $60 around the market each time, and that the days on market followed (1) exactly. Even if the market rent were known each time, a maximum-likelihood fit of n exponential waiting times would estimate the slope $1/\theta$ with a standard error of $1/(\$60\sqrt{9})$. That is 67% of the slope. A 95% interval for the slope runs from below zero, meaning no sensitivity at all, to 2.3 times the true slope. And real past asks are worse data than this: the landlord raised them in good markets and cut them in bad ones, which is the confounding of Sec. VI again.

B. Price changes within a vacancy

A cut during a vacancy is a better experiment, because it compares the same unit at two asks. Take two weeks at the best ask and two weeks at $100 less, in an average season. The apartment expects 4.3 inquiries in the first two weeks and 10.0 in the second two. The log ratio of the counts estimates $100/\theta$ with a standard error of $\sqrt{1/n_1 + 1/n_2}$, which here is 69% of the value. One cut is barely informative. To bring the standard error down to 20% takes about 12 such cuts, pooled.

C. A workable estimate

One unit cannot pin down θ, but it does not need to. Three facts make the problem workable.

- **The top is flat.** An error of 25% in θ costs a few dollars a year (Fig. 5). A range taken from comps with good controls, or from a property manager who pools many units, is enough.
- **The level matters more.** An error of $100 in the market rent costs $632 a year when the ask is too high. The market rent is also the number the vacancy can teach.
- **Pooling works.** A manager with many units sees enough cuts to estimate θ for a market. A small landlord can borrow that estimate as a prior.

The rest of this tutorial therefore takes θ as known and learns the market rent during the vacancy.

VIII. Learning Demand During the Vacancy

Once the unit is listed, each week brings new counts: how many people looked, wrote, came to see the unit and applied. This section turns those counts into a belief about the market rent.

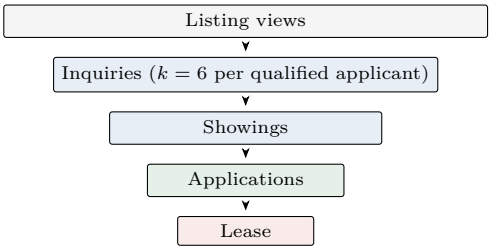


Fig. 6. The funnel of a rental listing. Every stage counts the same underlying demand; the model learns from the inquiries.

A. The funnel

A listing produces a funnel of events, shown in Fig. 6. Each stage is a count, and each count rises with the applicant rate. Views are plentiful but noisy, because many viewers are browsing. Applications are the best signal but too rare to learn from in a few weeks. Inquiries sit in between: they come several times a week, and they come from people who are interested. The model uses the inquiries. A landlord can add showings the same way, as a second count with its own ratio to the applicant rate.

B. The update

Write the unknown market rent as a multiplier on the applicant rate the comps predict. If the comps say $M_0 = \$2{,}400$ and the truth is M, then by (1) the true rate is u times the predicted rate $\hat{\lambda}_t = \lambda_M s(w_t)\, e^{-(p_t - M_0)/\theta}$ for week t's ask p_t, with $u = e^{(M - M_0)/\theta}$. In week t the landlord sees n_t inquiries, a Poisson count with mean $ku\hat{\lambda}_t$, and sees that the unit did not lease, which has probability $e^{-u\hat{\lambda}_t}$.

Both pieces of evidence are Poisson in u, so a Gamma prior stays Gamma [8]. With prior shape a and rate b, after T weeks

$$u \sim \text{Gamma}\Big(a + \textstyle\sum_t n_t,\ b + (k+1)\sum_t \hat{\lambda}_t\Big). \tag{5}$$

The first parameter counts the inquiries. The second counts the inquiries the comps predicted, plus the applicants they predicted, because a week without a lease is evidence too. Their ratio is the posterior mean of u. If the applicants were a subset of the inquirers, as in Fig. 6, the factor $k+1$ would become k. A price cut changes $\hat{\lambda}_t$ for the weeks after it and leaves the formula as it is. A Gamma prior with shape 0.58 and rate 0.35 has the same mean and variance as u when M is normal around M_0 with standard deviation $120. The program keeps the exact posterior on a grid of market rents, which is the same update without that approximation.

C. A worked example

The apartment is listed in mid-Sep. Its comps say $2,400, but they are of better units, and its true market rent is $120 lower. With the ask held at the comps' rent, the landlord would expect 1.8 inquiries a week; the truth gives 0.7.

Figure 7 shows the posterior on the market rent when the inquiries arrive at their expected number and the unit does not lease. The ask follows the markdown policy of Sec. IX, which cuts it as the belief moves, so the inquiries rise a little each week. Before listing, the belief is the comps': mean $2,400, standard deviation $119. After two weeks of thin inquiries the mean is $2,303 and the standard deviation $71. After four weeks they are $2,279 and $51. The landlord has learned the market rent to within about 2%, one standard deviation, from about 4 inquiries

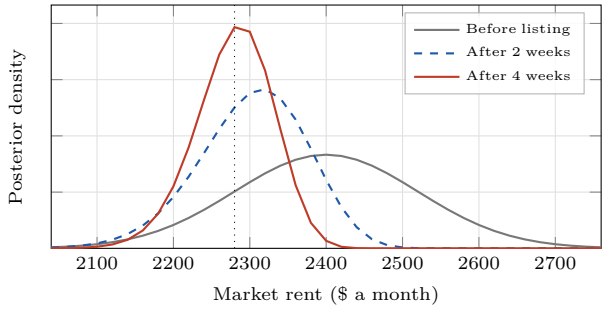


Fig. 7. The landlord's belief about the market rent in the worked example. The dotted line is the truth, $120 below the comps.

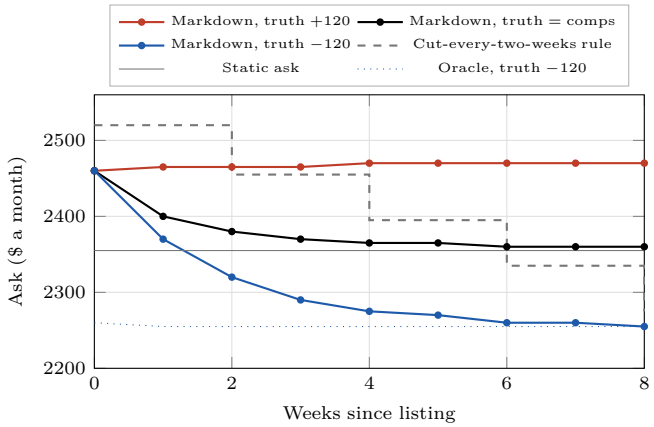


Fig. 8. The ask week by week for a unit listed in mid-Sep that has not yet leased, with the inquiries at their expected number.

and four weeks without an applicant. The weeks without an applicant count as evidence too.

The example is favourable in one way: the inquiries come in at their expected number. Real weekly counts scatter around it, so the belief moves in jumps, and a slow week can mislead. The simulations of Sec. X draw random counts and include that noise.

IX. When to Cut, and by How Much

A landlord who knew the market would set the ask once and wait. A landlord who is learning should cut when the evidence says demand is weaker than believed, and by about as much as the belief has moved. This section shows why, and what the resulting schedule looks like.

A. Why a known market needs no cuts

The applicant process has no memory (Sec. III). After six empty weeks, the chance of a lease next week is the same as it was in week one, apart from the season. So the best ask in week six is the best ask in week one. Days on market are a sunk cost. By themselves they give no reason to cut. A schedule that cuts the price because time has passed, such as "cut 2.5% every two weeks," is right only by accident.

A cut is right for two reasons. The first is the season: a vacancy that runs from October into winter faces fewer applicants each week, and the model lowers its ask slightly as that happens. The second is learning. Thin inquiries mean the market rent is lower than the comps said, so the best ask is lower too. That is a markdown in the sense of revenue management [4]: a sequence of prices driven by what the seller learns about demand.

B. The markdown policy

Each week the policy takes the current posterior on the market rent and asks the rent with the highest expected value for the coming week. That value counts the holding cost, the chance of a lease and the value of the lease if one is signed. If none is, it counts the value of the vacancy next week, computed as if the market rent were then known. This is a one-step approximation of the full Bayesian problem: it treats the next week's information as complete, so it overrates how fast it will learn. It is cheap, and Sec. X measures how close it comes to the best possible.

Figure 8 shows the ask it sets week by week in the worked example of Sec. VIII, for three true market rents, with the inquiries at their expected number and no lease signed.

- It starts at $2,460, above the static ask of $2,355. In its one-step value a high ask leases mostly in the markets where the rent is high, and in the others it costs one week before the policy, believing it then knows the market, cuts. That belief is optimistic, so the high start is partly an artefact of the approximation.
- When the truth is $120 below the comps, the ask falls to $2,320 by week 2 and $2,275 by week 4. It then settles near the ask an oracle that knows the market would set, $2,260.
- When the comps are right, the ask falls to $2,380 by week 2 and stays near $2,365.
- When the truth is $120 above the comps, the inquiries are plentiful and the ask stays at $2,470.

The cuts are largest early, when the belief moves most, and shrink as it settles. Cut when the inquiries fall short of what the current belief predicts. Cut by about the amount the posterior mean of the market rent has moved.

C. The cost of starting high

Starting high has a price. In the week before the evidence arrives, the unit is listed above what the posterior mean would support. The policy accepts that because one or two weeks at a high ask cost less than signing a whole tenancy below the market when the comps were too low. The fixed cut schedule pays that price too, and keeps paying it. It cuts whether or not the inquiries say so, and it stops cutting only when it reaches its floor.

X. Comparing the Policies

This section runs five asking-rent policies on the same 3,000 simulated vacancies and measures what each loses against the best possible.

A. The setup

Each vacancy starts in a random week of the year. Its true market rent is drawn from the comps' uncertainty: normal around $2,400, with standard deviation $120. Each week the program draws whether the unit leases and how many inquiries arrive, from the same random numbers for every policy. The policies are:

- **Comps**: ask the comps' rent, $2,400, and hold it.
- **Rule**: ask 5% over the comps, and cut 2.5% after every two weeks without a lease, down to a floor of 90%.

TABLE II

FIVE POLICIES ON 3,000 SIMULATED VACANCIES. LOSS IS AGAINST THE ORACLE, IN PRESENT-VALUE DOLLARS PER VACANCY (WITH ITS STANDARD ERROR) AND PER YEAR. GAP IS THE SIGNED RENT'S AVERAGE SHORTFALL BELOW THE TRUE MARKET RENT.

Policy	Loss	/yr	Days	Gap	>8 wk
Comps	\$1,541 (71)	\$642	39	\$5	19%
Rule	\$1,181 (40)	\$492	37	\$3	19%
Static	\$1,461 (61)	\$608	30	\$44	13%
Markdown	\$536 (29)	\$223	28	\$20	10%
Oracle	\$0	\$0	24	\$20	7%

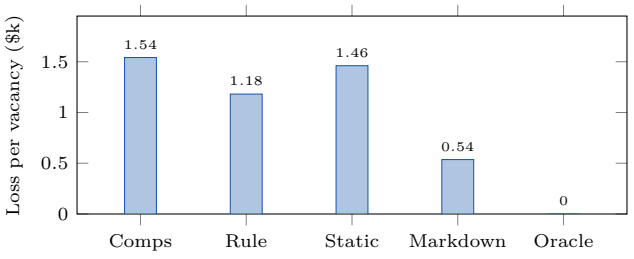


Fig. 9. Loss per vacancy against the oracle, in thousands of present-value dollars.

- **Static**: the best ask under the comps' uncertainty, chosen for the listing week and held.
- **Markdown**: the Bayesian policy of Sec. IX.
- **Oracle**: the best ask for the true market rent, week by week. No landlord can do this; it is the reference.

A vacancy's value is the present value of its holding costs, the tenancy it ends with, and every later cycle priced by the oracle. The loss of a policy is the oracle's value minus its own, averaged over the vacancies. Dividing by the 2.4 years from one vacancy to the next turns it into dollars a year.

B. The results

Table II gives the results, and Fig. 9 plots the losses. Three findings stand out.

Learning. The markdown policy loses \$536 per vacancy, or \$223 a year. Every policy that does not learn loses more than twice as much, between \$1,181 and \$1,541 per vacancy. The markdown policy also has the shortest vacancies, 28 days on average against 24 for the oracle. Learning is worth more than any of the fixed rules.

Cutting with time. The rule loses \$1,181 per vacancy, less than the static ask's \$1,461, although the static ask is optimal for what the landlord knows at listing. Cutting with time is a crude way of learning: a unit still empty after a month is probably overpriced. In the 950 vacancies whose market rent is within \$40 of the comps, the rule loses \$854 per vacancy, against \$105 for the comps' rent and \$237 for the markdown policy. The rule pays for its crudeness when the comps were right.

Bad comps. In the 512 vacancies where the comps overstate the market by \$120 or more, holding the comps' rent loses \$4,775 per vacancy. The markdown policy loses \$1,196. Where the comps understate it by as much, the comps' rent loses \$2,367 and the markdown policy \$895. Most of each policy's loss comes from these cases.

C. How big is this

The apartment earns about \$26,253 a year. The markdown policy saves \$418 a year over pricing at the comps, which is 1.6% of that income. The policy needs only a weekly count of messages. The sum is modest because the value curve of Fig. 3 is flat near its peak, so moderate errors cost little. Most of the saving comes in the vacancies whose comps are off by 5% or more, which a landlord cannot spot from the comps alone.

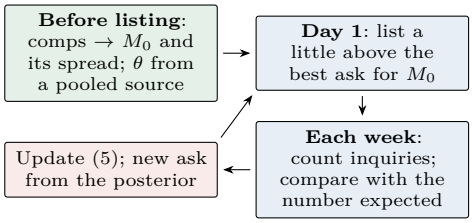


Fig. 10. The weekly loop for one vacancy.

XI. A ROUTINE FOR ONE VACANCY

This section turns the model into steps a landlord can follow with a spreadsheet. Figure 10 shows the loop.

Recipe 1 (Before listing)**.** Collect five to ten comps of the same size and area, preferring signed rents to asking rents. Take their median as M_0, adjusted for the unit's differences. Take the spread as the uncertainty, about 5% of the rent for good comps. Take θ from a pooled source, such as a property manager's records. If there is none, start from the apartment's value: an ask 5% over the market divides the applicant rate by e, about 2.7. Figure 5 shows that an error of 25% costs little. Learn the ratio of inquiries to qualified applicants from past listings.

Recipe 2 (Day 1)**.** Compute the best ask for M_0 from (3), or from the program with your own numbers. List somewhat above it: in the worked example the markdown policy started \$85 above the best ask for M_0. A high first ask costs little if you cut quickly when the inquiries are thin. Consider the length of the first lease: in autumn and winter, a lease that ends in late spring puts the next vacancy in the busiest weeks.

Recipe 3 (Each week)**.** Count the inquiries and compute the number the comps predicted at this week's ask. Update the posterior with (5). The program's markdown policy then maximizes the value expected under the posterior. A spreadsheet can use a shortcut: the best ask for the posterior mean of the market rent. If that is lower than the current ask, cut to it. If inquiries run ahead of expectations, hold the ask; the markdown policy raises it only slightly in that case (Fig. 8). Ignore how long the unit has been empty, and look only at what the counts say.

Recipe 4 (When an applicant offers less)**.** Accept any offer above the reservation rent of Sec. IV, about \$101 below the best ask for the apartment. Screen the applicant the same way at any rent. A tenant who cannot pay costs more than any discount.

Two numbers are worth keeping from every vacancy: the inquiries per week at each ask, and the signed rent. Over a few vacancies, and across units if the landlord has several, they become the pooled data that makes the next estimate of θ better.

XII. WHAT THE MODEL LEAVES OUT

Every number in this tutorial comes from a simulated market whose structure the program chose. Where the real market

differs, the numbers change. These are the differences that matter most.

The shape of demand. The applicant rate falls exponentially with the ask, with one sensitivity at every price. Real demand may fall faster above some threshold, such as a round number that renters use in search filters, or a voucher limit. Raising the ask from $2,395 to $2,405 may drop the unit out of many searches.

One unit, no competition. The model prices one unit against a fixed market. A landlord with several similar units competes with itself, and neighbours change their asks in response to the same season.

Known parameters. The markdown policy learns the market rent but takes θ, the inquiry ratio k and the season as known. If k is wrong by 30%, every update is biased by the same factor, and the policy cuts too much or too little.

Tenant behaviour. The tenure and default effects are simple exponential curves with assumed parameters. They lower the best ask by $69 and $10, so their parameters matter. The lease-length lever assumes renters accept 6 to 18 months as readily as 12.

Listing age and history. Listing sites push new listings to the top and show how long a listing has been up and every price cut. A new listing meets a backlog of searchers in its first week, and a history of cuts may invite low offers. Both break the model's memoryless arrivals.

Applicants one at a time. The first qualified applicant signs. In practice screening takes days, approved applicants sometimes back out, and at a low ask several apply at once, which lets the landlord choose the strongest and lowers the default risk.

Real terms only. Rents and costs are in today's dollars. The model does not forecast rent growth, and its rent-cap scenarios are stated in real terms.

Law. The summary of AB 1482 in Sec. V is not legal advice. Fair housing law also applies to every step of this routine: the ask, the screening and the response to offers must treat all applicants alike.

No real data. Nothing here was measured on a real unit. The model shows how to reason and what to measure. A real unit needs local data.

XIII. Conclusion

The asking rent of a vacant unit is one decision whose effects last for years. This tutorial valued it as the present value of the vacancy, the tenancy and every cycle after it. It found the best ask in four steps, each checked on a simulated market whose truth is known.

First, the best ask balances rent over the tenancy against rent lost while empty. With applicants arriving at a rate that falls exponentially with the ask, the balance has a closed form, (3). For the apartment it is $62 above the market rent. Tenants who leave sooner at higher rents, and defaults that rise with the rent, move it to $22 below.

Second, asking $100 too much costs $632 a year, and $100 too little costs $355. Overpricing is the costlier error. A rent cap makes the starting rent matter more, because a low start cannot be fixed at renewal.

Third, comps give the market rent to within about $120, but a regression across them overstates the price sensitivity 3.3 times in this sample. One landlord has about 9 vacancies in twenty years, far too few to estimate it. Fortunately an error of 25% in the sensitivity costs only a few dollars a year.

Fourth, the weekly inquiries, counted against the number the comps predict, update the belief about the market rent with a Gamma-Poisson formula. A markdown policy built on that update loses $536 per vacancy against an oracle. Holding the comps' rent loses $1,541, and cutting every two weeks loses $1,181.

The practical answer is short. Start a little above the best ask for the comps, count the inquiries every week, and cut when they fall short of what the comps predicted, by as much as the belief has moved. Let the inquiry count and the season decide every cut.

References

[1] *Pricing a Short-Term Rental: Estimating Demand, and Turning the Estimate into a Nightly Rate*, Mathematics & Physical Sciences series, this collection, 2026.

[2] S. M. Ross, *Introduction to Probability Models*, 12th ed. London: Academic Press, 2019, ch. 7 (renewal theory and renewal-reward processes).

[3] M. L. Puterman, *Markov Decision Processes: Discrete Stochastic Dynamic Programming*. New York: Wiley, 1994.

[4] K. T. Talluri and G. J. van Ryzin, *The Theory and Practice of Revenue Management*. Boston, MA: Kluwer, 2004.

[5] G. Gallego and G. van Ryzin, "Optimal dynamic pricing of inventories with stochastic demand over finite horizons," *Management Science*, vol. 40, no. 8, pp. 999–1020, 1994.

[6] R. L. Phillips, *Pricing and Revenue Optimization*. Stanford, CA: Stanford University Press, 2005.

[7] S. A. Lippman and J. J. McCall, "The economics of job search: A survey," *Economic Inquiry*, vol. 14, no. 2, pp. 155–189, 1976.

[8] A. Gelman, J. B. Carlin, H. S. Stern, D. B. Dunson, A. Vehtari and D. B. Rubin, *Bayesian Data Analysis*, 3rd ed. Boca Raton, FL: CRC Press, 2013.

[9] S. Rosen, "Hedonic prices and implicit markets: Product differentiation in pure competition," *Journal of Political Economy*, vol. 82, no. 1, pp. 34–55, 1974.

[10] J. D. Angrist and J.-S. Pischke, *Mostly Harmless Econometrics*. Princeton, NJ: Princeton University Press, 2009.

[11] U.S. Department of Housing and Urban Development, "Fair Market Rents." [Online]. Available: https://www.huduser.gov/portal/datasets/fmr.html

[12] Rentometer. [Online]. Available: https://www.rentometer.com/

[13] Zillow Research, "Zillow Observed Rent Index (ZORI): Methodology." [Online]. Available: https://www.zillow.com/research/data/

[14] California Assembly Bill 1482 (2019), the Tenant Protection Act of 2019, codified at Cal. Civ. Code §§ 1946.2 and 1947.12.

[15] Costa-Hawkins Rental Housing Act, Cal. Civ. Code §§ 1954.50 to 1954.535.