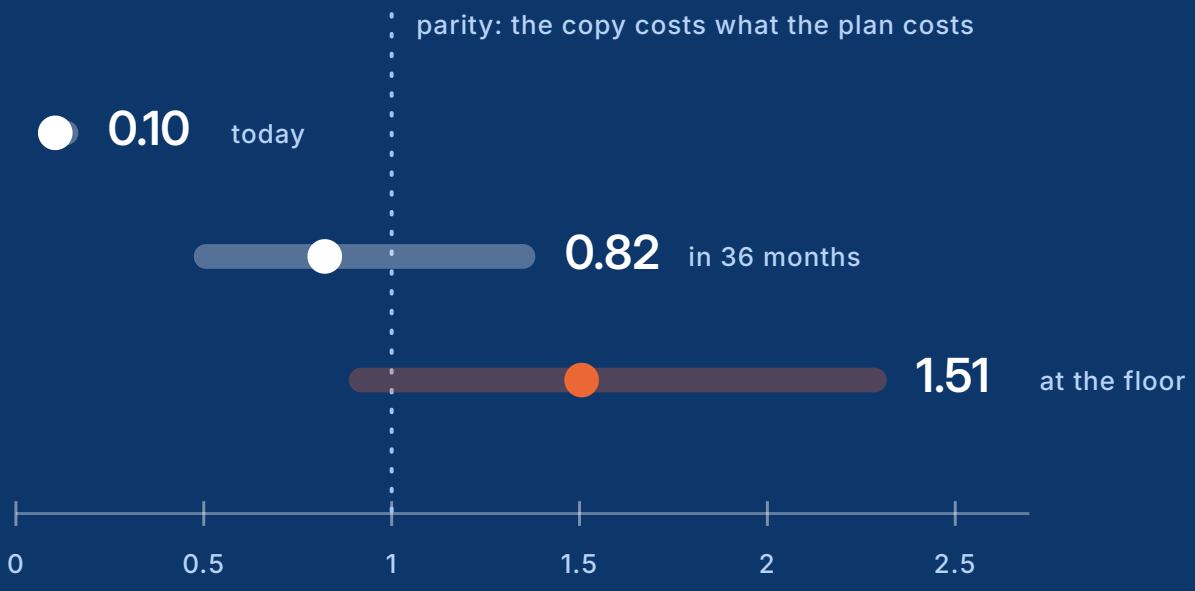
COTBase Premium

What a $14 subscription costs against its own replica: measured from public evidence, and what the owners' sealed test could change



P/R: what the plan costs, divided by what a lawful replica would cost

An illustration at assurance level A1, not a certified score

Contents

SECTIONS

	Summary	3
1	The subject and the question	4
2	The method in brief	5
3	Step 1: size	7
4	Step 2: the replication price today	9
5	Step 3: three years ahead	11
6	The label	13
7	The sealed test: what the owners' run could change	14
8	Price and plans	18
9	Beyond one buyer: groups, Pro and iCOT	19
10	What the owners can do	20
11	Limits and reproduction	20
	Sources	21
	Glossary	22

TABLES

1	Published facts	4
2	Standard conditions	4
3	Inputs and their ranges	7
4	Capability manifest	8
5	One worked draw	9
6	Readings along the path	12
7	How to read the label	14
8	Proof time by noise	17
9	The four list prices	18
10	Levers the standard recognises	20

FIGURES

1	The three readings	3
2	The six layers at COTBase	5
3	The measurement in two tiers	6
4	Anatomy of the replication price	10
5	Distribution of today's replication price	10
6	Where the range comes from	11
7	The P/R path	11
8	Hold probability by month	12
9	The floor and the durable price	13
10	The public label	13
11	Seven evidence scenarios	15
12	What the seal would have to find	16
13	Proof time and the gate	17
14	Price and the gate	18
15	A shared replica by group size	19

OVERVIEW

Summary

On public evidence COTBase Premium costs about a tenth of what a lawful personal replica would cost today. The advantage holds for 27 months at 90% confidence, which is not yet long enough to earn a Value Score.

Replication price today
a personal copy of the core, over 36 months

$4,800

80% range $3,000 to $7,900. The plan costs $504.

P/R today
price divided by replication price

0.10

Range 0.06 to 0.17. Below 1, subscribing is cheaper.

Value holds until
at 90% confidence

month 27

Month 43 at the median. It never ends in 16% of outcomes.

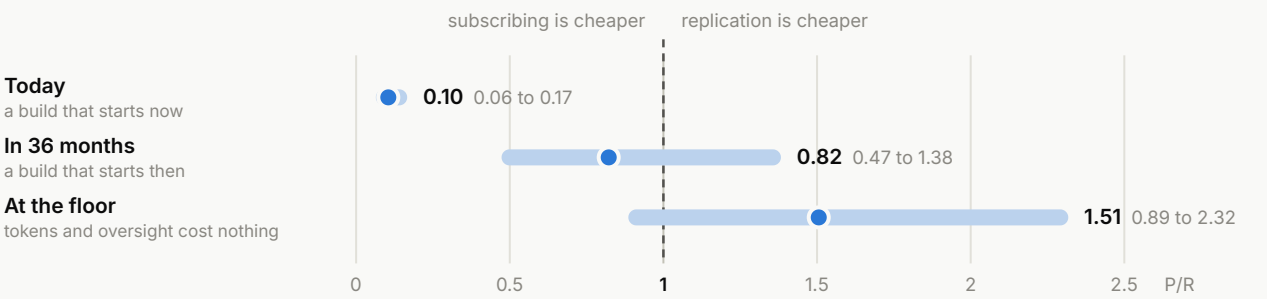
Value Score
needs 90% at month 36

Not issued

The plan is still cheaper at month 36 in 68% of outcomes.

FIGURE 1

A tenth of its replica today, four fifths in three years, above it at the floor



P/R is the outlay divided by the replication price. Dots: central readings. Bars: 80% ranges. Value model v0.2, level A1.

WHAT THE CASE SHOWS

1 Subscribe
A personal copy costs about ten times the plan today. Its upkeep and attention alone cost more than the plan.

2 People's time protects it
Tokens are 18% of the replication price. The rest is the time of the people who build, maintain and run the copy.

3 The protection erodes
AI-priced cost halves every 6 to 12 months. The top of the P/R range passes parity at month 27.

4 A seal moves the level
It replaces two guesses with measurements and can raise the estimate or lower it. It narrows the range only a little.

5 Proof time buys the score
Visible finesse adds months, and AI removes them again. Five to seven months of proof time pass the gate.

6 Price is the other lever
The test passes at any price under $10 a month. The standard reports that and recommends no price.

An illustration, not a certified score

Level A1: nobody has inspected the COTBase code. Size, unit costs, hidden work and finesse are declared ranges of evidence grade C, and every sealed finding in this document is invented to show the mechanics.

SECTION 1

The subject and the question

The case asks one question of one plan: is $14 a month cheaper than rebuilding what it buys, today and three years from now?

The subject

COTBase.com presents the Commitments of Traders reports of the US Commodity Futures Trading Commission as charts and analysis.[1] The site covers 60 futures markets and 11 market complexes, with history back to 1986 for some markets.[2] It sells a free tier, Premium, Pro and two iCOT plans. This case measures Premium.

TABLE 1 **Published facts (grade A)**

FACT	VALUE	SOURCE
Premium price	$19.99 a month; $17 on the six-month plan; $14 on the yearly plan; $12 on the two-year plan. The page banner quotes the six-month plan at $99, which is $16.50 a month	cotbase.com[1]
What Premium adds	All COT history and every analysis feature on the web: corrected and combined series, market complexes, options-only data, custom COT Index, extreme tracker, volume and open-interest scoring, highlights	Features page[2]
Other plans	Pro and iCOT at $99 a month, $83 on the yearly plan. iCOT with AfterSignal at $149 a month. All outside this case	Home and indicators pages[1, 3]
Underlying data	CFTC reports, public and free, released on Fridays with Tuesday's positions. A twice-weekly schedule is announced for the end of 2026	CFTC,[4] World Grain[5]
Free substitutes for basic charts	The CFTC's own reports, Barchart, and COTBase's free tier	Barchart[6]
Wage	$65.38 an hour, from a median of $135,980 a year for US software developers	BLS[7]

TABLE 2 **Standard conditions of this measurement**

CONDITION	VALUE IN THIS CASE
Buyer	One user with the plan's standard profile: a weekly look at their own markets with the Premium tools
Fidelity and quality	The core that buyer uses, at quality parity: the same outputs as the original from the same public reports, within rounding
Horizon	36 months, no discounting
Price	$14 a month, the lowest list price for a commitment of 12 months or less. The outlay is $504
Index	None is measured yet. Unit costs and the pace of AI are declared ranges
Statistic	The median, with the 10th and 90th percentiles of 100,000 seeded draws

Why the core

A single buyer does not need the whole service. They need the part they use, kept working for themselves. That is the cheapest lawful replica a subscriber

could choose, so it is the hardest test of the price. Groups that share one replica and the Pro plan are treated in Section 9.

What the price buys

A subscription sells six layers, and tokens can buy only some of them. The score prices what can be replicated and reports the rest beside it.

FIGURE 2

At COTBase the score prices four layers and reports two beside it

LAYER	AT COTBASE	CLASS	IN THE SCORE
Function	Net-position charts, COT Index with any lookback, extreme tracker, volume and open-interest scoring, highlights	Token-priced	Priced
Data	CFTC reports, with a corrected and combined history and 11 market complexes	Token-priced	Priced: the pipeline, not the data
Operations	Weekly ingest, hosting and support	Market-priced	Priced as upkeep and attention
Stream	New tools and fixes as they ship	Market-priced	Priced as upkeep
Content and judgement	The On COT guides, the blog, the YouTube channel, consultations	Not for sale	Beside the score
Trust and network	More than a decade in operation, and an audience	Not for sale	Beside the score
iCOT, a separate plan	An intraday estimate computed from market data by a proprietary algorithm	Hidden	Outside this case

Classes as in version 0.1 of the standard. Content, trust and anything hidden are never folded into the number, which keeps it conservative.

SECTION 2

The method in brief

The replication price is built from a size, a unit cost and the time people must spend. It is then projected forward on the pace of AI and set against the outlay.

The work

The build work is the counted size, enlarged for the work an outside view misses and for finesse, and priced per function point in tokens and oversight hours.

$$W = Z\,(1+\zeta)\,\varphi\,(\tau + \eta\,w) \tag{1}$$

Today

A build can fail, so its cost is divided by the probability of success. A copy also has to be kept working and looked after for as long as it is used. Proof time is zero unless a sealed run establishes it.

$$\mathrm{ORC}_0 = \underbrace{\frac{W + h_0\,w + n\,p}{\pi}}_{\text{build, risk-adjusted}} + \underbrace{\mu\,W\,\frac{T}{12}}_{\text{upkeep}} + \underbrace{a\,w\,T}_{\text{attention}} + \underbrace{L\,(p + a_v\,w)}_{\text{proof time}} \tag{2}$$

The floor

At the floor tokens and oversight cost nothing. What remains is an hour to commission the build, a month of access to observe the original, and the user's own attention.

$$\mathrm{ORC}_{\infty} \;=\; \frac{c\,w + n\,p}{\pi} \;+\; a_{\infty}\,w\,T \;+\; L\,(p + a_{v,\infty}\,w) \tag{3}$$

The path

Between the two, every AI-priced cost halves each h months.

$$\mathrm{ORC}(t) \;=\; \mathrm{ORC}_{\infty} + (\mathrm{ORC}_{0} - \mathrm{ORC}_{\infty})\,2^{-t/h} \tag{4}$$

The unit

The unit is price-to-replication. The hold probability is the share of outcomes in which subscribing is still cheaper.

$$\mathrm{P/R}(t) \;=\; \frac{S}{\mathrm{ORC}(t)}, \qquad S \;=\; p\,T, \qquad H(t) \;=\; \Pr\!\left[\,S < \mathrm{ORC}(t)\,\right] \tag{5}$$

The score

The Value Score is the discount to replication price that remains at month 36, at the cautious end of the range.

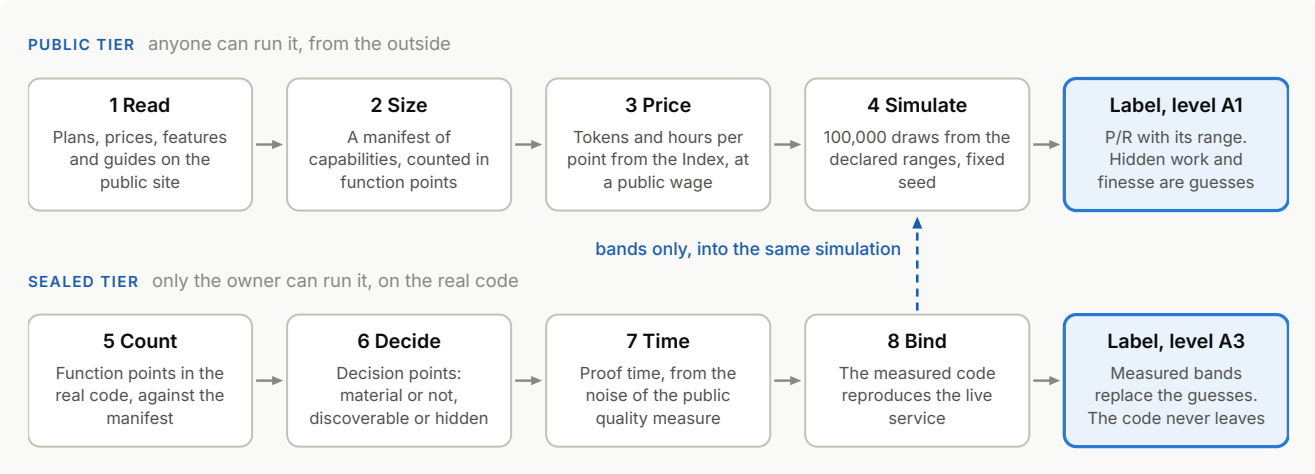
$$\text{Value Score} \;=\; 100\left(1 - \frac{S}{\mathrm{ORC}_{36}^{P10}}\right), \qquad \text{issued only if } H(36) \geq 0.90 \tag{6}$$

Two tiers

The public steps need nothing from the owner. The sealed steps replace the two largest guesses, hidden work and finesse, with measured bands.

FIGURE 3

Four public steps give the label. Four sealed steps can replace its two largest guesses



The public tier needs nothing from the owner. The sealed tier runs an open instrument inside the owner's environment and publishes bands, ranges and a hash.

TABLE 3 **Inputs: every uncertain one is a range**

SYMBOL	INPUT	10TH TO 90TH PERCENTILE	GRADE	SETTLED BY
Z	Size of the standard buyer's core	80 to 125 function points	C	A second count, or a sealed count
τ, η	Tokens and oversight per point	$2 to $6 and 0.08 to 0.24 hours, moving together	C	A measured Index
ζ	Hidden work	5% to 40% of the counted size	C	A sealed run
φ	Finesse factor	1.1 to 1.5, grade Φ2	C	A sealed run
a, a_∞	Attention to run a personal copy	0.25 to 1 hour a month today; 0.05 to 0.2 at the floor	C	A measured Index
h	Halving time of AI-priced cost	6 to 12 months	C	Time
L, a_v	Proof time, and attention while verifying	0 months on public evidence	C	A sealed run
w, p	Wage; price	$65.38 an hour; $14 a month	A	Published
π, h_0, n, μ, c	Success probability, setup, access, upkeep, commissioning	0.8; 3 hours; 1 month; 15% a year; 1 hour	C	Declared constants

The simulation

Each range is read as the 10th and 90th percentile of a lognormal distribution. The checker draws every input 100,000 times with the fixed seed 20261113 and recomputes the result, so anyone who runs the same inputs gets the same numbers. The unit costs bracket published AI rebuilds: a C compiler built by agents,[8] a web framework rebuilt in a week[9] and a benchmark of rebuilt programs,[10] at about 50 lines of code per function point.[11] The halving time brackets measured falls in the price of a fixed level of AI capability[12] and the growth of the tasks agents complete.[13]

SECTION 3

Step 1: size

Counted from the outside, the core a Premium subscriber uses is about 100 function points, with a range of 80 to 125.

The manifest

Size is counted in COSMIC function points: one point for each movement of data into, out of, to or from storage in a functional process.[14] The count below was made from the public feature pages. It is a rough count with evidence grade C, and two counters would not agree exactly. The range of 80 to 125 points allows for that.

TABLE 4 **Capability manifest of the standard buyer's core**

CAPABILITY	WHAT A REPLICA MUST DO	POINTS	SEEN
Ingest of CFTC reports	Fetch each release, parse four report families in two variants, store the positions	14	O1
Market catalogue and history	60 markets, 14 data sets and their history, back to 1986 for some	8	O1
Corrected series	Rescale history for past changes in contract size	6	O2
Combined series	Merge mini and standard contracts into one series	5	O2
Market complexes	Convert positions to dollar equivalents and sum them for 11 groups	8	O2
Options-only view	Derive option positions from the two published variants	4	O1
Net positions and changes	Per trader group, for every report family	8	O1
COT Index, any lookback	Percentile-style index with a lookback set per trader group	5	O1
Extreme tracker	Reports since the last equal reading, with an all-time flag	5	O1
Volume and open-interest score	A signed score with an adjustable lookback	6	O2
Analysis table	One row per market, sortable, with the signals above	7	O1
Highlights	Long and short candidates picked from the signals	5	O2
Interactive charts	Series selection, crosshair, scrolling, a time range from one month to all history	12	O1
Schedule and settings	A weekly update job and the user's saved defaults	7	O1
Total		**100**	

Seen: O1, specified by public documents and data; O2, revealed by the outputs a subscriber can observe.

Four remarks

1. **Observability is high.** Every output is a fixed function of public reports. Nine capabilities follow from public definitions. Five are undocumented but show in the outputs.
2. **That sets the finesse grade.** The corrected series, the complexes and the scoring rules are calibration that outputs reveal. That is grade Φ2, tuned, with a declared factor of 1.1 to 1.5.
3. **Hidden work is a guess.** Edge cases, data repairs and error handling do not show in a feature list. The public tier declares 5% to 40% and waits for a sealed count.
4. **Content and trust are left out.** The guides, the videos and a decade of operation are not counted. They appear beside the score, not in it.

SECTION 4

Step 2: the replication price today

A personal copy of the core costs about $4,800 over 36 months, with an 80% range of $3,000 to $7,900. The plan costs $504.

TABLE 5 **One worked draw: every input at its median**

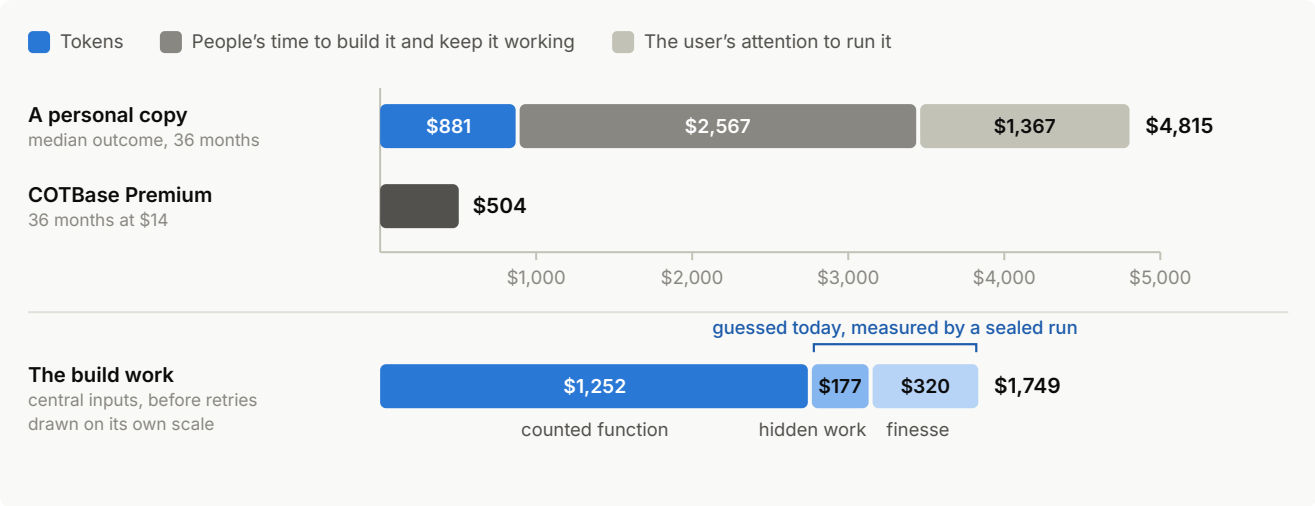
LINE	CALCULATION	AMOUNT
Counted function	100 points at $12.52: $3.46 of tokens and 0.139 hours at $65.38	$1,252.27
Hidden work	14.1% more, the median of 5% to 40%	$177.10
Finesse	22.4% on top, the median factor of 1.22	$319.62
Build work W	The three lines above	**$1,748.99**
Setup	3 hours at $65.38	$196.12
Access to observe the original	1 month at $14	$14.00
Build, risk-adjusted	Work, setup and access, divided by 0.8	$2,448.89
Upkeep	15% of the build work a year, for 3 years	$787.04
Attention	0.5 hours a month at $65.38, for 36 months	$1,176.75
Replication price ORC_0	Build, upkeep and attention	**$4,412.69**
Outlay S	36 months at $14	$504.00
P/R	Outlay divided by replication price	**0.11**

From one draw to 100,000

The table follows one draw through equation (1) and equation (2). The simulation repeats it 100,000 times with the inputs drawn from their ranges. Its median is $4,815, above the $4,413 of the central draw, because the ranges are skewed: hours and tokens can overshoot by more than they can undershoot.

FIGURE 4

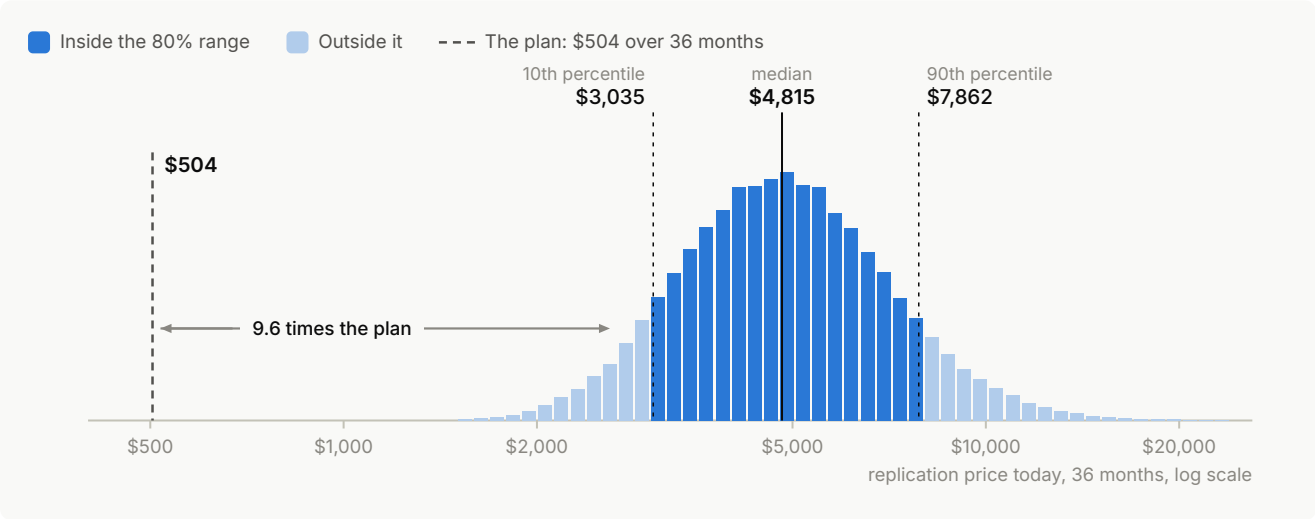
Tokens are 18% of the replication price. People's time is the rest



Top: the components of the median outcome, averaged over the simulated outcomes nearest the median. People's time is oversight of the build and its upkeep, with setup, at the developer wage. Bottom: the build work of the central draw, split by what is built. Value model v0.2, grade C.

FIGURE 5

In every simulated outcome a personal copy costs more than the plan today



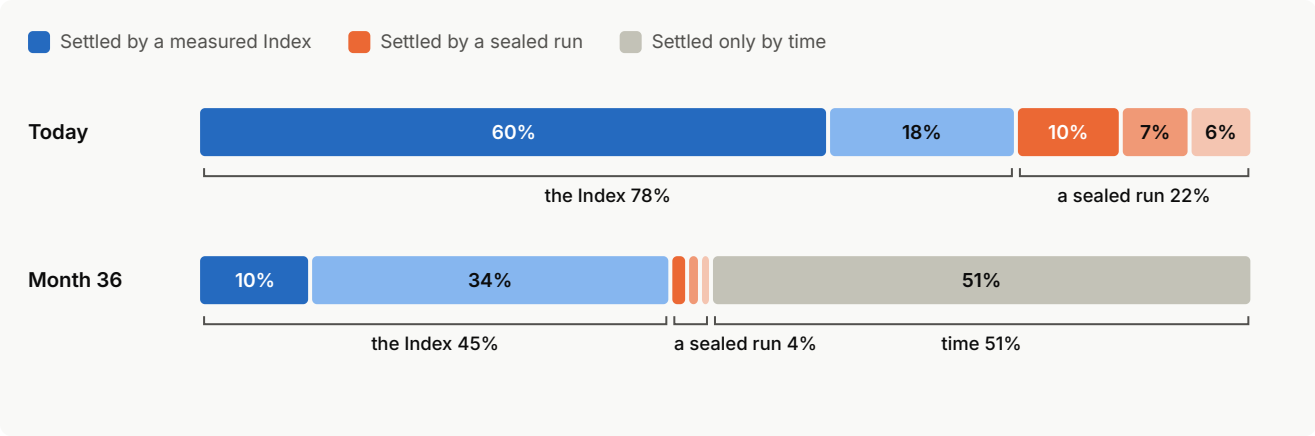
100,000 outcomes. The 1st percentile is $2,139, more than four times the outlay of $504. The scale is logarithmic, so equal distances are equal ratios.

Where the range comes from

The range is wide, a factor of 2.6 from its 10th to its 90th percentile, and the standard says where the width comes from. 78% of today's uncertainty is in unit costs and attention, which a measured Index would settle. 22% is in size, hidden work and finesse, which a sealed run would settle.

FIGURE 6

Most of today's range is the missing Index. Three years out, half of it is the pace of AI



Shares of the variance of the logarithm of the replication price, each input group varied alone and the rest held at their medians. Rounded, so a bar may not sum to 100.

SECTION 5

Step 3: three years ahead

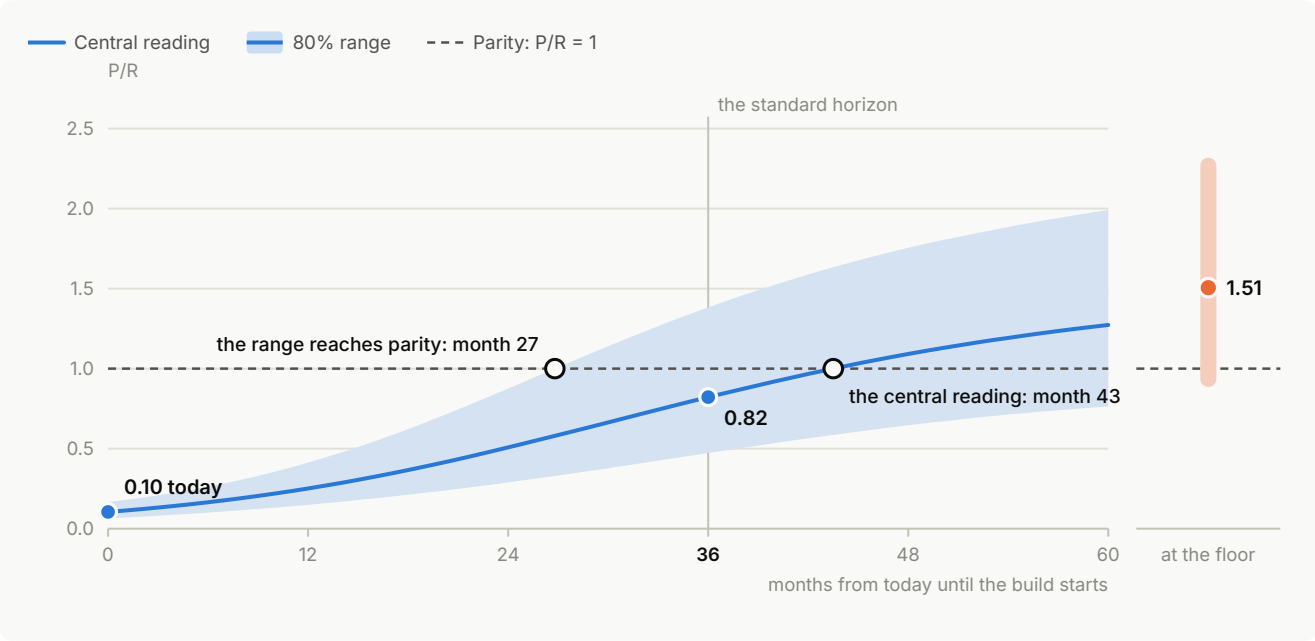
The same build gets cheaper every month. P/R rises from 0.10 today to 0.82 at month 36, and the cautious end of its range passes parity at month 27.

The question

A subscriber can switch to a replica in any month, so the long-term question is whether some later build date beats staying subscribed. The path below prices a build that starts t months from now, using equation (4).

FIGURE 7

P/R passes parity at month 43 on the central reading, and at month 27 at the cautious end



Each point is the outlay divided by the replication price of a build that starts in that month. The orange mark is the terminal reading, taken at the floor. Value model v0.2, assurance level A1.

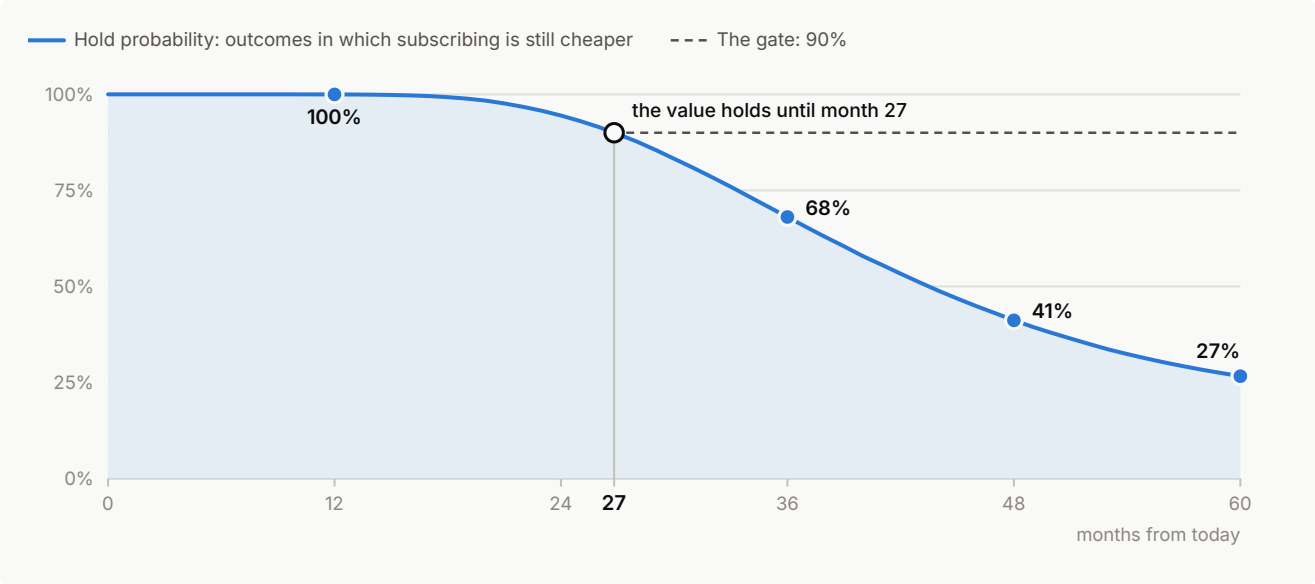
TABLE 6 **Readings along the path**

BUILD STARTS	REPLICATION PRICE	80% RANGE	P/R	80% RANGE	HOLD PROBABILITY
Today	$4,815	$3,035 to $7,862	0.10	0.06 to 0.17	100%
Month 12	$2,005	$1,221 to $3,373	0.25	0.15 to 0.41	100%
Month 24	$992	$576 to $1,742	0.51	0.29 to 0.88	94%
Month 36	$613	$365 to $1,065	0.82	0.47 to 1.38	68%
Month 48	$462	$287 to $780	1.09	0.65 to 1.75	41%
Month 60	$396	$253 to $660	1.27	0.76 to 1.99	27%
At the floor	$335	$217 to $569	1.51	0.89 to 2.32	16%

The outlay is $504 throughout. Hold probability: the share of outcomes in which subscribing is still cheaper.

FIGURE 8

The plan is certain to be cheaper for a year, and likely to be for three



The curve crosses the 90% gate at month 27. At month 36 it stands at 68%, so the test for a Value Score is not passed.

The runway

The month in which a replica first becomes cheaper is the runway. It is 27 months at the 10th percentile, 33 at the 25th and 43 at the median. In 16% of outcomes it never comes, because the floor already lies above the outlay.

$$t^{*} \;=\; h\,\log_{2}\frac{\mathrm{ORC}_{0}-\mathrm{ORC}_{\infty}}{S-\mathrm{ORC}_{\infty}} \tag{7}$$

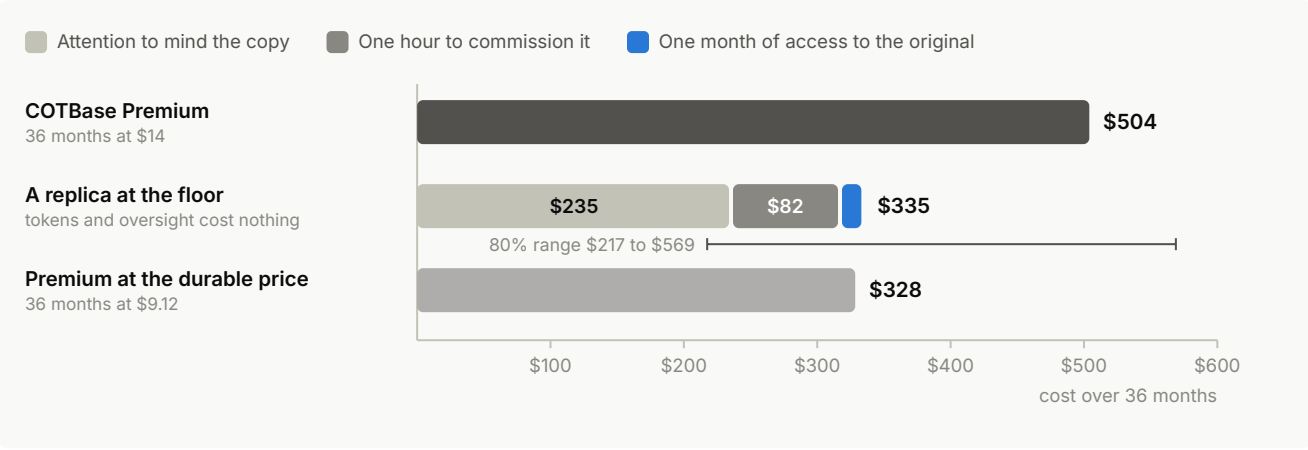
The floor

At the floor a replica costs $335: $235 for the attention to mind it, $82 for the hour that commissions it and $17.50 for a month of access. That is below the outlay of $504, which is why the terminal P/R is 1.51. The price the floor supports is the durable price.

$$p_{\mathrm{durable}} \;=\; \frac{c\,w/\pi \;+\; a_{\infty}\,w\,T \;+\; L\,a_{v,\infty}\,w}{T \;-\; n/\pi \;-\; L} \tag{8}$$

FIGURE 9

When building is free a replica still costs $335, and that supports $9.12 a month



The floor rests on what a user's attention is worth. Across the declared range of that attention the durable price runs from $5.74 to $15.90 a month.

SECTION 6

The label

The public label says subscribe, names the month until which that holds, and issues no Value Score.

FIGURE 10

The label for COTBase Premium on public evidence

ORC VALUE LABEL · ILLUSTRATION

LEVEL A1 · PUBLIC EVIDENCE

COTBase Premium

One user · the core at quality parity · 36 months · $14 a month on the yearly plan · outlay $504

Subscribe

the advice runs in one direction only

P/R 0.10 today

range 0.06 to 0.17: the plan costs a tenth of its replica

Replication price today **$4,800** range $3,000 to $7,900	P/R in 36 months **0.82** range 0.47 to 1.38	P/R at the floor **1.51** range 0.89 to 2.32
Hold probability **100% · 94% · 68%** at 12, 24 and 36 months	The value holds until **month 27** at 90% confidence; month 43 at the median	Finesse and seal **Φ2, declared** no sealed share evidenced

Value Score: not issued. The plan is still cheaper at month 36 in 68% of outcomes, and the test requires 90%.

Standard 0.2 · Index: declared ranges, no epoch · Manifest: outside count, 13 November 2026 · Seed 20261113 · Lapses after 90 days

An illustration of the label format of version 0.2. It is not a certified score: nobody has inspected the code.

TABLE 7 **How to read it**

LINE	IT SAYS	IT MEANS
Advice	Subscribe	P/R is below 1 today. Above 1 the label would report the premium and what it buys. It never says clone
P/R today	0.10	The 36-month outlay is a tenth of the price of a personal replica built now
P/R in 36 months	0.82, range to 1.38	Still cheaper on the central reading. The range reaches above 1, so it may not be
P/R at the floor	1.51	When building is free a personal copy costs less than the plan in most outcomes
Hold probability	100%, 94%, 68%	The share of outcomes in which subscribing is still cheaper at each date
The value holds until	Month 27	The last month in which that share is at least 90%
Finesse and seal	Φ2, declared	Know-how is a guess from the outside. Nothing has been measured under seal
Value Score	Not issued	It needs 90% at month 36. The plan reaches 68%

SECTION 7

The sealed test: what the owners' run could change

A sealed run replaces two guesses with measurements. It can move the estimate up or down. On its own it earns a score only if the code is far more crafted than the outside suggests. Proof time is what changes the result.

What the run does

The owners would run an open, read-only instrument inside their own environment. The code is never copied or sent.

1. **Count.** It counts function points in the real code and compares them with the public manifest. The difference is hidden work.
2. **List.** It lists decision points in the quality-critical code: constants, thresholds, weights, rule rows and special cases.
3. **Test.** It perturbs each one and reruns the quality harness. Only those that move the public quality measure count.
4. **Classify.** It checks whether the change shows in what an outsider can observe. If so the decision is discoverable. If not it is hidden.
5. **Time.** It computes the proof time from the noise of the public quality measure.
6. **Bind.** It confirms that the measured code reproduces the live service on a random probe set.

What leaves

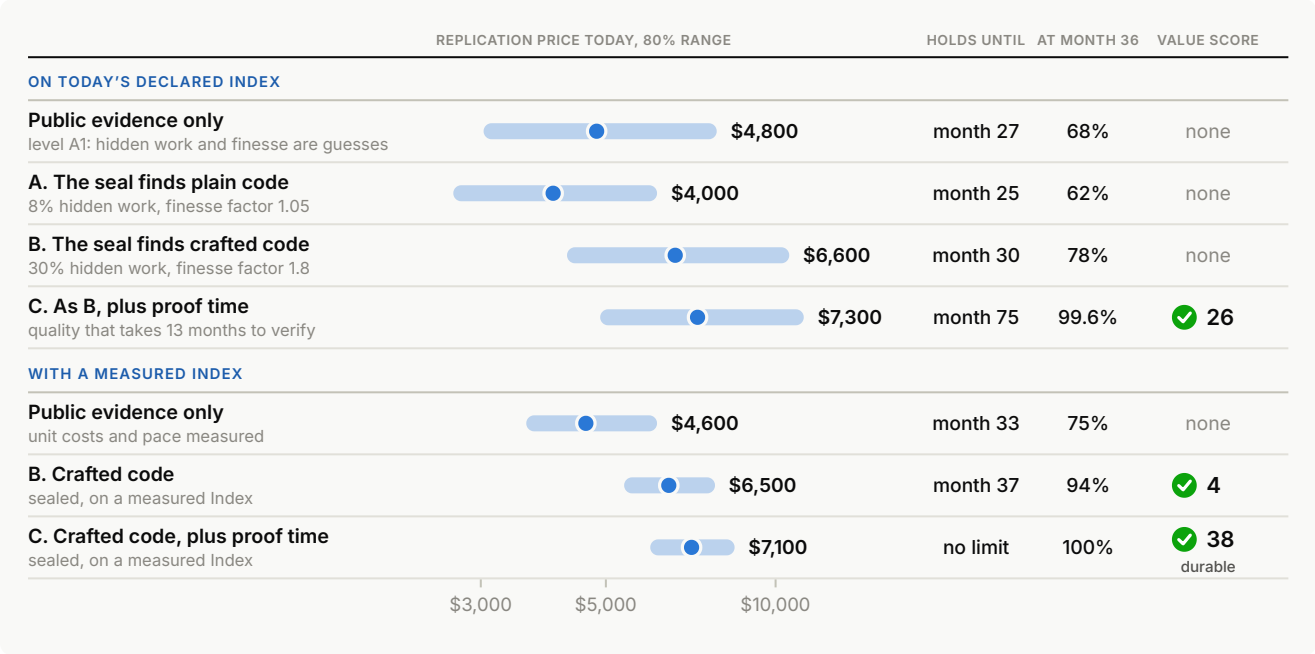
Bands, ranges and a hash of the sealed record leave the owners' environment. The record itself stays with them and is opened only to an arbiter in a dispute.

What it can find here

For the Premium core the public anchor is strong. Its outputs are functions of public reports, so a replica can be checked against the original on the day it runs. That makes its finesse discoverable and its proof time zero. A sealed run on the core can therefore find hidden work and discoverable finesse. Proof time arises only for a component that is judged against data published later, such as an estimate or a forecast.

FIGURE 11

A sealed run can lower the estimate, raise it, or earn the score



The findings in rows A to C are invented to show the mechanics. Each row is a full simulation of 100,000 outcomes. Holds until: the month to which subscribing stays cheaper at 90% confidence. Value model v0.2.

Four lessons

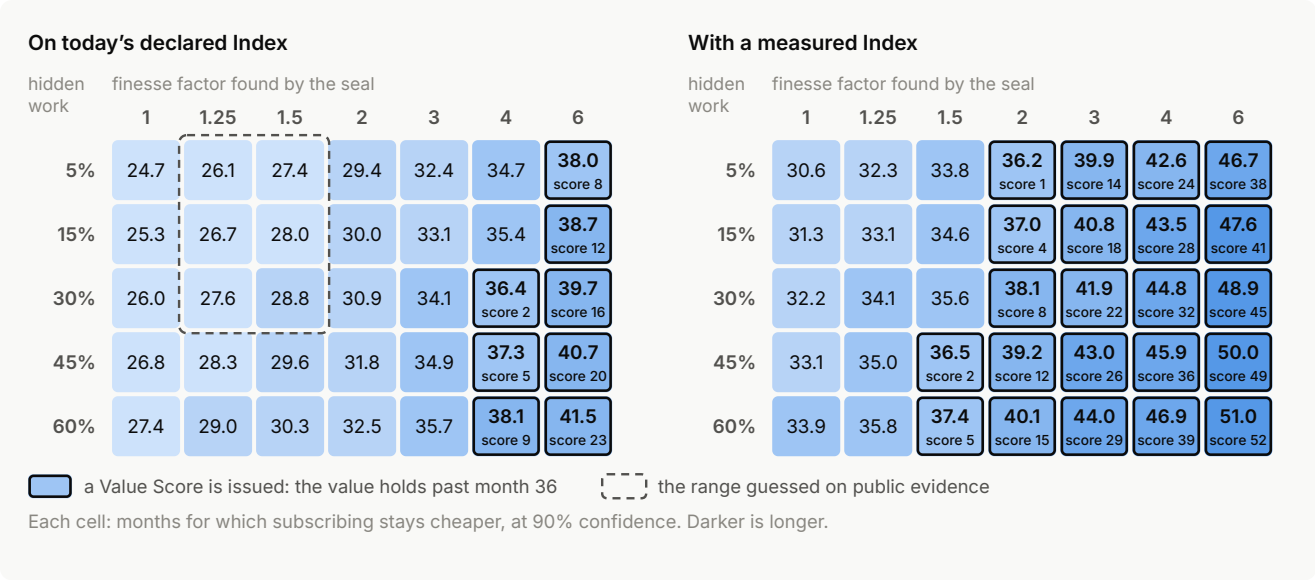
1. **A seal cuts both ways.** Scenario A lowers the price of a replica to $4,000 and shortens the value to month 25. An owner who seals commits to publish before the result is known.
2. **Visible finesse buys months.** Scenario B raises the replication price by more than a third and adds three months. The hold probability at month 36 rises from 68% to 78%, still short of the gate.
3. **Proof time changes the result.** In scenario C the floor rises from $335 to $618, above the outlay of $504, and a score of 26 is issued.
4. **Measurement compounds.** A measured Index cuts the spread of today's price from 2.6 times to 1.7 times. With a sealed run on top it falls to 1.4 times, and scenario B then passes narrowly with a score of 4.

The whole envelope

Whatever the run finds, the result can be read off before it is run. On today's declared Index the estimate could land anywhere from $3,800, for plain code with little hidden work, to about $22,000, for code six times its plain size. The month to which the value holds moves much less: from 25 to 41.

FIGURE 12

On finesse alone, the code would need four to six times its plain build work to earn a score



Each cell is a simulation with the size, the hidden work and the finesse factor measured to a narrow band and no proof time. Left: unit costs and the pace of AI as declared today. Right: with the spreads a measured Index would leave.

Reading the map

1. **Each doubling of finesse adds about five months.** At 30% hidden work the value holds to month 29 at a factor of 1.5, month 34 at 3 and month 40 at 6. The cost doubles, and AI removes the difference in about one halving time, at the fast end of its range.
2. **Hidden work matters less.** Going from 5% to 60% of hidden work adds about three months at any factor.
3. **The Index matters more than either.** With a measured Index a factor of 2 passes at 30% hidden work, where a factor of 4 is needed today. A narrower range lifts its cautious end.

Proof time

Finesse that an outsider cannot observe has to be re-created and then proven equal. Proof takes observations, and observations arrive at the pace of the calendar. The number needed follows from the noise of the public quality measure.

$$L = \Delta \left(\frac{(z_\alpha + z_\beta)\,\sigma}{\varepsilon} \right)^2 \tag{9}$$

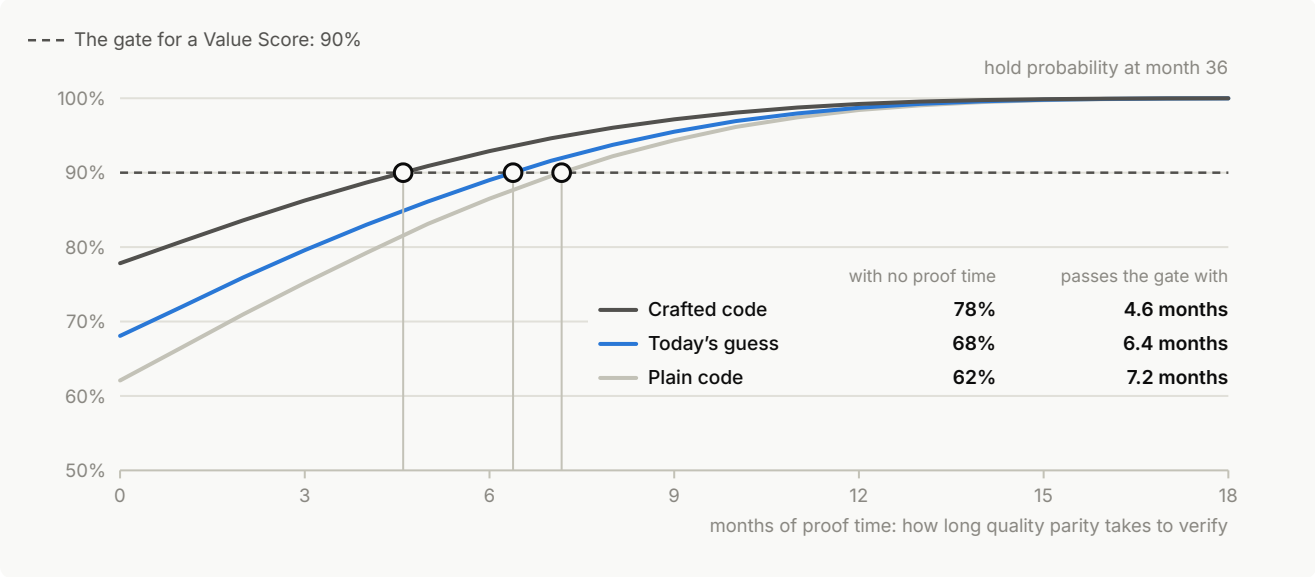
Δ is the time between independent public observations, σ their noise and ε the tolerance within which parity must be shown, at 5% significance and 80% power. During that time a replicator pays for a seat on the original and for the attention to compare.

TABLE 8 **Proof time by the noise of the quality measure**

NOISE OVER TOLERANCE	OBSERVATIONS NEEDED	WEEKLY OBSERVATIONS	TWICE-WEEKLY OBSERVATIONS
1 to 1	7	1.4 months	0.7 months
2 to 1	25	5.7 months	2.8 months
3 to 1	56	12.8 months	6.4 months
4 to 1	99	22.8 months	11.4 months

FIGURE 13

Five to seven months of proof time pass the gate, whatever the finesse



Hold probability at month 36 as a function of proof time, for three findings on finesse. From 15 months the floor itself lies above the outlay in nine outcomes out of ten, and the score is marked durable.

What it changes

Proof time works on the floor, which AI does not lower. With crafted code 4.6 months are enough to pass. With plain code 7.2 months are. For a quality measure observed weekly, that is a noise of a little over twice the tolerance. This is the rule of version 0.2 at work: discovery gets cheaper with AI, and proof does not.

SECTION 8

Price and plans

The same test passes on public evidence at any price under $10 a month. At $14 it reaches 68%.

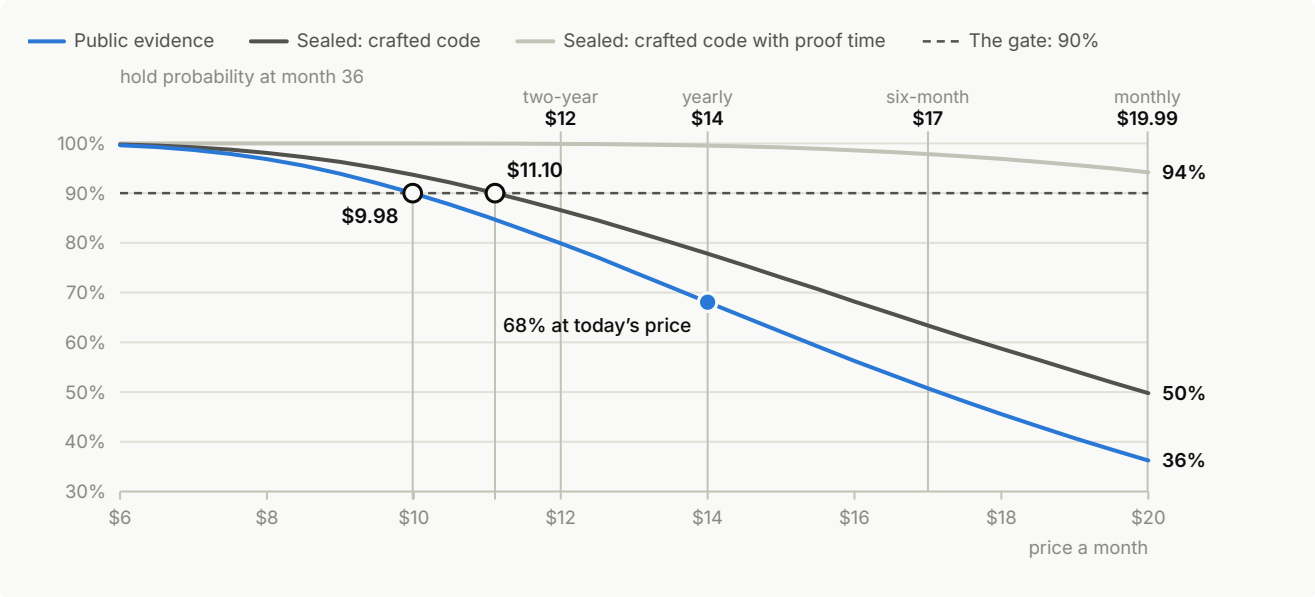
TABLE 9 **The four list prices of Premium**

PLAN	PRICE A MONTH	OUTLAY, 36 MONTHS	P/R TODAY	HOLDS UNTIL	HOLD AT MONTH 36
Two-year	$12	$432	0.09	month 30	80%
Yearly, the standard condition	$14	$504	0.10	month 27	68%
Six-month	$17	$612	0.13	month 23	51%
Monthly	$19.99	$720	0.15	month 20	36%

Every price is taken over the same 36 months so that the plans can be compared.

FIGURE 14

Under $9.98 a month the plan would pass on public evidence



Hold probability at month 36 by monthly price, each point a full simulation. The vertical lines are the four list prices of Premium. The standard reports where the gate is passed. It does not recommend a price.

What the prices show

1. **Every plan is cheaper than its replica today.** Even at $19.99 a month the outlay is 0.15 of the replication price.
2. **The two-year plan is covered by the result.** Its commitment ends at month 24. The value holds at 90% confidence to month 30 at that price.
3. **The gate and the floor are close.** The 36-month test passes under $9.98. The durable price, at which the plan also sits at the median floor, is $9.12.
4. **Evidence raises the passing price.** With crafted code under seal the test passes under $11.10. With proof time it passes at every list price.

SECTION 9

Beyond one buyer: groups, Pro and iCOT

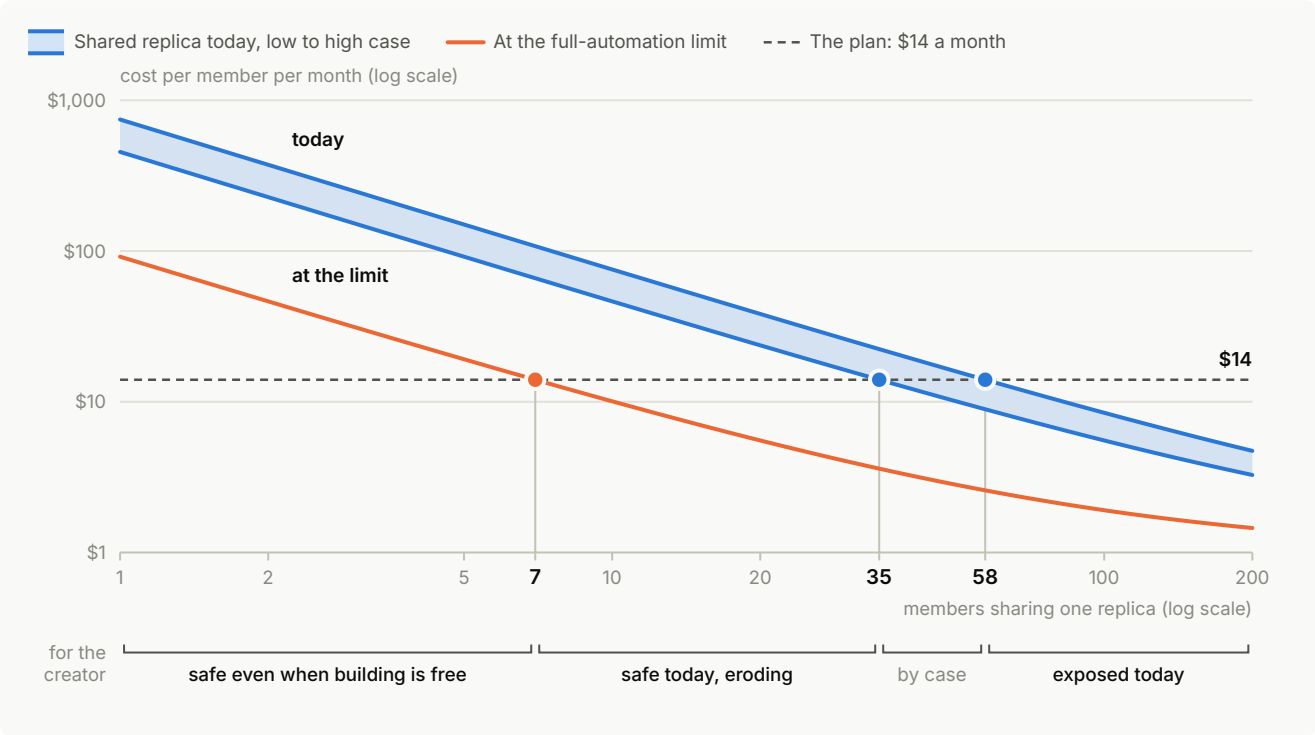
A single subscriber is the easy test. A group sharing one replica reaches the price at 35 to 58 members, the Pro step is priced above its replica, and iCOT stands outside the score.

Groups

Version 0.1 of the standard tests three replicators, and a price is safe only if it passes all three. The group test asks how many members must share one living replica, with hosting, an operator and tracking of new features, before each pays no more than the plan.

FIGURE 15

A shared living replica matches the $14 plan at 35 to 58 members today, and at 7 at the limit



Cost per member per month of one shared living replica over 36 months, from the reference model of version 0.1, low and high case, grade C. Groups too small to reach the lower curve stay safe even when building is free.

What the other tests add

1. **Groups are the nearer exposure.** A few dozen subscribers could share a replica at the same cost per head today, and 7 could at the limit. The personal test does not show this.
2. **Pro is priced above its replica.** The step from Premium to Pro costs 2 to 7 times what its added capabilities would cost a technically able user to add. On the label that is a premium, with a statement of what it buys.
3. **iCOT stands outside the score.** The site says the team reverse-engineered COT data into an intraday indicator whose values do not come from the CFTC.[1] Its algorithm is hidden, so a rebuild cannot be verified from the outside and no replication price is quoted.

4 **iCOT has the profile of scenario C.** If each published report can serve as a public check on an intraday estimate, iCOT has a quality measure, weekly observations and a proof time. A twice-weekly schedule would halve that time. This is an inference from the public description, not a finding.

SECTION 10

What the owners can do

Six levers move the result, and only two of them lift the floor.

TABLE 10 **Levers the standard recognises**

LEVER	WHAT IT DOES	EFFECT ON THIS CASE
Publish a declaration	The owners correct the manifest and name the quality measure. Level A2	Corrects scope. Does not narrow cost
Run the sealed instrument	Measured size, hidden work and finesse replace the guesses. Level A3	Moves the value between month 25 and month 41. Can lower it
Give the product a public quality measure	Without one, hidden finesse counts for nothing	A condition for the next lever
Build what takes time to prove	An estimate or a forecast scored against data published later	Lifts the floor. Passes the gate from 5 to 7 months of proof time
Supply what a replica must keep paying for	Licensed data, live operations, continuous access to the original	Lifts the floor by its running cost
Price	The outlay falls and the replication price does not	Passes under $9.98 a month on public evidence

Runway and floor

A longer runway comes from finesse that an outsider can observe, and AI shortens it again. A higher floor comes from what tokens do not buy: proof that takes calendar time, and inputs that keep costing money. The first two levers measure the product. The next three change it. The last changes only the label.

SECTION 11

Limits and reproduction

The case rests on an outside count, declared unit costs and a forecast of the pace of AI. Every number can be recomputed.

Limits

- The size is one counter’s reading of the public pages. A second count may differ by a quarter.
- No Index exists yet. Tokens and hours per point are declared ranges that bracket a handful of published rebuilds.

- The forward readings depend on a halving time of 6 to 12 months, which is declared, not measured. At month 36 it is half of the uncertainty.
- The floor rests on what a user's attention is worth. For a buyer who values that time at zero it is lower.
- Quality parity is assumed to be checkable on public data. That holds for the core and not for iCOT.
- The sealed findings are invented. A real run could land anywhere on the map of Figure 12.
- The score prices replication only. A competitor who builds something better is outside it.

Reproduction

Three files are attached to this PDF. The value model, `orc_value_model.py`, holds the formulas and the inputs and prints the label's numbers. The script `data.py` runs the sweeps behind the figures, and `case_data.json` is its output. All use the seed 20261113, so a second run gives the same numbers.

REFERENCES

Sources

Numbers match the superscripts in the text, in order of first appearance. The method is the ORC Value Standard, version 0.2, and its predecessor, Open Replication Cost, version 0.1. Every figure in this case study comes from the value model attached to this PDF.

1 COTBase, home page with plans and prices, read 13 November 2026. cotbase.com/

2 COTBase, "Features", read 13 November 2026. cotbase.com/features/

3 COTBase, "Indicators", read 13 November 2026. cotbase.com/indicators/

4 US Commodity Futures Trading Commission, "Commitments of Traders". cftc.gov/MarketReports/CommitmentsofTraders/index.htm

5 World Grain, "Twice-weekly COT reports coming". world-grain.com/articles/23385-twice-weekly-cot-reports-coming

6 Barchart, "Commitment of Traders charts". barchart.com/futures/commitment-of-traders

7 US Bureau of Labor Statistics, "Software developers", Occupational Outlook Handbook. bls.gov/ooh/computer-and-information-technology/software-developers.htm

8 Anthropic, "Building a C compiler with a team of parallel Claudes", Anthropic Engineering, February 2026. anthropic.com/engineering/building-c-compiler

9 Cloudflare, "How we rebuilt Next.js with AI in one week", The Cloudflare Blog, February 2026. blog.cloudflare.com/vinext

10 Epoch AI, "MirrorCode: preliminary results", 2026. epoch.ai/blog/mirrorcode-preliminary-results

11 QSM, "Function point languages table". qsm.com/node/68

12 Gundlach, Lynch, Mertens and Thompson, "The Price of Progress", arXiv 2511.23455. arxiv.org/abs/2511.23455

13 METR, "Time Horizon 1.1", 29 January 2026. metr.org/blog/2026-01-29-time-horizon-1-1/

14 COSMIC, "The basics of COSMIC sizing", on ISO/IEC 19761. cosmic-sizing.org/cosmic-sizing/intro/

TERMS

Glossary

Replication price
What a lawful rebuild of the same thing would cost: build, upkeep and the attention to run it. ORC in the formulas.

Outlay
What the buyer pays over the horizon: price times months.

P/R
Price-to-replication: the outlay divided by the replication price. Below 1, subscribing is cheaper.

Floor
The replication price when tokens and oversight cost nothing.

Hold probability
The share of simulated outcomes in which subscribing is still cheaper at a given date.

Runway
The month in which a replica first becomes cheaper. The label quotes it at 90% confidence.

Hidden work
Function points in the real code that the public manifest missed.

Finesse factor
Total build work divided by the work of the plain function. Grades run from Φ1, plain, to Φ5, singular.

Proof time
The months of public observations needed to show that a replica is as good as the original.

Assurance level
A1 public, A2 declared by the owner, A3 sealed, A4 attested by an independent assessor.

Value Score
The discount to replication price at month 36, at the cautious end of the range. Issued only at 90% confidence.

Durable
A score that also passes at the floor: tokens could be free and the plan would still be cheaper.