

THE SUNDAY NIGHT WORKFLOW

Seven modules on building three scheduled AI systems that prep your week. Market arithmetic, prompt architecture, and scoring-model design — with worked exercises and answer keys.



 **Bogen.AI**

THE SUNDAY NIGHT WORKFLOW

ONE RUN PREPS YOUR ENTIRE WEEK
WEDNESDAY, NOVEMBER 4, 2026 • 9:00 AM ET • 60 MINUTES

- 1 • THE WEEK BRIEF
- 2 • THE CONTENT BATCH
- 3 • THE MONDAY CALL LIST

TYPE 11 NUMBERS | THREE TASKS RUN | REVIEW IN 15 MIN | MONDAY IS READY

How to use this course study

Three systems, and the three skills underneath them: reading market arithmetic, writing a prompt that holds up unattended, and designing a scoring model you can defend.

Most AI training hands you a prompt. A prompt stops working the moment your market, your database or the model changes. This course teaches the reasoning underneath each of the three systems so you can repair and extend them yourself. The prompts are in Appendix C; they are the *last* thing in the book on purpose.

The course map

#	Module	The skill it teaches	Pages
1	How unattended runs work	The three execution models, what a scheduled run can and cannot see, and why most first builds fail silently	3–5
2	Data you may and may not use	Classifying your data against your MLS licence, and detecting proxy variables before they reach a model	6–9
3	Market arithmetic	What each metric measures, the two ways to compute supply, the small-sample problem, and the two-population error	10–17
4	Prompt architecture	The six components of a prompt that survives unattended execution, and the chaining pattern	18–23
5	Scoring-model design	Choosing inputs, setting weights, why caps exist, and how to tune a model that disagrees with you	24–30
6	Scheduling and diagnosis	Building the three tasks, and diagnosing the six ways they fail	31–33
7	The weekly rhythm	Operating the system, and the feedback loop without which it decays	34

EXERCISES AND ANSWER KEYS

Modules 1–5 each end with a worked exercise. **Every answer key is in Appendix A (pages 30–33).** Do the exercise before you look. The arithmetic ones are the ones that matter — if you cannot do them by hand, you cannot tell when the model has got them wrong, and it will.

APPENDICES

- **A** — Answer keys (30–33)
- **B** — Metric reference card (34)
- **C** — The three prompts in full (35–37)
- **D** — Glossary and sources (38)

WHAT YOU NEED

- A ChatGPT or Claude account. The \$0 tier runs the whole workflow — Module 1 explains the limits.
- Your MLS member market report for **one** farm, for the last three weeks, not one.
- A contact export: relationship, purchase date, purchase price, estimated value, last contact date.
- Your Voice Print file from the September 2 session (Module 4 explains why it is load-bearing).
- A calculator. You will use it.

READ THIS FIRST

Nothing in this course sends anything on your behalf. Every system produces a draft in a review queue. No post publishes, no email sends and no call is made until you have read it. The review step is not a safety blanket you can drop once you trust the system — Module 3 shows you an output that is internally consistent, confidently written, and wrong.

EDMUND'S MASTERMIND

THE 3 SYSTEMS

BUILT LIVE • NOVEMBER 4

- 1 THE WEEK BRIEF**
11 numbers in → the market brief + your action list
- 2 THE CONTENT BATCH**
3 posts, a reel script, the sphere email → in your voice
- 3 THE MONDAY CALL LIST**
5 ranked calls, each with its reason and opening line

EB Bogen.AI

AGGREGATES, NEVER LISTING RECORDS

⚡ Why this session matters

You have built the systems. The voice print, the prompt library, the listing machine, the database triggers. Every one of them only runs when you remember to run it — so they run in good weeks and quietly stop in busy ones.

This is the session where they start running without you. Unattended scheduling reached every account tier in August 2026. It was not available to this room earlier in the year.

The seller update that slips to Thursday. The post that never gets written. The client from eleven years ago nobody has rung. Those are not time problems. They are consistency problems, and a schedule is the only known fix for a consistency problem.

🎯 What you'll learn

We will master 3 systems live, together:

1 • The Week Brief — eleven numbers you type in become a client-ready market brief for your farm, plus a private action list that ranks your own listings against the market.

2 • The Content Batch — the same run turns that brief into the week's content: three posts, a reel script, and the sphere email, in your voice, about this week.

3 • The Monday Call List — your database ranked on tenure, equity and contact gaps, returning five calls with the reason and the opening line for each.

MODULE 1

LESSON 1.1

Three execution models, and why it matters which one you are in

Almost every failed build in week one is a person running a task in a context that cannot see what the task needs.

BY THE END OF THIS MODULE YOU CAN

- Name what each execution model can and cannot see
- Predict, before you build, whether a given task will have access to your files
- Explain why a scheduled task has no memory of last week unless you give it one

Model	Can see	Cannot see	Use it for
1 • Plain chat a one-off conversation	Only what you paste into that thread	Your files, your other chats, anything from last week	Testing a prompt once, nothing more
2 • Project chat a chat inside a project	The project's files and instructions, plus the thread	Other projects; other threads in the same project	Developing and debugging your prompts
3 • Scheduled task fires on a cadence	Its own prompt text, and — only if attached to a project — that project's files and instructions	Your screen, your email, last week's run, anything you did not write into it	The three systems, once they work

1.2 The fact that catches everyone: each run starts cold

A scheduled task does not remember its own previous output. Sunday the 11th does not know what Sunday the 4th produced. Every run begins with an empty head and whatever you have handed it.

This has a consequence people discover in week three and misdiagnose as the model "getting worse": **anything you want compared across weeks has to live in a file, not in the conversation.**

THE PRACTICAL RULE

If a system needs last week's number to say something useful this week — and System 1 does, because year-over-year and week-over-week comparison is most of its value — then last week's number must be written into the project as a file. Keep a running log. One row per week, eleven columns. The task reads the file; the file is the memory.

1.3 What "unattended" honestly means here

Two of the three systems are genuinely unattended. System 1 is not, and the course will not pretend otherwise.

System	Needs from you each week	Why
1 • Week Brief	Eleven numbers, typed in	No AI can read your MLS member report for you, and Module 2 explains why you would not want it to
2 • Content Batch	Nothing	It reads System 1's output from the project
3 • Call List	Nothing, if your export is current	It reads the contact file from the project

The honest description: ten minutes of data entry, three systems that run themselves, fifteen minutes of review. Anyone selling "fully automated" here is selling a system that pastes your MLS data into a chatbot.

MODULE 1

LESSON 1.4

What your account actually allows

Two of these limits change what you can build. The rest are detail.

Capability	ChatGPT	Claude
Scheduled recurring prompts	Yes, every tier including \$0	Yes, paid plans (Pro, Max, Team, Enterprise)
Weekly cadence	Yes	Yes — hourly, daily, weekdays or weekly
Runs with your machine shut	Yes	Yes, remotely, since July 2026
Active tasks allowed	3 on \$0 and Go · 5 Plus · 10 Business · 15 Pro	Not published
Choosing the exact hour	Paid only; \$0 runs in a window (morning / afternoon / night)	Set per task
Custom GPT on a schedule	Not supported	Skills and connectors available

Platform documentation, verified October 2026. Allowances change — re-check before you build. Sources, page 38.

1.5 The two that change your design

Three active tasks on the \$0 tier. Hence three systems. This is a real constraint, not a teaching device: if you later add a fourth idea, you are choosing which of these three to turn off. Decide that deliberately rather than discovering it when a task silently stops firing.

A CUSTOM GPT CANNOT BE SCHEDULED

If your work currently lives inside a Custom GPT, that work does not move to a schedule. It has to be rebuilt as **project instructions plus a saved prompt**. These are not the same object and the migration is not automatic.

Transfers: instruction text, knowledge files. **Does not:** conversation starters, name and icon, Actions or API calls, and anything configured through the builder's chat rather than its Configure tab.

1.6 Exact time versus window, and why it matters for chaining

System 2 reads System 1's output. If System 2 runs first it reads nothing and writes generic content from an empty brief — confidently, without erroring.

If you are on...	Chain them like this
A paid tier	Stagger by exact time: System 1 at 6:00pm, System 2 at 6:30pm, System 3 at 6:15pm. Thirty minutes is ample.
The \$0 tier	Order inside a window is not guaranteed. Merge Systems 1 and 2 into one task producing brief then content in a single run — two slots, and the ordering problem disappears.

MODULE 1

REVIEW

Module 1 review

ANSWER FROM MEMORY

1. A scheduled task is attached to a project. Name two things it can see and two it cannot.
2. Why does a weekly comparison require a file rather than a conversation?
3. Which of the three systems is not truly unattended, and why is that a deliberate choice rather than a limitation?
4. You are on the \$0 tier. Explain, in one sentence, why you merge Systems 1 and 2.
5. What part of a Custom GPT does *not* transfer when you rebuild it as a project?

EXERCISE 1 — PREDICT THE FAILURE

An agent builds System 3 as a scheduled task in a plain chat — not attached to a project — with their contact export pasted into the prompt itself. It works perfectly the first Sunday. Four weeks later it is surfacing people they have already closed.

1. Why did it work the first week?
2. Why is it wrong in week four?
3. What is the one-word fix, and what is the better fix?

Answer key: Appendix A, page 30.

Set up your project now

Everything after this point assumes the project exists. Build it before Module 2.

1. Create a project named **Sunday Night Workflow**.
2. Upload your **Voice Print** file and your **contact export**.
3. Create an empty file called **week-log.csv** with a header row of the eleven metric names plus a date column. This is the memory from Lesson 1.2.
4. Put your market name, farm name, farm size and banned-phrase list into the project instructions, so no prompt has to repeat them.
5. Add this line to the instructions, verbatim: *Never request addresses, MLS numbers or listing records. Work only from the aggregate figures provided.* Module 2 explains why that sentence is load-bearing.

YOUR SETUP RECORD

Platform and tier I am on

Active scheduled tasks my tier allows

Exact time, or window only?

Therefore: three separate tasks, or merge 1 and 2?

Anything of mine currently stuck inside a Custom GPT

Before Module 2: the project exists, the three files are in it, and you know whether you are building three tasks or two. Everything from here is about what goes inside them.

MODULE 2

LESSON 2.1

The constraint is the licence, not the technology

The model will happily accept your hot sheet. That is not the same as you being allowed to give it.

BY THE END OF THIS MODULE YOU CAN

- Classify any piece of data you hold into one of four categories, and know what each permits
- Run a three-question test on a field before it goes anywhere near a model
- Detect a proxy variable in your own CRM export

2.1 What the rules actually say

Three facts, each from a published source, and together they define the boundary:

- **Redistribution is the prohibited act.** MLS participants may not redistribute listing data to parties the MLS has not authorised, and IDX displays must remain under the broker's actual control.
- **Your licence probably does not mention AI at all.** Industry analysis through 2026 notes most IDX agreements predate language models and contain no sublicensing language covering them. Silence is not permission.
- **Where rules do exist, they are specific.** ARMLS, for example, permits MLS content in closed non-public AI systems; licenses aggregated data for statistical and analytical purposes; prohibits using even aggregated data for AI training without written authorisation; and requires participants to disclose intended AI use in advance.

NAR's position is that innovation is welcome *provided* use complies with MLS rules and the data licence. The duty sits with you.

2.2 The four-category classification

Sort every piece of data you are tempted to paste:

Category	Examples	What it is	Into a public model?
A — Your own	Your listings, your showings, your pipeline, your notes	Your brokerage's business record	● Yes — subject to your brokerage policy
B — Aggregates you derived	Counts, medians, ratios, months of supply	A statistic about many homes; no individual record	● Yes — this is what the course uses
C — Others' listing records	Hot sheet rows, IDX feed, another broker's listing detail, MLS numbers, addresses	Licensed data belonging to the MLS compact	● No
D — Client personal data	Names, contact details, financial position	Personal information you hold in trust	● Only in a closed system, minimised — see 2.4

THE DESIGN DECISION THAT FOLLOWS

System 1 takes **Category B only**: eleven numbers you read off your member report and type in. There is nothing to redistribute because no record is ever transmitted. System 3 takes Categories A and D, minimised to initials (Lesson 2.4).

This is not legal advice. MLS rules and data licences differ by market and several are being rewritten this year. Confirm your own MLS's position with your broker or MLS staff before you automate anything touching listing data.

MODULE 2

LESSON 2.3

Three questions, asked of every field

A procedure you can run in about fifteen seconds per column, before anything is uploaded.

Q1. WHO OWNS IT? If the answer is "the MLS" or "another broker" → Category C → stop.

Q2. IS IT ONE HOME OR MANY? One identifiable property → a record, not an aggregate → stop.

Q3. COULD I PUBLISH IT? If you could legitimately post the number in a market update, it is safe to compute with.

All three must pass. Q3 is the useful one: "6.1 months of supply" passes. "123 Palm Ave, listed 4/12, reduced twice" fails all three.

2.4 Minimising Category D before it goes in

System 3 needs your database. It does not need your clients' identities. Reduce the export before you upload it:

Field	Keep?	Why
Full name	● Replace with initials	The ranking never needs identity; you re-attach names from your own CRM after
Contact ID	● Keep	Your key for rejoining the output to real people
Email / phone / address	● Drop	Never used by the model; pure exposure
Purchase date and price	● Keep	Tenure and equity both derive from these
Estimated current value	● Keep	The other half of equity
Last contact date, touch count	● Keep	The neglect signal
Free-text notes	● Review first	Notes are where protected characteristics hide. See 2.5.

2.5 Proxy detection: the procedure

Nobody reading this would add a column called "race". The risk is a field that *correlates* with a protected characteristic without naming it. Run this on every column you plan to score:

- 1. Say the column name aloud, then say "...and I am using this to decide who to contact."** If that sentence needs explaining, flag it.
- 2. Ask what else the column predicts.** "Bought a 4-bedroom" predicts household size. "Moved from out of state" can proxy national origin. "Has owned since 1998" is age-adjacent only if you combine it with other fields — tenure alone is fine.
- 3. Check for a substitute.** Nearly always there is a behavioral field that does the same work: replace "likely wants more space" with "viewed larger homes", which is something they actually did.
- 4. Make the model report it.** The System 3 prompt instructs the model to name any column it ignored as a possible proxy. That turns the check into something you receive rather than something you remember.

FREE-TEXT NOTES ARE THE REAL EXPOSURE

A notes column scraped across eleven years will contain "getting divorced", "new baby", "mother moving in", "church group", "wife's visa". Those are protected characteristics in prose. Either strip the notes column, or restrict it to a single curated field — which is exactly what the sample database's *stated_or_behavioral_signal* column is: one short line, written deliberately, nothing inherited.

MODULE 2

LESSON 2.6

Ranking people is a regulated act

System 3 sorts human beings into an order and then treats them differently based on that order. That is exactly the shape of activity Fair Housing law governs.

The rule: personalize on behavior and stated preference, never on protected characteristics. Targeting that sorts or steers contacts by neighborhood demographics, family status or any protected class — even indirectly — can support a steering complaint.

✓ Safe to rank on	⊘ Never rank on
Tenure — how long they have owned	Inferred race, religion, national origin
Equity position on their own property	Familial status or disability
Months since you last made contact	Age, or any field standing in for it
Listings they actually viewed or saved	"Neighborhood fit" logic
Preferences they stated themselves	School-district language used as a family proxy
Purchase anniversary and other time events	Anything you could not explain to a regulator

2.7 Why the distinction is behavior versus inference, not sensitive versus ordinary

The line is not "avoid personal topics". It is about where the information came from.

Defensible

"They told me in July they were thinking about downsizing." You are acting on a statement they made to you, about their own intentions, recorded as they said it.

Not defensible

"They are probably ready to downsize." You inferred it. From what? If the honest answer involves their age or their children leaving home, you have built a protected-class model by accident.

Both sentences could produce the same call. Only one survives the question "why did you contact them and not someone else?"

2.8 The six inputs this course scores on

Complete list, and the reason each one is defensible:

Input	Why it is a behavior or a fact about the property, not a characteristic of the person
Equity gain	A fact about their property's value against what they paid. Says nothing about them.
Tenure	Elapsed time since a transaction you were part of. A date, not a demographic.
Months since contact	A fact about <i>your</i> behavior, not theirs.
Stated or observed signal	Something they said or did. The strongest and most defensible input there is.
Touch neglect	Again a fact about your own conduct.
Relationship type	The role they played in a past transaction with you.

Read that table aloud. Nothing in it describes who a person *is*. Four of the six describe either a property or your own conduct. That is what makes the model explainable, and explainability is the whole defence.

MODULE 2

EXERCISE

Exercise 2 — audit a real export

Below is a column list from an actual CRM export. Classify every column, then decide what to do with it.

#	Column as exported	Category (A/B/C/D)	Keep, drop, or replace?	Safe to score on?
1	contact_full_name			
2	purchase_price			
3	mls_number_of_purchase			
4	years_owned			
5	bedrooms_at_purchase			
6	agent_notes_freetext			
7	last_contact_date			
8	preferred_language			
9	competitor_listing_they_viewed			
10	estimated_current_value			
11	lead_source			
12	household_size_estimated			

THEN ANSWER

1. Three of these columns are hard stops — they must not be scored on under any framing. Which three, and what is the distinct reason for each?
2. One column looks innocuous but is the single most dangerous field in the export. Which, and why?
3. One column is Category C and most people would not notice. Which?
4. For each column you rejected, name the behavioral substitute that does similar work legitimately.

Answer key: Appendix A, page 30.

MODULE 2 REVIEW

1. State the three-question field test without looking.
2. Why is "silence in your IDX agreement about AI" not the same as permission?
3. What is the difference between a stated signal and an inferred one, and why does only one survive scrutiny?
4. Four of the six scoring inputs describe something other than the person. What do they describe?

MODULE 3

LESSON 3.1

What each of the eleven numbers actually measures

You are about to hand these to a model that will reason from them confidently whether or not they mean what you think.

BY THE END OF THIS MODULE YOU CAN

- State what each metric measures and the one thing that most commonly distorts it
- Compute months of supply two different ways and explain a disagreement between them
- Say how many weeks of data you need before a change is a trend
- Explain why the ask-to-close gap is not a discount

Metric	Measures	What distorts it	Watch for
New listings	Fresh supply entering	Relistings of the same home read as new; seasonal clustering at month start	Same address twice in a quarter
Price changes	Seller capitulation	Counts increases too unless you separate them; one home cutting 3× counts 3×	Record cuts and rises separately
Went pending	Demand actually converting	Pendings that later fall through still counted; contingent vs pending differ by MLS	Your MLS's status definitions
Closed	Completed transactions	Lags the market by 30–60 days — it describes contracts signed last month	Never use alone for "right now"
Active inventory	Current competition	Whether your MLS counts contingent/coming-soon as active	Keep the definition fixed
Median list (active)	What sellers are asking	Skewed by a few luxury outliers in a small farm	Median not mean; note outliers
Median DOM (active)	How long current stock has sat	Only measures <i>unsold</i> homes — survivorship: fast sellers already left	Rises when good stock sells
Median closed (30d)	What buyers paid	A different set of homes from the actives — see Lesson 3.4	Not comparable to list
Close-to-list (30d)	Negotiation room on sold homes	Against <i>final</i> list, so reductions are invisible; a home cut 10% then sold at ask reads 100%	Original vs final list
Months of inventory	Market balance	Two formulas that disagree — Lesson 3.2	Which one your MLS uses
Absorption (3-mo)	Pace of sale	Three-month average smooths real turns as well as noise	Lags a genuine shift

THE TWO THAT MISLEAD MOST OFTEN

Median DOM on actives suffers survivorship: it only describes homes that have *not* sold. A hot market where everything good sells in nine days leaves the stale stock behind, and your median DOM goes *up* while the market gets stronger.

Close-to-list is usually computed against the final asking price. A home listed at \$1.2M, reduced to \$1.0M, sold at \$1.0M reports 100% — a perfect negotiation, concealing a 17% reduction. This is why close-to-list alone cannot tell you whether sellers are getting their price.

MODULE 3

LESSON 3.2

Months of supply, computed two ways

Your MLS reports one. Your model will compute the other. When they disagree, the disagreement is the information.

METHOD A – ABSORPTION BASED

months of supply = active inventory ÷ average monthly closings

worked: $31 \div 5.0 = 6.2$ months

Backward-looking. Uses closings, which describe contracts signed 30–60 days ago.

METHOD B – PENDING BASED

months of supply = active inventory ÷ (pendings this month)

worked: $31 \div (3 \times 4.33 \text{ weeks}) = 31 \div 13.0 = 2.4$ months

Forward-looking. Uses pendings, which describe demand right now — but overstates if fall-through is high.

3.3 Reading the disagreement

In the worked example the two methods give 6.2 and 2.4 months. That is not an error. It is the single most useful signal in the dataset, and here is how to read it:

Pattern	What it means	What you tell a seller
Pending-based much LOWER than absorption-based	Demand is accelerating. Recent closings understate what is happening now.	"The closed numbers you are reading online are describing August. Right now homes are going under contract faster than that."
The two roughly agree	Stable market, no turn in progress	Quote either with confidence
Pending-based HIGHER than absorption-based	Demand is decelerating. The closed figures are flattering.	"Closings still look healthy, but that is last month's market. What is going under contract now has slowed."

WHY THIS IS WORTH THE TEN MINUTES

Every agent in the farm has the same MLS report. Almost none of them compute both methods, because the MLS hands them one number. The disagreement between the two is available to anyone and used by nearly nobody — and it is the difference between describing last quarter and describing this week.

BUT: ONE WEEK OF PENDINGS IS THREE HOMES

Method B above multiplied a single week's three pendings up to a monthly rate. On a 412-home farm that is a very small number doing a lot of work — and it is exactly the error Lesson 3.4 is about. Use a trailing four-week pending count, not one week $\times 4.33$, before you say any of this out loud.

MODULE 3

LESSON 3.4

Four listings is not a trend

This is the lesson that stops you confidently telling a seller something false, in writing, on a Tuesday.

A 412-home farm produces about four new listings and two or three closings a week. At that size ordinary variation looks exactly like a market shift — and a model narrates either with equal confidence.

THE ROUGH TEST

For a weekly count of n events, expect natural wobble of about $\pm\sqrt{n}$.

$n = 4$ new listings $\rightarrow \sqrt{4} = 2 \rightarrow$ anything from **2 to 6** is a normal week

$n = 9$ new listings $\rightarrow \sqrt{9} = 3 \rightarrow$ anything from **6 to 12** is a normal week

$n = 2$ closings $\rightarrow \sqrt{2} \approx 1.4 \rightarrow$ **1 to 3** is a normal week

So "new listings doubled from 2 to 4" is almost certainly noise. "New listings went 4, 4, 5, 9, 11 over five weeks" is a signal.

3.5 What you may and may not claim, by sample size

Weekly count	You may say	You may not say
Under 5	"Four homes came on this week." A count, stated flatly.	Anything comparing it to last week. No percentages. No direction.
5 to 15	Compare to a <i>four-week average</i> , never to last week.	"Up 40% on last week." On these counts that is one extra house.
Any size, 4+ weeks	Direction of a rolling average. This is where trends legitimately live.	—

THE FIX: COMPARE TO A ROLLING AVERAGE, NOT TO LAST WEEK

This is why Module 1 had you create **week-log.csv**. With four weeks of history the brief can say "price cuts have run 6, 5, 7, 6 — steady at about six a week" instead of "price cuts fell 14% this week", which on these counts means one seller changed their mind.

3.6 Which metrics are safe at small n , and which are not

Metric	Weekly reading	Because
Active inventory	● Reliable	A stock count of ~31, not a flow of 4. Large and directly observed.
Months of inventory	● Reliable	Built on inventory plus a 3-month average — both smoothed
Median list, active	● Reliable	Median of 31 homes
Median DOM, active	● Careful	Reliable as a level; survivorship makes the <i>direction</i> treacherous
Median closed & close-to-list, 30d	● Careful	A 30-day window helps, but that is still only ~8 homes here — plus the final-list distortion from 3.1
New listings, pendings, closings, price cuts	● Noisy	Flows of 2–4. Report as counts; trend only on 4-week averages. One indecisive seller can cut three times.

Stocks are reliable, flows are noisy. Anything counting what exists now is trustworthy on a small farm; anything counting what happened this week is not. Build the argument on stocks, use flows as colour.

MODULE 3

LESSON 3.7

The ask-to-close gap is not a discount

The most quoted number in any market update, and the most commonly misread. Get this right and you sound like the only person in the room who has thought about it.

median list (active) = \$1,145,000

median closed (30d) = \$985,000

gap = \$160,000, or 16% above

The error: reading this as "homes sell for 16% under asking".

3.8 Why that reading is wrong

These two medians describe **two different sets of houses**:

- **Median list** describes the 31 homes that have *not* sold. By definition this group over-represents homes the market has already declined at their current price.
- **Median closed** describes roughly eight different homes that *did* sell — which tended to be the better-priced ones.

No individual home appears in both figures. Subtracting one from the other does not produce a discount on anything; it compares the asking prices of the unsold to the achieved prices of the sold.

THE NUMBER THAT ACTUALLY MEASURES DISCOUNT

Close-to-list is the discount figure, because it compares each sold home to its own asking price. At 96.2%, homes that sold did so within 3.8% of ask. That is the same-home comparison; the gap is not.

3.9 So what does the gap mean?

Put the two figures together and a real conclusion appears:

Wrong

"There's a \$160,000 gap, so prices are falling and sellers are taking 16% haircuts."

Right

"Homes that are priced correctly still close within about 4% of ask. The \$160,000 gap exists because the unsold pool is full of homes asking last year's price. The gap measures mispricing, not weakness."

Both statements come from the same two numbers. The second one is defensible, more useful to a seller, and is the argument the Week Brief is instructed to make. It also positions you correctly: you are not the agent delivering bad news, you are the agent explaining why the bad news everyone is reading is misread.

THE SAME ERROR, ELSEWHERE

This two-population mistake recurs constantly: comparing average DOM of actives to average DOM of solds; comparing this year's closed median to last year's list median; comparing your listings' performance to "the market" when your listings are all in one price band. **Before any comparison, ask whether both numbers describe the same homes.** Usually they do not.

MODULE 3

LESSON 3.10

The six derivations, worked

These are what the model is instructed to compute. Do them by hand once so you can audit them forever.

Derivation	Arithmetic	Reading	Reliability
Year-over-year supply	$6.1 \div 4.3 - 1 = +42\%$	42% more competition than last fall	● Stock-based, solid
Cross-check	$31 \div 5.0 = 6.2$ vs reported 6.1	Inputs internally consistent	● Always run this
Ask-to-close gap	$1,145,000 - 985,000 = \$160,000$ (16%)	Mispricing in the unsold pool — not a discount (3.7)	● Small closed base
Negotiation room	$100\% - 96.2\% = 3.8\%$	Priced right, homes still sell near ask	● Final-list caveat (3.1)
Seller pressure	$6 \text{ cuts} \div 4 \text{ new} = 1.5$	More sellers cutting than entering	● Both flows tiny
Net absorption	$(3+2) - 4 = -1$	Inventory shrank this week	● One home either way flips the sign

NOTICE WHAT THE RELIABILITY COLUMN IS TELLING YOU

The two most dramatic-sounding derivations — seller pressure and net absorption — are the two least reliable on a farm this size. Net absorption of -1 means five homes left and four arrived. **One extra listing on Saturday would have made it zero; two would have flipped it positive.**

So in week one: note it, log it, say nothing public. If it stays negative three or four weeks running, *then* you have something — and you will be the only agent in the farm who saw it, because you were the only one logging.

3.11 The hierarchy of confidence

When your brief has to lead with one thing, lead with the most reliable thing that is also interesting:

1. **Year-over-year supply** — stock-based, large base, genuinely meaningful, and almost nobody else quotes it.
2. **The gap plus close-to-list together** — the argument from Lesson 3.9. Interesting and defensible *as a pair*, neither alone.
3. **Inventory level and direction** — reliable, if less dramatic.
4. **Flows** — counts only, as colour. "Six sellers cut their price this week" is a fact. "Price cuts are up 50%" is noise wearing a suit.

This hierarchy is the difference between a brief that is right and a brief that merely sounds right. The model will not impose it — it will happily lead with the most dramatic number you gave it. That is why the prompt in Appendix C explicitly instructs it to lead with the relationship and to flag anything it cannot support.

MODULE 3

EXERCISE

Exercise 3 — four weeks of a different farm

A 680-home farm. Work these by hand. A calculator is fine; the model is not.

Metric	Wk 1	Wk 2	Wk 3	Wk 4
New listings	7	5	9	8
Price cuts	4	6	5	11
Went pending	6	7	5	4
Closed	5	4	6	5
Active inventory	48	47	50	53
Median list, active	\$845k	\$848k	\$839k	\$831k
Median DOM, active	41	44	46	52
Median closed, 30d	\$792k	\$795k	\$788k	\$790k
Close-to-list, 30d	97.8%	97.5%	97.1%	96.4%
Absorption, 3-mo avg	20.0	20.0	19.5	19.0
Months of inventory (reported)	2.4	2.4	2.6	2.8

QUESTIONS

1. **Cross-check week 4.** Compute months of supply from inventory and absorption. Does it match the reported 2.8? What does the result tell you about the inputs?
2. **Price cuts went 4, 6, 5, 11.** Apply the $\pm\sqrt{n}$ test to the week-4 figure. Is 11 inside or outside normal variation against the prior three-week average of 5? What may you say publicly?
3. **Compute the ask-to-close gap for week 1 and week 4.** Did it widen or narrow? Now compute close-to-list change over the same period. Do the two agree about direction, and what does the combination mean?
4. **Median DOM rose 41 → 52 while inventory rose 48 → 53.** Give the straightforward reading, then give the survivorship reading from Lesson 3.1. Which additional single metric would settle it?
5. **Net absorption for each week.** Compute all four. Is the four-week total a signal or noise? Justify with the test from 3.4.
6. **Write the one-sentence lead** for week 4's client-ready brief, using the confidence hierarchy from 3.11. It must be defensible from this data alone.

Answer key: Appendix A, page 31. Question 6 has a model answer with commentary.

MODULE 3

WORKED EXAMPLE

What a good run produces

From the eleven Coral Ridge inputs. Read it as the standard, then read page 17 before you trust it.

PART A — CLIENT-READY

Coral Ridge, this week. Four new homes came on, six sellers cut their price, three went under contract and two closed. Inventory now sits at about six months of supply — a year ago it was closer to four. That is the whole story in one line: there is roughly 40% more competition for the same buyer than there was last fall.

Here is the number that matters if you are thinking about selling. Homes currently asking a median of \$1,145,000 are closing at a median of \$985,000. That is a \$160,000 gap between what sellers hope for and what buyers are actually paying. When a home is priced correctly it still sells close to ask — the typical closing lands within about 4% of list. The homes creating that gap are the ones that came on high and are now working their way down in public.

***If you are selling in the next six months:** the first fourteen days carry the offer. Price into the market you are actually in and you keep your negotiating position. Price above it and you spend the next two months buying it back in reductions, in public, where every buyer can watch.*

***If you are buying:** you have something you did not have a year ago, which is choice and time.*

PART B — AGENT-FACING (NEVER SENT)

96-DOM listing — 38 days past the 58-day market median. This week's price conversation. Evidence: six cuts in the farm this week, inventory up 42% year over year, a \$160,000 ask-to-close gap.

47-DOM listing — approaching the median. Pre-empt it. Show this seller the trend now, so a reduction at day 60 is their idea rather than your bad news.

12-DOM listing — healthy. Four showings this week. Report and reassure.

What I would otherwise have missed: net absorption was negative this week — five properties left, four came on.

3.12 The script Part B hands you

Notice what this lets you say on the 96-day call. Not "we need to drop the price." Instead:

"Six sellers in Coral Ridge reduced this week. Inventory is up 42% over last year. Homes asking \$1.145M are closing at \$985K. We are on the wrong side of that gap, and every week we stay there costs us buyer attention. Here are the two options."

The market made that argument. You are carrying it, not inventing it — which is why it lands as counsel rather than as pressure.

MODULE 3

LESSON 3.13

An output that is fluent, confident and wrong

This came from the same eleven inputs. Four errors. Find them before you read the annotations.

Coral Ridge Market Update

The Coral Ridge market shifted decisively this week. Price reductions surged 50% as sellers scrambled to meet a cooling market, and net absorption turned negative for the first time — a clear signal that inventory is building rapidly.

Homes in Coral Ridge are now selling for 16% below asking price, down sharply from last year. With days on market climbing to 58 days, sellers who do not act quickly risk being left behind in what is becoming a buyer's market.

If you are considering selling, the window is closing. Contact me today for a free home valuation.

#	The claim	Why it is wrong
1	"Price reductions surged 50%"	From the eleven inputs you cannot compute a change in price cuts at all — there is no prior week in the data. The model invented a comparison. Even with a prior week, 6 against 4 is inside $\pm\sqrt{n}$ noise (Lesson 3.4).
2	"Selling for 16% below asking price"	The two-population error from Lesson 3.7. The gap compares unsold asks to sold prices. The actual discount figure is close-to-list: 3.8% , not 16%. This sentence is four times wrong and will be repeated back to you by a seller.
3	"Net absorption turned negative for the first time"	"For the first time" is unsupported — no history was provided. And a -1 reading is within one listing of zero (Lesson 3.10). Presented as decisive; is statistically nothing.
4	"Days on market climbing to 58"	"Climbing" implies a direction nobody supplied. 58 is a level. And per Lesson 3.1, DOM on actives can rise <i>because</i> good homes sold — the survivorship trap.

THE PATTERN BEHIND ALL FOUR

Every error is the model **supplying a direction where you only gave it a level**. "Surged", "climbing", "turned negative for the first time", "down sharply" — four invented trends, because a single week of data contains no trends, and narrative prose wants verbs.

This is not a model defect you can prompt away entirely. It is why the review step is permanent, why the prompt says "do not invent any figure I did not give you", and why week-log.csv exists — once the model has four weeks, its directional language becomes true instead of decorative.

EXERCISE 4 — YOUR REVIEW PASS

Take your own first brief. Underline every verb implying direction or change: rose, fell, surged, cooling, climbing, shifting. For each one, write next to it the two numbers that justify it. **Any verb you cannot support with two numbers gets cut.** Count how many survive. Most agents' first brief loses between four and seven.

MODULE 4

LESSON 4.1

The six components of a prompt that survives unattended

A prompt you supervise can be sloppy — you correct it in the next message. A scheduled prompt gets no second message.

#	Component	What it does	What breaks without it
1	Role + scope	Who it is and what it reasons about	Generic framing; writes for no one in particular
2	Data contract	Exactly what it will receive, and what it must never ask for	It requests your hot sheet, or silently assumes missing fields
3	Derivation instructions	The computations to perform, named explicitly	It summarises your inputs back at you instead of reasoning
4	Interpretation rules	How to read a result once computed — the conditional logic	The page-17 failure: fluent, confident, wrong
5	Output contract	Sections, lengths, audience split, format	Your pricing strategy ends up in the client email
6	Guardrails	What it must not do, and what to do when unsure	Invented figures, invented directions, smoothed-over gaps

4.2 Component 4 is the one almost everybody omits

Most prompt advice covers 1, 3 and 5. Components 2, 4 and 6 separate a prompt that works while you watch from one that works while you sleep.

The interpretation rule from the Week Brief prompt, in full:

"If the ask-to-close gap is wide but close-to-list is still high, say explicitly that this is not falling prices and explain what actually causes the gap."

That sentence is Lesson 3.7 compiled into an instruction. Without it the model takes the dramatic reading. You are not asking for care in general — you are encoding one specific, known failure.

THE TRANSFERABLE METHOD

For any system: **list the three ways the output could be confidently wrong, then write one interpretation rule per failure.** That is where your expertise enters. Anyone can write components 1, 3 and 5; only you know what this output looks like when it is wrong in your market.

4.3 Guardrails that actually work

Weak guardrail	Strong guardrail
"Be accurate."	"Do not state any figure I did not give you. If something cannot be derived from these inputs, say so."
"Use my voice."	"Write in the voice described in the Voice Print file in this project" — plus the file.

Weak guardrails name a virtue; strong guardrails name a test. A model cannot check itself against "be accurate". It can check a sentence against "would this be true in a different market?"

MODULE 4

LESSON 4.4

The Week Brief prompt, by component

The full text is in Appendix C. Here is why each part of it exists.

Component	The actual line	The reasoning
1 Role	"You are my market analyst for [NEIGHBOURHOOD], a farm of [N] homes."	Analyst, not copywriter. Sets the register and makes it reason before it writes.
2 Data contract	"I will give you aggregate figures only. You will never ask me for addresses or listing records."	Module 2 compiled into one sentence. Also stops the helpful-assistant failure of suggesting you paste more.
3 Derivations	"Work out: year-over-year supply change; the gap between median list and median closed; what close-to-list says about correctly priced homes; cuts per new listing; net absorption."	Named explicitly, in order. Unnamed derivations do not happen.
4 Interpretation	"If the gap is wide but close-to-list is still high, say explicitly that this is not falling prices."	Encodes Lesson 3.7. Prevents error #2 on page 17 — the one that gets quoted back at you.
5 Output	"Produce TWO parts. PART A — client-ready, ~250 words, no bullet lists. PART B — for me only."	The audience split. Without it your pricing strategy reaches 400 inboxes.
5b Output	"Finish with a line headed 'What I would otherwise have missed'."	Asks for the overlooked thing explicitly. You get what you name.
6 Guardrail	"Do not invent any figure I did not give you. If something cannot be derived, say so."	The last line, deliberately — recency weighting. This is the one that would have caught three of page 17's four errors.

4.5 What is deliberately missing, and why

The prompt does not say "be insightful", "write compellingly" or "act as a world-class analyst". Adjectives of quality do nothing — the model has no dial for insight. Every instruction in the prompt is either a computation, a conditional, a format, or a prohibition. **All four are checkable.**

EXERCISE 5 — DIAGNOSE A WEAK PROMPT

An agent writes this and schedules it:

"You are an expert real estate market analyst. Every Sunday, analyse my neighbourhood's market data and write a compelling, insightful market update for my clients that positions me as the local authority. Make it engaging and professional."

1. Which of the six components are present, and which are missing?
2. Name three specific failures this prompt will produce on its first unattended run.
3. It will fail the same way every week without erroring. Why is that worse than crashing?
4. Rewrite it with all six components. Aim for 150–200 words.

Answer key: Appendix A, page 33.

MODULE 4

LESSON 4.6

Chaining: why System 2 receives output, not a topic

The architectural choice that makes content specific, and the prerequisite that makes it sound like you.

UNCHAINED "Write three posts about the Coral Ridge market this week."

→ the model has no figures, so it writes plausible-sounding generalities

CHAINED "Using the Week Brief you produced for Coral Ridge this week..."

→ every sentence is anchored to \$160,000, 42%, 1.5 cuts per listing

Generic content is an input problem, not a writing problem. No amount of "be specific" fixes an empty input.

4.7 The specificity test, and why it is mechanical

The strongest line in the System 2 prompt is a test the model can actually apply:

"If a sentence would still be true in a different market, cut it."

Fails the test

"The market is shifting and buyers have more options than they did last year. Now is a great time to make a move."

True in any market, any year.

Passes

"Homes here are asking a median \$1,145,000 and closing at \$985,000 — and correctly priced ones still close within 4% of ask."

False in almost every other neighbourhood.

4.8 The Voice Print is load-bearing, not decorative

"AI writes in your voice" is only true if it has a corpus of your voice. Without that file System 2 produces competent, interchangeable copy — and the agent blames the model rather than the missing input.

BUILD A MINIMUM VOICE PRINT IN FIFTEEN MINUTES

1. Collect six to ten things you genuinely wrote: emails to past clients, listing descriptions you liked, a post that did well. Real ones — not what you wish you wrote.
2. Paste them into one file. Save it into the project.
3. Ask the model: *"Describe my writing voice from these samples in ten specific bullet points. Cover sentence length, formality, how I open, how I close, and words I use often."*
4. **Correct it where it is wrong.** This step is the actual work — the model's first read is usually 70% right and the 30% is what makes you sound like yourself.
5. Save the corrected description at the top of the same file, above the samples.

The corrected description matters more than the samples. Samples show the model what you did; the description tells it what is essential versus incidental.

A VOICE PRINT DOES NOT FIX A BAD BRIEF

If System 1's output is thin, System 2 produces well-voiced thin content. Chaining propagates quality in both directions. Fix the brief first — this is why Module 3 is twice the length of Module 4.

MODULE 4

WORKED EXAMPLE

Four angles off one brief

Same three numbers, four different jobs. Note that none of the four repeats another.

Piece	When	Angle	The job it does
Number post	Thursday	The \$160,000 gap, corrected	Argues — makes you the one who read it right
Buyer post	Saturday	Supply as negotiating room	Reframes — reaches the half of the audience the first post ignored
Authority post	Tuesday	The Sunday ritual itself	Demonstrates — proves diligence instead of claiming it
Reel	—	Three numbers as a countdown	Reaches people who will never read 150 words
Sphere email	Tue 7am	Part A, plus two lines	Goes to people who already know you

POST 1 — THE NUMBER POST

Coral Ridge sellers are asking a median of \$1,145,000. Coral Ridge buyers are paying a median of \$985,000.

That \$160,000 gap is the single most important number in this neighborhood right now, and it is not a story about prices falling. Correctly priced homes here are still closing within about 4% of asking.

The gap is made almost entirely of homes that launched at last year's price and are now walking themselves down in public, one reduction at a time. Six of them cut this week.

If you are listing before spring: the first two weeks carry your offer. Price into the market you are actually in, and you keep your leverage.

Structure worth copying: **two numbers → the gap → the correction → what to do**. The correction is the whole post. Without it this is the same market update everyone else posted.

POST 3 — THE AUTHORITY POST

Every Sunday night I run the same report on one neighborhood. New listings. Price changes. Pendlings. Closings. Months of supply against the same week last year. It takes about fifteen minutes and it is the reason I can tell a seller what happened in their market this week instead of what happened last quarter.

This week it caught something I would have missed: five properties left the market and only four came on.

You cannot advise anybody on a market you are only reading about.

This post sells the process, not a listing — which is why it is the one that brings referrals rather than enquiries. Note it survives the specificity test only because of the second paragraph.

MODULE 5

LESSON 5.1

Designing a model that ranks people

Half your business comes from your database. This module is about sorting it defensibly.

BY THE END OF THIS MODULE YOU CAN

- Choose scoring inputs and justify each as a behavior or a property fact
- Set weights deliberately, and explain why each cap exists
- Score a contact by hand and audit the model's arithmetic
- Tune a model that is surfacing the wrong people

28%

FROM REPEAT CLIENTS

22%

FROM PAST-CLIENT REFERRALS

9

MEDIAN SIDES A YEAR

The 2026 NAR Member Profile reports that the typical REALTOR® earned 28% of business from repeat clients and 22% from past-client referrals in 2025. Half the business is already in the CRM. The same report shows the split widening with experience — 49% repeat plus 32% referral at sixteen-plus years, against zero on both for agents with two years or less.

5.2 Why a model beats your instinct here

Your attention has a structural bias. Recent clients feel live; old clients feel closed. But the two strongest predictors of a move — **tenure and equity** — both increase with time, which means the people most likely to move are precisely the ones who feel most finished.

THE TEST OF A USEFUL MODEL

A scoring model that returns the five people you would have called anyway has told you nothing and cost you an hour. **If the output never surprises you, the weights are wrong** — you have encoded your own bias into arithmetic and given it the authority of a number. Expect to disagree with two of the five. Those two are what you paid for.

5.3 Choosing inputs: the two tests

Every candidate input must pass both:

Test	The question	Fails look like
Predictive	Does this plausibly relate to whether they will transact?	Lead source from nine years ago. Interesting, not predictive.
Defensible	Is it a behavior, a stated preference, or a fact about the property — never a characteristic of the person? (Module 2)	Household size. Predictive, and disqualified.

Note the asymmetry: an input that is defensible but weakly predictive just adds noise. An input that is predictive but not defensible is a liability no matter how well it performs. **The second test is a veto, not a tiebreaker.**

MODULE 5

LESSON 5.4

The six inputs, and why each weight is what it is

Weights are editorial judgements, not discoveries. What matters is that yours are deliberate and written down.

score =

```
min(equity_gain, 1.50) × 20 ← up to 30 pts
+ 15 if tenure ≥ 7, else 8 if ≥ 5 ← up to 15 pts
+ min(months_since_contact, 24) × 0.8 ← up to 19 pts
+ 25 if stated_or_observed_signal ← 25 pts, the largest single term
+ 10 if touches_12mo = 0 ← 10 pts
+ 5 if they have been a seller ← 5 pts
```

Theoretical maximum ≈ 104. Absolute scores are meaningless; only the ordering matters.

5.5 The reasoning behind each weight

Input	Max	Why that size
Stated signal	25	The largest single term, deliberately. Someone telling you they are thinking of moving outranks every inference you could make. It should be able to lift a contact several places on its own.
Equity gain	30	Highest ceiling, but it takes a 150% gain to max it. Equity enables a move; it does not motivate one. Big range, gentle slope.
Months since contact	19	Deliberately sizeable. This is the correction for your attention bias — it actively pushes forgotten people up.
Tenure	15	A step, not a slope: 7 years is a meaningful threshold, 14 years is not twice as meaningful as 7.
Touch neglect	10	A bonus on top of the contact gap, catching people on a mailing list who have had no real contact.
Seller relationship	5	Small on purpose. A tiebreaker, not a driver.

5.6 Why the two caps exist

CAPS PREVENT ONE VARIABLE EATING THE MODEL

Equity capped at 150%: a client who bought in 1998 might show a 400% gain. Uncapped, that is 80 points and they top the list forever regardless of every other signal. The cap says "beyond 150%, more equity does not make a move more likely" — which is true.

Contact gap capped at 24 months: someone untouched for eight years is not four times more urgent than someone untouched for two. Past about two years the signal saturates — they are cold, and the degree of cold stops mattering.

The general rule: cap any input where the real-world relationship flattens out. Uncapped linear terms let outliers dominate, and your list becomes the same five people every week — which is how these systems get abandoned.

Write your weights down with a sentence of reasoning each, as above. In six months you will want to change one, and you cannot sensibly change a number whose justification you have forgotten.

MODULE 5

LESSON 5.7

Score one contact by hand

Do this once and you can audit every list the model hands you.

CONTACT A.D. — THE FULL ARITHMETIC

Given: past client (seller+buyer) · bought 2012-11-19 for \$362,000 · value \$1,220,000 · tenure 14.0 yrs · last contact 2026-01-30 · 1 touch in 12 months · signal: "asked what the house would bring today"

Term	Working	Points	Note
Equity gain	$(1,220,000 - 362,000) \div 362,000 = 2.370 \rightarrow \text{cap}$ 1.50×20	30.0	Cap binds
Tenure	14.0 years ≥ 7	15.0	Top step
Months since contact	9.1 months $\times 0.8$	7.3	Under the cap
Stated signal	Signal present	25.0	Full term
Touch neglect	1 touch in 12 months, not 0	0.0	Does not trigger
Relationship	Has been a seller	5.0	Tiebreaker
TOTAL		82.3	Rank 1 of 40

5.8 Reading the decomposition

A.D. tops the list on **three** strong terms, not one — equity at the cap, tenure at the top step, a stated signal. Compare rank two:

Contact	Equity	Tenure	Gap	Signal	Neglect	Rel.	Total
A.D.	30.0	15.0	7.3	25.0	0.0	5.0	82.3
P.R.	30.0	15.0	19.2	0.0	10.0	5.0	79.2

P.R. gets there differently: no signal, but the contact gap is at its cap and the neglect bonus fires. A.D. asked a question you never answered; P.R. you lost touch with entirely. Different calls — which is why the prompt asks for a reason per contact, not just a ranking.

EXERCISE 6 — SCORE THESE THREE

Today is 2026-11-01. Show your working.

Contact	Relationship	Bought	Paid	Value	Last contact	Touches	Signal
R.K.	past client (buyer)	2021-03-15	\$710,000	\$905,000	2026-09-02	4	none
T.W.	past client (seller)	2014-06-20	\$405,000	\$1,015,000	2024-08-14	0	none
N.B.	sphere	—	—	—	2026-10-06	6	asked about rates

- Score all three. Which ranks highest?
- N.B. has a signal but no property data. What does that reveal about how the model treats sphere contacts — flaw, or correct?
- T.W. has no signal and no contact in 26 months. Which two terms carry the score, and what do the caps do?
- Instinct says call R.K. — recent and engaged. The model disagrees. Who is right, and what does Lesson 5.2 say?

Answer key: Appendix A, page 33.

MODULE 5

LESSON 5.9

Tuning a model that is surfacing the wrong people

Diagnose by symptom. Change one weight at a time, and write down why.

Symptom	Likely cause	The adjustment
Same five people every week	An uncapped or over-weighted term; or you are not logging outcomes	Check your caps first. Then confirm you are writing call outcomes back — the model cannot know you rang them (Module 7).
All long-tenure, none recent	Tenure plus equity together dominate — they correlate, so you are effectively double-weighting the same thing	Drop tenure's top step from 15 to 10. They are not independent inputs and should not be weighted as if they were.
The list never surprises you	Contact-gap weight too low; you have encoded your own attention bias	Raise the gap multiplier from 0.8 to 1.2. Per Lesson 5.2, a model that only confirms you is worthless.
Sphere contacts never appear	Correct — they have no equity or tenure terms available	Not a bug. If you want sphere ranked separately, run them as a second list rather than distorting the weights.
People you just closed keep appearing	Stale export, or no exclusion rule	Refresh the export; add "exclude anyone who transacted with me in the last 6 months" to the prompt.
A signal from 2019 still scoring	No decay on the signal term	Add: signals older than 12 months score 10 instead of 25. A stated intention has a shelf life.

THE TUNING DISCIPLINE

One change per week, and write down what you changed and why. Change three weights at once and you cannot attribute the improvement. You also cannot explain the model any more, and explainability is the Fair Housing defence from Module 2 — a model you cannot narrate is a model you cannot justify.

5.10 The audit line, and what it measures

The prompt's last section counts contacts meeting the **ripe test** — past client, 7+ years tenure, 60%+ equity gain — and reports how many have been contacted once or less in twelve months.

15	15	\$9.17M
MET THE RIPE TEST	TOUCHED ONCE OR LESS	UNREALIZED EQUITY

Be precise about what those three numbers are. They describe the 40-contact sample database built for this course, not the market. The citable market fact is the NAR one on page 22: 28% repeat plus 22% past-client referral. The sample is constructed so the mechanism is visible on one page — and the mechanism is real even though the figures are illustrative: **the people most likely to move are the ones you have stopped talking to**, because tenure and equity both rise with exactly the time that makes a client feel finished.

MODULE 5

WORKED EXAMPLE

The ranked list

40-contact sample — 28 past clients, 12 sphere. Initials in place of names, per Lesson 2.4.

#	Who	Score	Tenure	Equity	Why they surfaced
1	A.D.	82.3	14 yrs	237%	Asked what the house would bring today
2	P.R.	79.2	13 yrs	208%	No contact in 24 months
3	J.L.	77.7	13 yrs	176%	Downsizing comment in July
4	D.F.	75.3	14 yrs	246%	No contact in 19 months
5	M.N.	73.4	14 yrs	199%	No contact in 17 months

Scores cluster between 73 and 82 — the gaps are small. Treat the top five as a set, not a strict order; the difference between rank 4 and rank 5 is two points and well inside the noise of your own value estimates.

THE OPENING LINES

1 • A.D. — stated signal. Quote it back.

"You asked me a while back what the house would bring today. I finally have a straight answer for you, and it changed this month — asking prices in Coral Ridge are running about \$160,000 ahead of what's actually closing. I'd rather tell you what you'd really get than what you'd like to hear. Ten minutes this week?"

2 • P.R. — no signal, 24-month gap. Own the silence.

"It's been far too long and that's on me. I'm not calling with a pitch — I'm calling because you've been in that house thirteen years and I don't think you know what it's worth now. Want me to send you the real number?"

Both lines come from the same model. They differ because the *reason* differs — which is why the prompt requires a reason per contact. A ranking without reasons produces five identical calls, and identical calls to people in different situations is how you waste the best five names in your database.

Seven more go to drafted email rather than calls. That keeps the call list at five — achievable on a Monday — while still touching twelve people. One draft carried a flag: a contact who mentioned relocation in February, a time-sensitive signal now nine months stale. The flag exists because the prompt asks for it.

MODULE 6

THE BUILD

Building the three tasks

Twenty minutes once. Do it only after the prompts produce output you would use.

DO NOT SCHEDULE IN WEEK ONE

A scheduled bad prompt produces bad output reliably, every week, and you stop reading it by week three. Run by hand until the output is worth automating. This ordering is the single biggest predictor of whether the system is still running in March.

6.1 The three tasks

#	Task	When	Needs from you
1	Week Brief	Sunday evening	This week's eleven numbers, and it appends them to week-log.csv
2	Content Batch	30 min after task 1	Nothing — reads the brief
3	Monday Call List	Sunday evening	Nothing, if the export is current

On the \$0 tier, merge 1 and 2 into a single task (Lesson 1.6) — you cannot guarantee ordering inside a window.

6.2 The six failure modes

Symptom	Cause	Fix
Output ignores your files entirely	Task attached to a plain chat, not the project	Recreate inside the project. The commonest week-one failure.
Content batch is generic	Task 2 ran before task 1 finished — it read an empty brief and did not error	Stagger by 30 min, or merge the two tasks
Call list repeats last week exactly	No outcomes logged; the database is unchanged	Module 7. The system's only maintenance cost.
A task silently stopped firing	You added a fourth and hit the tier limit	Check the active-task count before adding anything
Brief asserts trends in week one	No history in week-log.csv yet	Expected. Add "state levels only until four weeks are logged" to the prompt.
Output quality drifted over months	Project files went stale — old export, outdated Voice Print	Refresh the export monthly; re-check the Voice Print quarterly

Note how many of these fail silently. Four of the six produce confident, well-formatted, wrong output rather than an error. That is the defining hazard of unattended AI, and the only countermeasure is that you read the output every week.

MODULE 7

OPERATING IT

Running the system, and the loop that keeps it alive

Step	Who	Time
Read your market report, type in the eleven aggregates	You	10 min
Three tasks run	Scheduled	Unattended
Review: figures trace, verbs supported, voice right	You	10–15 min
Monday: five calls, log every outcome	You	~45 min + 30 sec/call
Total Sunday time		20–25 min

7.1 The review pass, in three checks

1. **Trace two figures.** Pick two numbers from the brief and find them in your eleven inputs or derive them. If either is invented, re-run and tighten the guardrail.
2. **Underline the direction verbs.** Exercise 4's test: every "rose", "fell", "climbing" needs two numbers behind it. Cut what you cannot support.
3. **Read one post aloud.** If you would not say it that way to a client, the Voice Print needs more samples.

7.2 The feedback loop, which is the whole maintenance cost

LOG EVERY CALL OUTCOME

System 3 reads database fields. If you call A.D. on Monday and write nothing down, next Sunday surfaces A.D. again with the same opening line, and the week after that too. Thirty seconds of logging per call is the entire upkeep of this system.

Minimum to write back: **date contacted, outcome, and any new stated signal.** That third field is how a passing remark in November becomes the top-ranked call in March.

7.3 The 30-day path

Week	What you do
1	Build the project. Run System 1 <i>by hand</i> . Fix the brief until you would send it. Log week 1 in week-log.csv. Schedule nothing.
2	Add Systems 2 and 3, still by hand. Run the ripe test and write the number down. Make the five calls Monday and log every outcome.
3	Now schedule. Confirm all tasks fire. Compare week 3's brief to week 1's — the first comparison you could not make before.
4	Tune one weight (Lesson 5.9), add your banned phrases, and leave it alone. With four weeks logged the brief can finally speak about direction honestly.

The honest promise. Twenty minutes on a Sunday does not transform a 35-hour week. What changes is that the seller update, the post and the eleven-year client all happen on the week your closing falls through and you are not feeling it. That is not a time saving. It is a consistency guarantee, and consistency is what compounds.

ASSESSMENT

COURSE STANDARD

Can you do these fourteen things?

The ones you cannot do point at the module to revisit. Be honest; nobody is marking this but you.

✓	I can...	Module
☐	Name what a scheduled task can and cannot see, and why it has no memory of last week	1
☐	Say which of my three systems is not truly unattended, and why that is deliberate	1
☐	Classify any field into Category A/B/C/D and state what each permits	2
☐	Run the three-question field test from memory	2
☐	Detect a proxy variable and name its behavioral substitute	2
☐	Compute months of supply two ways and interpret a disagreement	3
☐	Apply the $\pm\sqrt{n}$ test and say what I may claim at a given sample size	3
☐	Explain why the ask-to-close gap is not a discount, and name the figure that is	3
☐	Spot an invented direction verb in an AI-written market update	3
☐	Name the six prompt components and what fails when each is missing	4
☐	Write an interpretation rule that encodes a failure I know from my own market	4
☐	Score a contact by hand and audit the model's arithmetic	5
☐	Explain why each cap exists and what happens without it	5
☐	Diagnose a list that never surprises me, and fix it with one weight change	5

MY COMMITMENT

The farm I am running this on

The Sunday I start

My ripe-test number, and the one I will fix first

The week I will schedule (not week 1)

APPENDIX A

EXERCISES 1 & 2

Answer key — Exercises 1 and 2

EXERCISE 1 — PREDICT THE FAILURE (PAGE 4)

- 1. Why it worked the first week.** The contact data was pasted as literal text inside the task prompt, so it was physically present in the instruction. The task needed no file access because the data *was* the prompt.
- 2. Why it is wrong by week four.** That paste is a frozen snapshot. A scheduled task has no memory and, outside a project, no file access — so every Sunday it re-reads the same October data. Closed deals, new contacts and logged outcomes are invisible. The output is not stale-looking; it is confidently current-sounding and a month out of date.
- 3. The fixes.** One word: **project** — attach the task to a project holding the export as a file. Better: that plus refresh discipline, since a file is only as fresh as your last export. Refresh monthly, write call outcomes back weekly (Module 7) — otherwise a file-based system fails the same way, just slower.

EXERCISE 2 — AUDIT A REAL EXPORT (PAGE 9)

#	Column	Cat.	Action	Reason
1	contact_full_name	D	Replace with initials	Ranking never needs identity; rejoin names from your CRM afterwards
2	purchase_price	A/D	Keep · score	Half of the equity calculation
3	mls_number_of_purchase	C	Drop	An MLS record identifier — licensed data, and it does nothing for the model
4	years_owned	A	Keep · score	Tenure. A date difference, not a characteristic
5	bedrooms_at_purchase	D	Do not score	Proxies familial status
6	agent_notes_freetext	D	Strip or curate	Eleven years of prose containing protected characteristics
7	last_contact_date	A	Keep · score	A fact about your own conduct
8	preferred_language	D	Do not score	Proxies national origin
9	competitor_listing_they_viewed	C/A	Reduce	Keep the behavior ("viewed a home in the \$900k band"), drop the listing detail
10	estimated_current_value	A/B	Keep · score	Other half of equity
11	lead_source	A	Keep	Defensible but weakly predictive — noise, not harm
12	household_size_estimated	D	Do not score	Familial status, and "estimated" means you inferred it

- 1. The three hard stops** are #5, #8 and #12 — and the reasons are genuinely distinct: bedrooms proxies *familial status*; preferred_language proxies *national origin*; household_size is familial status *arrived at by inference*, which is the Lesson 2.7 failure in its purest form.
- 2. The most dangerous innocuous-looking column** is #6, agent_notes_freetext. Structured columns are easy to audit because you can read the header. Eleven years of free text will contain "going through a divorce", "expecting in March", "mother moving in", "waiting on her visa" — protected characteristics in prose, invisible to a header scan and fully visible to a model.
- 3. The Category C column most people miss** is #3, mls_number_of_purchase. It looks like harmless metadata about your own past deal. It is an MLS record identifier.
- 4. Behavioral substitutes.** bedrooms → "has viewed larger homes" (something they did). preferred_language → a preference they stated to you and you recorded as a preference. household_size → nothing inferred; only a stated signal such as "mentioned needing more space".

APPENDIX A

EXERCISE 3

Answer key — Exercise 3, the 680-home farm

Q1 — Cross-check week 4. $53 \div 19.0 = 2.789$, reported 2.8. They agree, so the inventory and absorption inputs are mutually consistent. Run this every week: when it *disagrees*, one of the two numbers has been redefined — usually the active count quietly starting or stopping including contingents.

Q2 — Price cuts 4, 6, 5, 11. Prior three-week average = 5.0. $\sqrt{5} = 2.24$, so a normal week runs roughly **2.8 to 7.2**. Week 4's 11 sits about **2.7 standard deviations** above — outside normal variation and worth attention.

What you may say: "Eleven sellers reduced their price this week, against about five in each of the previous three weeks." A count, with its comparison base stated. **What you may not say:** "Price cuts more than doubled" as evidence of a market shift — one outlier week is not a trend, however large. Log it and watch week 5.

Q3 — The gap, and the trap. Week 1: $\$845k - \$792k = \$53k$ (6.7% above). Week 4: $\$831k - \$790k = \$41k$ (5.2% above). The gap **narrowed**. Close-to-list over the same period **fell** from 97.8% to 96.4%.

Do they agree? They appear to contradict — a narrowing gap sounds like convergence, falling close-to-list sounds like weakness. Resolve it by asking *which* median moved: median list fell $\$845k \rightarrow \$831k$ while median closed was flat $\$792k \rightarrow \$790k$. **The gap narrowed because sellers cut their asking prices, not because buyers paid more.** Combined with close-to-list falling, both figures say the same thing: sellers are repricing downward *and* still conceding more at the table. A narrowing gap is not automatically good news, and this is exactly the two-population care from Lesson 3.7.

Q4 — DOM 41→52 while inventory 48→53. Straightforward reading: homes are taking longer and stock is building — softening. Survivorship reading (Lesson 3.1): DOM on actives can rise because fast homes sold and slow ones remain.

Here the survivorship explanation is weak, because inventory rose at the same time — if good stock were clearing, inventory would fall. The settling metric: **median DOM of sold homes**. If solds' DOM is flat while actives' DOM climbs, you are accumulating stale inventory rather than selling quickly. Pendings (6, 7, 5, 4 — falling) point the same way.

Q5 — Net absorption, and a contradiction worth finding. (pending + closed) – new, by week: **+4, +6, +2, +1**, total **+13**.

A positive total means more homes left than arrived, so inventory should have *fallen* by about 13. Inventory *rose* by 5. That is an **18-home discrepancy**, and finding it is the point of the question. The usual explanation: **your MLS keeps pending homes in the active count** until they close, so "went pending" is not an exit. Other candidates are relistings and expired listings returning.

Signal or noise? Each week's figure rests on flows of 4–11 events, so no single week means much. But the direction is consistently downward (+4, +6, +2, +1) across four weeks, and per Lesson 3.4 four-week direction is where trends legitimately live. Treat it as suggestive — and do not quote the absolute values at all until you have resolved the inventory contradiction.

Q6 — Model lead sentence. Confidence hierarchy says lead with stock-based figures: inventory 48→53 (+10%), months of supply 2.4→2.8 (+17%).

"Inventory here has risen from 48 homes to 53 over the past month, lifting months of supply from 2.4 to 2.8 — still firmly a seller's market by any historical measure, but with noticeably more competition than four weeks ago."

Why this and not the price cuts: eleven cuts is the dramatic number, and it is a single noisy week. Inventory and months of supply are stock-based, reliable at this sample size, and the "+17% but still a seller's market" framing is both true and more useful than alarm. Note it also refuses the easy narrative — which is what makes you the one worth listening to.

APPENDIX A

EXERCISE 5

Answer key — Exercise 5

EXERCISE 5 — DIAGNOSE A WEAK PROMPT (PAGE 19)

1. Components present and missing. Present: **1 (role)**, weakly — "expert real estate market analyst" is a job title, not a scope. Partially present: **5 (output)** — "for my clients" implies an audience but gives no sections, length or format. Missing entirely: **2 (data contract)**, **3 (derivations)**, **4 (interpretation rules)** and **6 (guardrails)**.

2. Three specific failures on the first unattended run.

- **It will invent data.** No data contract, and nothing is attached. Supervised, it would ask you for numbers; unattended it cannot ask, so it produces plausible-sounding figures for your neighbourhood.
- **No audience split.** One undifferentiated update, so either you get nothing actionable for yourself, or your pricing read reaches clients.
- **Invented directions.** "Compelling" and "engaging" push toward narrative, and narrative wants verbs — so you get page 17: "surged", "climbing", "cooling", with nothing behind any of them.

3. Why silent failure beats a crash. A crash is visible and gets fixed Monday. This produces confident, wrong output every week, and it goes to clients. Repetition breeds trust — by week six nobody reads it closely, which is when an invented figure gets quoted back to you.

4. Model rewrite — all six components:

You are my market analyst for Coral Ridge, a farm of 412 homes. [1]

I will give you aggregate figures only. Never ask for addresses, MLS numbers or listing records. [2]

THIS WEEK: [the eleven numbers]

Work out: year-over-year change in months of supply; the gap between median list and median closed; what close-to-list says about correctly priced homes; price cuts per new listing; net absorption. [3]

If the gap is wide but close-to-list is still high, say explicitly this is NOT falling prices and that the gap reflects mispriced unsold stock. Do not call any metric rising or falling unless I gave you two time points for it. [4]

PART A, client-ready, ~250 words, plain prose. PART B, for me only: my listings against the market median DOM, plus a line headed "What I would otherwise have missed". [5]

State no figure I did not give you or that cannot be derived. If something cannot be derived, say so. [6]

APPENDIX A

EXERCISE 6

Answer key — Exercise 6

EXERCISE 6 — SCORE THESE THREE (PAGE 24)

Contact	Equity	Tenure	Gap	Signal	Neglect	Rel.	TOTAL
T.W.	30.0	15.0	19.2	0	10.0	5.0	79.2
N.B.	0	0	0.7	25.0	0	0	25.7
R.K.	5.5	8.0	1.6	0	0	0	15.1

Working: T.W. equity $(1,015,000 - 405,000) \div 405,000 = 150.6\% \rightarrow$ capped $1.50 \times 20 = 30.0$ · tenure 12.4 yrs $\rightarrow 15$ · gap 26.6 mo $\rightarrow$ capped $24 \times 0.8 = 19.2$. R.K. equity $27.5\% \times 20 = 5.5$ · tenure 5.6 yrs $\rightarrow 8$ · gap 2.0 mo $\times 0.8 = 1.6$. N.B. gap 0.9 mo $\times 0.8 = 0.7$ · signal 25.

1. T.W. ranks highest at 79.2, more than three times R.K.

2. What N.B. reveals. N.B. has the single largest term in the model (a stated signal, 25 points) and still finishes at 25.7, because equity and tenure are unavailable for a sphere contact. Sphere contacts are structurally capped at roughly 44 points (signal + a maxed contact gap) and can never compete with a past client. **This is correct behaviour, not a flaw** — the model was designed to rank people whose property position you know. But it means a hot sphere lead will sit invisibly at rank 20. Per Lesson 5.9, run sphere as a *separate* list rather than distorting the weights to force them in.

3. What carries T.W., and what the caps do. Equity (30.0, at the cap) and the contact gap (19.2, at the cap), with tenure 15 and the neglect bonus 10 behind them. Both caps bind: raw equity is 150.6%, just over the 1.50 ceiling; the real gap is 26.6 months against a 24-month ceiling. Uncapped, T.W. would score **81.4** instead of 79.2 — the caps cost 2.2 points here, which changes nothing. That is the sign of well-set caps: they are invisible in ordinary cases and only bite on genuine outliers.

4. Instinct versus model. The model is right — the Lesson 5.2 case exactly. R.K. *feels* live: four touches, contacted eight weeks ago. But they are 5.6 years in with 27% equity — well-served, not move-likely. T.W. is twelve years in with 150% equity and has not heard from you in two years. R.K. feels current because you have been talking to them, which is a fact about your attention, not their intentions. **A list that never disagrees with you is your own bias wearing the authority of a number.**

APPENDIX B

KEEP THIS PAGE

Metric reference card

Everything you need at the desk on a Sunday. Formulas, reliability, and the trap in each.

Quantity	Formula	Small-n	The trap
Months of supply (absorption)	active ÷ avg monthly closings	● Solid	Backward-looking by 30–60 days
Months of supply (pending)	active ÷ monthly pendings	● Careful	Overstates if fall-through is high
Year-over-year supply	(MOI now ÷ MOI last year) – 1	● Solid	Needs the same active definition both years
Ask-to-close gap	median list (active) – median closed	● Careful	Two populations — not a discount
Negotiation room	100% – close-to-list	● Careful	Against <i>final</i> list — hides reductions
Seller pressure	price cuts ÷ new listings	● Noisy	One seller can cut three times
Net absorption	(pending + closed) – new	● Noisy	Pendings may still sit in active inventory
Noise band	$n \pm \sqrt{n}$	—	Below $n=5$, report counts only
Equity gain	(value – paid) ÷ paid	—	Only as good as your value estimate

THE SCORING MODEL

```
score = min(equity_gain, 1.50) × 20
+ 15 if tenure ≥ 7 yrs, else 8 if ≥ 5
+ min(months_since_contact, 24) × 0.8
+ 25 if stated or observed signal
+ 10 if touches in 12 months = 0
+ 5 if they have been a seller with me
```

Ripe test: past client, tenure ≥ 7 yrs, equity gain ≥ 60%.

Safe to rank on

Tenure · equity on their own property · months since contact · listings viewed or saved · preferences they stated · purchase anniversary

Never rank on

Inferred race, religion, national origin · familial status · disability · age · "neighborhood fit" · school-district proxies · anything inferred rather than observed

THE THREE-QUESTION FIELD TEST

1. Who owns it? MLS or another broker → stop. **2.** One home or many? One identifiable property → stop. **3.** Could I publish it? If yes, it is safe to compute with.

APPENDIX C

SYSTEM 1

The three prompts

Component numbers in brackets refer to Lesson 4.1. Replace the bracketed placeholders and save each as a task.

System 1 • The Week Brief

You are my market analyst for [NEIGHBOURHOOD], a farm of [N] homes. I will give you aggregate figures only – counts, medians and ratios. You will never ask me for addresses, MLS numbers or listing records. [1,2]

THIS WEEK, ending [DATE]:

New listings: [] • Price changes: [] (median change []%)

Went pending: [] • Closed: [] • Withdrawn/expired: []

Active inventory: [] • Median list (active): \$[]

Median DOM (active): [] • Median closed (30d): \$[]

Close-to-list (30d): []% • Months of inventory: [] (same week last year: [])

Absorption (3-mo avg): [] closings/month

MY OWN LISTINGS: [DOM for each] • Pending: [] • Showings last 7 days: [] • Seller updates due: []

Also read week-log.csv in this project for prior weeks, and append this week's figures to it. [memory, Lesson 1.2]

Work out: year-over-year change in months of supply; the cross-check of active ÷ absorption against the reported months of inventory; the gap between median list and median closed; what close-to-list says about correctly priced homes; price cuts per new listing; net absorption. [3]

INTERPRETATION RULES: [4]

- If the ask-to-close gap is wide but close-to-list is still high, say explicitly that this is NOT falling prices, and explain that the gap reflects mispriced unsold stock rather than a discount on any individual home.
- Do not describe any metric as rising, falling, surging or cooling unless week-log.csv gives you at least two time points for it. With fewer than four weeks logged, report levels only.
- For any weekly count below 5, report the count and do not compare it to another week.
- If the cross-check disagrees with the reported months of inventory, say so and stop treating either as reliable.

PART A – CLIENT-READY. About 250 words, plain language, no jargon, no bullet lists. Lead with the single most important relationship between these numbers, not a recap of them. Close with one short paragraph for sellers and one for buyers, then an invitation to reply for their own street. [5]

PART B – FOR ME ONLY. Rank my listings against the market median DOM and name which needs a price conversation this week, with the evidence to carry into that call. Flag any listing approaching the median. List the seller calls owed. Finish with a line headed "What I would otherwise have missed". [5]

Do not state any figure I did not give you or that cannot be derived from these inputs. If something cannot be derived, say so rather than estimating. [6]

APPENDIX C

SYSTEM 2

System 2 — The Content Batch

Runs after System 1 and reads its output. On the \$0 tier, merge this into the System 1 task (Lesson 1.6).

System 2 • The Content Batch

Using the Week Brief you produced for [NEIGHBOURHOOD] this week, and writing in the voice described in my Voice Print file in this project, produce the week's content. [1,2]

Every piece must be built from THIS week's figures in THIS farm. If a sentence would still be true in a different market, cut it. [4]

1. THE NUMBER POST (LinkedIn + Facebook, ~150 words). Lead with the sharpest number in the brief, then correct the obvious misreading of it, then say what a seller listing before spring should do differently.
2. THE BUYER POST (Instagram + Facebook, ~120 words). Reframe the supply figure from the buyer's side. Do not say prices are falling. Say what the buyer has now that they did not have a year ago.
3. THE AUTHORITY POST (LinkedIn, ~140 words). The subject is the weekly process itself, not the market. Describe what I run every Sunday, then the specific thing it caught this week.
4. THE REEL SCRIPT (32 seconds, vertical). A table: beat, on-screen text, spoken line. Build around the three strongest numbers as a countdown. Last beat is a follow prompt. Add a caption under 25 words.
5. THE SPHERE EMAIL. Subject under 50 characters built on the sharpest number. Body is Part A unchanged, plus one opening line and one closing line inviting a reply. [5]

Rules: no hashtag stacks, no emoji beyond one in the caption, no exclamation marks, no "the market is shifting", [YOUR BANNED PHRASES]. Never state a figure that is not in the brief. If the brief is missing or empty, stop and tell me rather than writing from general knowledge. [6]

APPENDIX C

SYSTEM 3

System 3 — The Monday Call List

System 3 • The Monday Call List

Read the contact file in this project. Score every contact and return my call list for the week beginning [DATE]. [1,2]

SCORE ON THESE SIX INPUTS ONLY: [3]

1. Equity gain – (estimated value - purchase price) ÷ purchase price. Score gain × 20, capped at 150%.
2. Tenure – years in the home. +15 if 7 or more, +8 if 5 to 7, otherwise 0.
3. Months since last contact – +0.8 per month, capped at 24 months.
4. Stated or observed signal – +25 if the signal field contains anything other than "none".
5. Touch neglect – +10 if touches in the last 12 months is zero.
6. Relationship – +5 if they have been a seller with me.

YOU MUST NOT score, sort, group or comment on age, race, religion, national origin, familial status, disability, or any field that could stand in for one of those. If a column in my file looks like a proxy for a protected characteristic, ignore it and tell me which column you ignored and why. [4,6]

Exclude anyone who transacted with me in the last 6 months. Treat a stated signal older than 12 months as worth 10 rather than 25, and say that you have discounted it. [4]

RETURN: [5]

A. THE FIVE CALLS. Ranked, with the score decomposed into its six terms so I can audit the arithmetic. For each: who, tenure, equity gain, the specific reason they surfaced, and a two-to-four sentence opening line written for that person. Where the reason is a long silence, open by owning the silence, not with a pitch. Where they stated something, quote it back. Never open with a market statistic.

B. DRAFT, DON'T CALL. The next seven by score, each to receive the Week Brief plus one personalised line tied to their own tenure and equity. Flag any contact whose stated signal is time-sensitive and now stale, and say how stale.

C. THE AUDIT LINE. Count contacts that are past clients with 7+ years tenure and 60%+ equity gain. Of those, how many have been contacted once or less in 12 months? Give the combined equity gain of that group and one sentence on what the number means.

Do not invent contacts. If a field is blank, treat it as unknown rather than guessing, and say how many records you had to skip. [6]

ONE CHANGE WORTH MAKING IN MONTH TWO

Add to section A: "Compare this week's five against last week's five, which are in call-log.csv, and tell me if the same name has appeared three weeks running." A name that keeps surfacing and never gets called is the model telling you something about your own avoidance, not about the contact.

Glossary

Absorption rate. Closings per month, usually a 3-month average. Divided into active inventory it gives months of supply.

Aggregate. A count, median or ratio describing many homes, carrying no individual record. Category B in Lesson 2.2.

Ask-to-close gap. Median active list minus median closed. Compares two different sets of homes — not a discount.

Cap. A ceiling on a scoring term, used where the real-world relationship flattens.

Close-to-list ratio. Median closed as a percentage of list. The genuine discount figure, because it compares each home to itself.

Months of supply. How long current inventory would take to clear at the current pace. Two formulas — Lesson 3.2.

Net absorption. Properties leaving minus arriving. Check whether your MLS keeps pendings in the active count.

Proxy variable. A field correlating with a protected characteristic without naming it. The main Fair Housing risk in any scoring model.

Ripe test. Past client, 7+ years tenure, 60%+ equity gain. Shorthand for move-likely.

Scheduled task. A prompt running on a cadence without you, starting cold each time.

Survivorship. Bias in any metric computed only on unsold homes — why active DOM can rise in a strengthening market.

Two-population error. Comparing two figures that describe different sets of homes as though they described the same ones.

Voice Print. A file of your own writing plus a corrected description of your voice, used as context so output sounds like you.

Sources

Every sourced figure in this course traces to one of these. Checked October 2026.

PLATFORM CAPABILITY (MODULE 1)

- Anthropic, "Schedule recurring tasks in Claude Cowork," support.claude.com — plan availability, cadences, remote execution with the machine closed.
- OpenAI, "Scheduled tasks in ChatGPT," help.openai.com — active-task limits by plan, flexible windows on the \$0 tier, exact times on paid plans, Custom GPTs unsupported.
- OpenAI Academy, "Tasks" — weekly recurrence, timing, delivery by notification or email.

MLS DATA AND AI (MODULE 2)

- National Association of REALTORS®, "Listing App on AI Platform Can Help Clients But Must Comply With MLS Policy," nar.realtor — innovation permitted provided use complies with MLS rules and the data licence.
- WAV Group, "Stop Giving AI Another Broker's Listing," July 3, 2026 — redistribution limits, broker control of IDX display, and the absence of AI language in current IDX agreements.
- ARMLS Content Access and Use Policy — closed non-public AI systems permitted; aggregated data licensing for statistical and analytical use; AI training prohibited without written authorisation; advance disclosure required.

AGENT BUSINESS DATA (MODULE 5)

- National Association of REALTORS®, 2026 Member Profile, "The Feature Sheet," nar.realtor — 28% of business from repeat clients and 22% from past-client referrals in 2025; nine median transaction sides; median income \$59,200; the 16+ years versus 2-years-or-less split.
- HousingWire, "NAR 2026 member profile shows Realtors more experienced" — median 35-hour week; 13 years median experience.

On the worked examples. The Coral Ridge aggregates, the 680-home farm in Exercise 3, and the 40-contact database are **sample data**, constructed so each mechanism is visible on a single page. They are illustrations and must never be quoted as market statistics. Every sourced figure is cited above. Confirm your own MLS's AI and data-licence position with your broker or MLS before automating anything that touches listing data; nothing in this course study is legal advice.