


Zones at *the Table*

A 16-page companion for noticing,
calculating and choosing.

Name the map.
Read the price.
Size the opportunity.

WSRF Heads-Up Lab · Fern and Rook
A compact reading companion to the 300-page field guide.

Edition 1 · 19 September 2026
40 / 100 BB practice tables · Color and ink-saving editions



A calm sequence for every turn

Step	What to record	What it helps you decide
1 · Visible state	Your cards, board order, seat and stacks	Which information the decision may use
2 · Map	Pocket label, dynamic label and near coverage	How this state is represented
3 · Price	Pot already present and additional call	The immediate continuation threshold
4 · Range	Hands that reach this point and continue	Which equity estimate is relevant
5 · Commitment	Extra cost, effective stack and future plan	Whether the investment is justified
6 · Review	Your reason now; outcome later	What the next experiment should examine

Take these steps in order until they become a familiar rhythm. You can pause the practice table for as long as you like. The map provides a compact description; the price and range connect that description to the action. A good explanation can fit in one sentence: “The price is modest, the continuing range is wide, and this call preserves a useful turn decision.”

Keep your pocket label fixed for whole-hand notes. Recalculate the dynamic map whenever a new street changes its last two board references. The coach shows both, along with the near rank count and an equity estimate. A label change asks you to look again at the public state; it leaves room for either aggression or patience.

Worked question

Can one repeatable sequence help you use an imaginative method more confidently?

Yes. It turns the idea into actions you can perform and records you can compare. The sequence stays useful even when the next hand surprises you.

For a first session, choose only one theme: recognize the references, calculate the price, or explain a sizing choice. Ten complete hand records are enough for a focused practice session. Use the large scheduled evaluation, rather than those ten outcomes, to judge long-run model performance.

The complete pocket calculation

Use A♠ K♠ as the pocket. In the thirteen-rank coordinates, A is 12 and K is 11. The sum modulo thirteen is 10, which names Q; the absolute difference is 1, which names 3. The reference set is therefore A, K, Q and 3.

LIKELY · 7 / 13 near ranks

2 1	3 0	4 1	5 2	6 3	7 4	8 4	9 3	T 2	J 1	Q 0	K 0	A 0
Reference ranks				Unknown-card distance counts				Uniform near share				
A · K · Q · 3				14 / 12 / 8 / 16				52.0%				

Read the rank above each distance. Zero and one are near; two and farther are outside the near region. The four counts are zero-away, one-away, two-away and three-plus-away cards.

The near ranks are the union of each reference and its immediate circular neighbors. Count each rank once, even when several references cover it. With seven near ranks, the pocket is **LIKELY**. This label comes from rank coverage. It is not a statement that the hand wins seven-thirteenths of showdowns.


Changing the suits from suited to offsuit preserves the entire rank map but changes flush opportunities. The model keeps a separate suited feature so the compressed representation can be combined with that ordinary poker information. The same distinction applies after a board arrives: the map and the evaluated hand describe different aspects of the state.

Worked question
Can you calculate the label using a short scratch line?
Yes. Write “12, 11 → 12, 11, 10, 1,” then mark each reference and its neighbors on the thirteen-rank row. Seven marked positions reach the **LIKELY** threshold. The field guide’s atlas lets you check every other rank pair.

At the table, follow the map with a range-and-price sentence. For example, decide which hands would call your chosen opening size and how the stack depth changes later commitment. Strong-looking hole cards still benefit from a plan for the responses they may receive.

An unlikely map can hold strength

Use A♠ A♥. The coordinate pair is 12 and 12. The derived sum modulo thirteen is 11, and the absolute difference is zero. The distinct references are A, K and 2. Reconstruct the near region before reading its label below.



Read the rank above each distance. Zero and one are near; two and farther are outside the near region. The four counts are zero-away, one-away, two-away and three-plus-away cards.

Pocket aces illustrate why a coverage label and poker strength deserve their own columns. The pair’s ordinary preflop strength comes from the hand and the opposing range. The map describes a rank-space construction. An UNLIKELY label therefore creates an opportunity to inspect context, rather than an instruction to treat every such hand as weak.

Suppose a favorable value opportunity appears in this map. Compare the hands that continue, the price of each size and the effective stack. The learned player has ordinary hand, range, board and action features alongside zones, so it can in principle distinguish strong opportunities inside either label. Its actual success is measured by held-out complete returns.

Worked question

Can a strong value opportunity belong to an UNLIKELY map?

Yes. The classification concerns how many rank positions are near. Hand strength and opponent responses determine the value opportunity. Keeping the terms precise makes the method easier to apply rather than less imaginative.

For your notes, write a separate strength description beside the map: premium pair, vulnerable pair, draw, bluff candidate or another concrete range-dependent explanation. Then write the commitment you are considering and the evidence that supports its cost.

Follow the board's changing references

Street	Public board in dealing order	Dynamic source pair	Zone
Preflop	—	A♠ K♠	LIKELY
Flop	7♦ T♣ 2♥	T♣ 2♥	UNLIKELY
Turn	7♦ T♣ 2♥ Q♣	2♥ Q♣	UNLIKELY
River	7♦ T♣ 2♥ Q♣ 9♠	Q♣ 9♠	LIKELY

The pocket map still begins with A♠ K♠ on every street. The dynamic map begins with those cards preflop, then follows the final two public board cards. Flop order matters to this implemented rule. Preserve dealing order in notes and screenshots instead of sorting the board by rank.

UNLIKELY · 6 / 13 near ranks



Reference ranks	Unknown-card distance counts	Uniform near share
2 · Q	6 / 14 / 7 / 19	43.5%

Read the rank above each distance. Zero and one are near; two and farther are outside the near region. The four counts are zero-away, one-away, two-away and three-plus-away cards.

The turn example changes both the public hand texture and the dynamic references. Inspect those changes separately. Ask which made hands and draws the turn introduces, then inspect the new near region. A transition can be a useful cohort for study without being a complete betting recommendation.

Worked question

Can you update the dynamic label while retaining the pocket label for the hand?

Yes. Record both. The pocket label makes whole-hand grouping stable, while the dynamic label follows each decision. Use the latter for action and commitment notes.

A single hand can visit multiple dynamic labels. Keep its final profit in one whole-hand record. If you want to study all-in situations, also record the first all-in decision, its zone and a distinct hand count. This prevents a hand with several actions from being counted as several independent results.

Track the cards that remain

On the example flop, your two hole cards and three public cards are visible. Forty-seven identities remain unknown to you. The opponent’s actual cards and burn cards stay hidden from the policy; their identities are integrated over possible worlds rather than secretly removed.

UNLIKELY · 6 / 13 near ranks

2
0

3
1

4
2

5
3

6
4

7
3

8
2

9
1

T
0

J
1

Q
2

K
2

A
1

Reference ranks	Unknown-card distance counts	Uniform near share
T · 2	6 / 15 / 15 / 11	44.7%

Read the rank above each distance. Zero and one are near; two and farther are outside the near region. The four counts are zero-away, one-away, two-away and three-plus-away cards.

Uniform unknown-pool event	Calculation	Probability
Next sampled card near	21 / 47	44.7%
Two sampled cards both near	$21 \times 20 / (47 \times 46)$	19.4%
At least one of two near	$1 - 26 \times 25 / (47 \times 46)$	69.9%

These are exact calculations for uniform unknown identities. Betting-conditioned opponent pockets can alter future-card removal probabilities. The live equity estimator uses a public-range model for that purpose. A near-card count remains distinct from a winning-out count: some near cards help the opponent, and some far cards help you.

Worked question

Can the simple unknown-pool count remain useful beside a richer range estimate?

Yes. It provides a transparent baseline and a check on visible-card removal. Label its assumptions, then let the range-conditioned analysis answer the showdown question.

When you write a probability in your notebook, include its event: next card near, flush completed, showdown won or opponent folds. Clear event names keep a compact set of numbers from becoming ambiguous at the table.

Keep three readings side by side

Reading	Question it answers	How to use it
Zone coverage	How many rank positions are near these references?	Describe the map and select a study cohort
Range equity	What share of showdowns does this range model predict?	Compare with the price and continuation
Decision confidence	How supported and stable is this action estimate?	Inspect samples, assumptions and alternatives

The table’s equity estimate comes from sampled opponent pockets and board completions. More samples reduce one source of Monte Carlo noise. They do not automatically repair an inaccurate opponent-range assumption. The worked studies use 8,192 samples per state to make the illustration steadier, while the live game uses 128 for responsive play.

A forest’s disagreement between trees is another quantity. It can reveal variation among fitted predictors, but this implementation does not calibrate that spread into a probability that the top action is correct. The book therefore labels predictions as estimated incremental returns and treats action gaps as model preferences. A large gap is useful to inspect, rather than a certified confidence bound.

Each worked study also shifts its equity input by four percentage points in both directions and ranks the same actions again. This is a local input-sensitivity probe. It holds other features fixed, so it is neither a new dealt hand nor a confidence interval. When the favorite action changes, the example deserves closer inspection of the equity and range assumptions.

Performance uncertainty belongs to completed matches. The extended test starts independent match blocks with fresh HUD priors, then summarizes their average returns with Student t intervals across blocks. This captures block-to-block variability more directly than treating every action as independent. Eight blocks still leave uncertainty, and the sensitivity pass uses only four blocks at the larger sampling budget.

Worked question

Can a familiar zone still contain an uncertain decision?

Yes. An unusual public line, deep commitment or poorly estimated continuing range can make a state difficult even when its map is common. Keep the rich state visible when deciding how much confidence to place in the recommendation.

A worked call you can remember

The river pot already including your opponent’s bet is 12 BB. Calling costs 4 BB. Assume the call closes the action and realizes the modeled showdown share. Your additional-return comparison is between folding for zero from this point and calling into a final contest of 16 BB.

4 / (12 + 4) = 25% break-even equity

Assumed equity	Expected incremental call return
20%	$0.20 \times 16 - 4 = -0.80$ BB
25%	$0.25 \times 16 - 4 = 0.00$ BB
35%	$0.35 \times 16 - 4 = +1.60$ BB

Now attach either zone label to the same arithmetic. **LIKELY** does not change the offered price; **UNLIKELY** does not erase a favorable range estimate. The zone features can influence a learned action forecast, but the numerical call baseline remains available as a check on the explanation.

Earlier streets add realization and future betting. If another expensive decision is likely, raw showdown equity can overstate what a call realizes. If position or a strong draw permits productive future action, the continuation may add value. Write the future plan rather than silently treating an early-street call as a river call.

Worked question

Can the 35% scenario support calling under these stated assumptions?

Yes. Its expected incremental return is +1.60 BB, above the zero fold baseline. The conclusion depends on the equity model and closed-action payoff, both of which are explicit here.

A useful live note is “P=12, C=4, price=25%, range estimate=35%, river.” Add the zone and the public line on the next line. This small record contains enough structure to revisit the decision after the outcome has lost its emotional immediacy.

Read a size as a commitment

Suppose you already have 3 BB on the current street and face a target of 7 BB. A call adds 4 BB. A raise to 15 BB adds 12 BB from your stack and increases the target by 8 BB. The button label, the call amount and the extra cost are separate quantities.

Displayed or derived amount	Value
Your existing street contribution	3 BB
Current target	7 BB
Additional call	4 BB
Candidate raise-to target	15 BB
Additional raise cost	12 BB
Increase over current target	8 BB

Compare that 12 BB investment with the money that can still be contested. When one player has a larger stack, excess chips beyond the opponent’s matching capacity cannot earn a call. The engine settles uncalled excess, and the model uses effective-stack features. A large displayed stack is not by itself a reason for a large effective risk.

The forest evaluates its legal presets. The custom slider allows intermediate legal sizes for human play, which can make the game more expressive. A displayed forecast for one preset belongs to that preset. Treat a custom size as its own response question and save interesting examples for later review.

Worked question

Can the same raise-to target cost different players different amounts?

Yes. The cost subtracts what the acting player has already placed on the street. Read both the target and the existing contribution before comparing commitment.

At the table, name the hands that call your larger size and the hands that fold. Then ask whether you welcome those responses. This connects sizing to your range read. The zone map can enrich the state description while the response model gives the investment its purpose.

Study Fern's selective pressure

Fern is the tight-aggressive reference opponent. Relative to Rook, Fern needs stronger estimates for several aggressive actions and uses lower bluff frequencies. Fern still responds to price, position and public betting history; tight describes a relative policy tendency rather than a promise that every bet means a premium hand.

Begin by observing which opportunities produce folds, calls and raises. VPIP measures voluntary preflop participation; PFR measures preflop raising. Their gap can suggest how much participation comes through calls. The displayed statistics use prior counts so the first few hands do not swing the estimates to extremes. An early percentage is partly a prior, and the hands-seen display helps you interpret it.

Against a selective continuing range, a hand that dominated the opening range can face a much closer decision after a large raise. Revisit the range instead of defending an earlier conclusion. A cheaper continuation may preserve flexibility; a fold may avoid a poor additional investment; a strong value hand may welcome the narrower range. The public line distinguishes these situations.

The first completed experiment gave the served zone player a positive point estimate against Fern with an interval spanning zero. Its controlled final-round ablation favored the zone-free fit against Fern. Those results make Fern an especially useful training target. The extended evaluation uses fresh seeded blocks to examine whether the pattern persists, rather than treating more hands as an automatic improvement.

Worked question: Can a difficult matchup show us where the next training improvement should focus? Yes. Group decisions by street, price, stack depth and zone, then compare returns with the same information available to competing models. Let a candidate earn promotion on separate validation data and reserve new matches for its final examination.

Study Rook's wider pressure

Rook is the loose-aggressive reference opponent. Wider participation and more frequent bluffs put more pressure on passive responses. Wider also means a continuing range can contain more hands you beat. The productive adjustment depends on which branches are wide: opening, calling, raising and bluffing can each have different distributions.

A simple reaction of calling every bet gives aggression too much control over price. A simple reaction of raising every likely map discards range strength and commitment cost. Instead, compare the actual offered price with the hands represented by the line. Identify where Rook's pressure adds bluffs and where a repeated raise filters the range toward value.

The original fresh-deal test gave the selected checkpoint a positive interval against Rook. The controlled final-round zone ablation against Rook was less decisive. These are different findings. One concerns the served policy's performance; the other asks what the distance features and zone routing added to a comparable fit. Keeping both visible makes the result more informative.

Try a review session with ten hands chosen before seeing their outcomes. Write one sentence about the opponent range and one about your commitment. Then reveal the result and add a separate outcome sentence. This protects useful learning from the temptation to rewrite the original read after a surprising river. A memorable loss can become an informative example without dominating the entire dataset.

Worked question: Can patience and selective aggression work together against a wider player? Yes. Patience means choosing a justified price, while aggression means charging or pressuring the range when the opportunity supports it. Use the dynamic map as one feature in that explanation and keep the complete line available for review.

A purposeful all-in decision

Question to answer	Useful next step
Which hands call this size?	Build a continuing range, separate from the opening range.
What folds does the bet earn?	Estimate response frequency at this particular size.
What can each player still contest?	Use current contributions and effective remaining stacks.
How stable is the action preference?	Compare nearby equity assumptions and legal alternatives.
What will count as evidence?	Record every scheduled hand and separate forecasts from returns.

Your commitment hypothesis is a meaningful experiment: find circumstances where stronger commitment produces better returns, and where smaller exposure controls losses. Make “circumstances” concrete using zone, price, range, street and stack. A state description rich enough to guide action is also rich enough to test.

For a fresh river bet of B into P , let f be the fold frequency and e the called showdown share. With matching stacks and a call closing the hand, incremental expected return is $fP + (1-f)[e(P+2B)-B]$. The formula makes both parts visible: immediate folds and the price of continuing against the calling range. A raise over an existing bet needs its own accounting.

Worked question

Can a large commitment become a clear, testable decision rather than a reaction to a label?

Yes. Write the continuing range and response assumptions, calculate the cost and compare alternatives. A LIKELY label can be part of that record; the combined evidence is what justifies the stake.

The current trained model can prefer all-in in some printed studies. Those are forecasts to inspect, not a guarantee that every shove is best. The field guide gives both the exact ranked values and the broader fresh-match results so you can examine the individual choice and the policy as a whole.

Use the little brain as a companion

At the table	What you can learn
Zone badge and rank row	Current dynamic classification and near ranks
Pocket label	Stable reference for the whole hand
Equity and price readings	Modeled share beside immediate call threshold
Ranked actions	WSRF's estimated incremental returns
One tree's path	A glimpse into features used by one predictor
Public HUD	Smoothed participation and raise observations
Hand history	A portable record of visible decisions and outcomes

Read your own preferred action before revealing the ranked list. Then compare explanations. A difference can reveal a missed range detail, a mistaken price or a poorly supported model preference. The tree path is a glimpse of one predictor; the entire forest determines the displayed score.

Session statistics use prior observations to stabilize early estimates. The displayed hand count helps you distinguish a mature read from a mostly-prior read. The current browser updates those public observations but serves a fixed WSRF model. Offline refitting belongs to a separate training experiment.

The opponent-card reveal is a practice option. It stays outside the policy's input boundary. You can use it to study a showdown or reconsider a range, while keeping pre-reveal decision notes intact. A fresh hand begins at the selected 40 or 100 BB depth; the session total preserves your running results.

Worked question

Can you learn from the adviser without copying every recommendation?

Yes. Use it as a concrete prediction to investigate. Ask what information it used, what alternatives it ranked and which assumptions deserve the next test. Curiosity becomes more useful when it has a specific forecast to examine.

A printable decision notebook

Fill the first six columns before the result. Add the outcome afterward. Use the note area to describe the opponent range and the reason for your chosen cost. A balanced review includes ordinary hands as well as memorable ones.

Hand	Street / seat	Pocket / dynamic	P / C	Equity / price	Choice / cost	Result
1						
2						
3						
4						
5						
6						
7						
8						
9						
10						

Range and continuation notes

Worked question

Can keeping the decision and outcome in separate fields improve the review?

Yes. It preserves what you believed when acting and what actually happened. That gives you two pieces of evidence instead of letting the result rewrite the original reasoning.

Optional study tags: first all-in, custom size, close forecast, zone transition, unusual public line. Choose one tag at a time for a small review set. Use complete scheduled samples for performance claims.

What the fresh test measured

Served checkpoint	Hands	BB/100	95% block interval
Fern · TAG	160,000	+10.47	+1.67 to +19.27
Rook · LAG	160,000	+18.13	+13.02 to +23.24

The served checkpoint is round three from a three-million-hand adaptive experiment. The new evaluation freezes that player and runs independent match blocks with fresh seeds and public HUD priors. Each main cell has eight blocks of 20,000 hands, with both seats and stack depths balanced.

Controlled representation comparison	Zone fit minus zone-free control	95% block interval
TAG	-24.52 BB/100	-44.17 to -4.88
LAG	-22.83 BB/100	-30.65 to -15.00

The controlled comparison uses final-round models fit to the same replay dataset with equal total tree count. Its negative differences favor the zone-free control in this experiment. The served player’s positive performance and the representation comparison answer different questions. The book keeps both results visible so the next training changes have a clear target.

A separate 256-sample sensitivity pass and all seeded block results appear in the large field guide and downloadable report. The larger sampling pass has four blocks per cell, so several intervals are wide. The opponents are transparent local sparring policies rather than independently benchmarked professional or equilibrium players.

Worked question

Can these results support useful practice while pointing toward better zone training?

Yes. They give us a functioning learned player, inspectable examples and a specific representation question to improve. Future candidates can test richer action evidence, shared learning across zones and new held-out opponents.

Keep this page beside the table

When you notice...	Try this next...
A LIKELY map	Check price, continuing range and effective commitment before increasing size.
An UNLIKELY map	Inspect the full hand strength and offered price; useful value can exist here too.
A zone transition	Rebuild the dynamic references and separately read the changed board texture.
A cheap call	Consider realized equity, possible later prices and what information may arrive.
An attractive shove	Model both folds and the hands that call this size.
A close action ranking	Inspect sensitivity and range assumptions before relying on the gap.
A surprising outcome	Preserve the decision record and review alternatives with comparable data.
A useful new idea	Write a measurable policy change and schedule fresh comparison hands.

Map → price → range → commitment → record

Pocket references: your two hole ranks. Dynamic references: your hole ranks preflop, then the final two public board ranks. Derived set: a , b , $(a+b)$ modulo thirteen, and $|a-b|$. Near means circular distance zero or one. Seven or more near ranks means LIKELY; six or fewer means UNLIKELY.

Call price: additional call divided by the pot after that call. Full-hand return: final stack minus starting stack. Decision forecast: estimated return from the current choice, with earlier contributions treated as sunk. Keep the units and the question next to every number.

Worked question

Can your next session be both playful and purposeful?

Yes. Choose one small theme, keep ten complete records and finish with a question worth exploring. The full 300-page guide is ready whenever you want another map, calculation or worked model decision.

Name the map.
Read the price.
Find your next good decision.