

Survey-Based Strategy Selection

Timothy Darrah, PhD
Mangrove Technologies

October 20, 2026

Abstract

Mangrove translates survey answers into explicit conditions for selecting automated digital-asset trading strategies. Goal and decline preferences determine a risk level; growth ambition and familiarity determine an independent reward level. An unsure decline answer assigns risk level 2 while preserving the user’s uncertainty. Each resulting profile requires a minimum recorded Sharpe ratio, maximum drawdown, minimum gain-to-pain ratio, and minimum return difference against the same asset. The method applies constraint-based recommendation to strategy onboarding through a fully specified answer-to-threshold mapping. Enumeration of all 108 answer combinations confirms coverage of the 25 profile cells. Evaluation of 2,812 unambiguous historical records confirms agreement between strategy placement and the four-condition matching rule in 70,300 comparisons. Three BTC evaluations illustrate a complete selection. These checks establish computational consistency; preference validity and the future performance of selected strategies require separate empirical evaluation.

1 Introduction

Mangrove’s onboarding survey asks users what they follow, what they want from a strategy, how familiar they are with investing, what growth they seek, and how much decline they would tolerate. The selection task is to turn those answers into conditions that a recorded digital-asset strategy can either satisfy or fail. An asset preference restricts the candidate pool; the remaining answers determine the requirements applied to each candidate.

This paper specifies how two independent ordinal levels select four simultaneous performance requirements. A risk-2 default makes an unsure decline answer usable without attributing a numerical loss preference to the user. The analysis tests answer-space coverage, agreement between matching and strategy placement, and a worked historical selection. Preference validity and future performance require separate evaluation.

1.1 Literature Review

Grable and Lytton’s 13-item instrument measures financial risk tolerance through a multistage empirical development process [1]. Mangrove’s shorter survey instead configures strategy-selection requirements. Its scores and default are product policy, without calibration against that instrument or evidence of equivalent measurement reliability. Sonning’s treatment of ordinal response scales further distinguishes numerical coding from evidence of equal distances between responses [2]. Mangrove makes its additive coding explicit so that this assumption can be inspected and tested.

Felfernig and Burke describe constraint-based recommendation as compatibility between user requirements and item properties [3]. Mangrove applies that framework to recorded strategy evaluations. Its contribution is a complete mapping from plain-language answers to two independent requirement levels, four simultaneous thresholds, and a placement rule checked against direct matching. A failed condition excludes a candidate; stronger performance on another metric does not compensate. This formalizes and verifies a domain application of constraint-based recommendation; a comparative user study is needed to establish improvement over another selection method.

Sharpe’s return-to-variability measure supplies one established performance perspective [4]. Mangrove combines its recorded Sharpe field with drawdown, gain-to-pain, and same-asset period outperformance; it introduces neither a new performance metric nor an empirically optimal threshold. Bailey et al. identify the overfitting risk created by searching historical backtests [5]. Four-condition screening does not remove that risk. Historical compliance therefore establishes neither preference validity nor current-market or future performance.

2 Methodology

The method separates survey scoring, profile requirements, and strategy matching. Assets restrict the candidate pool; four preference and experience answers determine two ordinal levels. The selected levels specify four thresholds, which every matching record must satisfy. Table 1 records the quantities used in these steps; Table 3 gives the full threshold policy.

Table 1. Model Parameters, Units, Domains, and Meanings

Scores and levels are ordinal; ratios are dimensionless; pp denotes percentage points.

Parameter	Unit	Domain	Meaning
Asset filter	Ticker	Supported assets or no preference	Restricts candidates without changing scores.
g	Score	1, 2, 3	Goal: consistency, balance, or growth.
f	Score	1, 2, 3	Familiarity: new, some, or experienced.
a	Score	1, 2, 3	Growth ambition: little, modest, or ambitious.
d	Score	1, 2, 3; undefined if unsure	Numerical decline-answer score.
ℓ	%	5, 15, 30; undefined if unsure	User’s stated decline limit.
c	Level	1, 3, 5; 2 if unsure	Decline cap, or assigned default if unsure.
r, w	Level	Integers 1 to 5	User risk and reward levels.
S_s	Ratio	Finite real number	Recorded Sharpe for evaluation s .
D_s	%	0 to 100	Recorded maximum drawdown.
P_s	Ratio	Nonnegative, finite; no sentinel	Recorded gain-to-pain for evaluation s .
R_s, B_s	%	Finite real numbers	Recorded strategy and same-asset returns.
α_s	pp	Finite real number	Same-window difference $R_s - B_s$.
$S_{\min}, P_{\min}$	Ratio	2, 1.75, 1.5, 1.25, 1	Floors selected by risk level.
$D_{\max}$	%	5, 10, 15, 20, 30	Drawdown ceiling selected by risk level.
$\alpha_{\min}$	pp	5, 7.5, 10, 15, 20	Outperformance floor selected by reward.
N	Count	$108 = 81 + 27$	Non-asset combinations: 81 numerical, 27 defaulted.

2.1 Survey Scoring

The survey asks about assets, goals, familiarity, growth ambition, and acceptable decline. Assets filter candidates without changing the profile. Table 2 gives the four scored inputs; Table 1 defines their units and domains. An unsure decline answer assigns risk level 2; it does not supply a numerical decline score.

Table 2. Survey Choices and Scores

Growth labels express everyday comparisons; unsure decline assigns risk level 2 without a numerical answer score.

Input	Score 1	Score 2	Score 3
Goal (g)	A smoother journey	A little of both	More from my assets
Familiarity (f)	I am finding my feet	I have invested before	I test my own ideas
Ambition (a)	A little extra: beat inflation	A modest boost: beat mortgage rates	Something ambitious: beat my credit card
Decline (d)	\$50 down / 5%	\$150 down / 15%	\$300 down / 30%

The growth comparisons express the user’s ambition. They do not set a live inflation, mortgage, or credit-card rate as the strategy benchmark. Strategy outperformance is measured against buying and holding the selected asset over the same evaluation period.

The worked example uses BTC, “A little of both,” “I have invested before,” “A modest boost,” and “\$150 down / 15%.” Thus $g = f = a = d = 2$. Figures 1–3 carry these answers through the profile and strategy selection. Dollar amounts in this worked example use a \$1,000 reference value; the scoring uses the percentage decline.

Risk level sets the permitted historical drawdown and the corresponding Sharpe and gain-to-pain floors. The user’s

decline limit caps this level: $c = 1$ for 5%, $c = 3$ for 15%, and $c = 5$ for 30%. With goal score g and decline score d , risk level r is

$$\underbrace{r}_{\text{risk level}} = \begin{cases} 2, & \text{if the decline answer is unsure,} \\ \min(g + d - 1, c), & \text{if a numerical decline answer is given.} \end{cases} \quad (1)$$

Subtracting 1 maps the two input scores to five levels; the cap prevents the goal from overriding a smaller decline limit. The example gives $r = \min(2 + 2 - 1, 3) = 3$. Choosing 5% instead gives $r = \min(2 + 1 - 1, 1) = 1$.

Reward level controls how much a strategy must have outperformed its asset. With ambition score a and familiarity score f , the user's reward level w is

$$\underbrace{w}_{\text{reward level}} = \underbrace{a}_{\text{ambition score}} + \underbrace{f}_{\text{familiarity score}} - 1. \quad (2)$$

The example gives $w = 2 + 2 - 1 = 3$. Raising ambition from 2 to 3 raises reward to 4 while leaving risk unchanged. Familiarity's contribution is a product hypothesis requiring user validation.

An unsure decline answer assigns $r = 2$ for every goal answer. The original answer remains recorded as unsure, and d , ℓ , and the uncapped sum remain undefined. The default uses the ordinary risk-2 requirements in Table 3 and permits matching and paper exploration. It expresses a starting policy, not a measured preference for a 10% decline.

Survey Answers and Profile Scores

Assets BTC Filter only	What matters most? A little of both Goal score = 2	How familiar? Invested before Familiarity score = 2	Growth ambition A modest boost Ambition score = 2	Comfort with decline \$150 down / 15% Decline score = 2
-------------------------------------	---	--	--	--

Risk = $\min(2 + 2 - 1, 3) = 3$ Reward = $2 + 2 - 1 = 3$

Figure 1. Survey Answers and Profile Scores

The selected scores produce risk 3 and reward 3. The 15% decline answer caps risk at level 3; it does not change the resulting score.

2.2 Profile to Strategy Requirements

Recorded Sharpe summarizes return relative to variability [4]; maximum drawdown is the largest peak-to-trough decline, expressed as a positive percentage. Under Mangrove's existing definition, gain-to-pain compares accumulated positive trade P/L with the absolute accumulated negative trade P/L. These values are read from the strategy record.

For this method, alpha means the strategy's total return minus its same-asset buy-and-hold return over the same run. It is expressed in percentage points (pp). The method uses Mangrove's existing return-difference field or its existing expression:

$$\underbrace{\alpha_s}_{\text{asset outperformance, pp}} = \underbrace{\text{total_return}_s}_{\text{strategy return, \%}} - \underbrace{\text{benchmark_asset_return}_s}_{\text{same-asset return, \%}}. \quad (3)$$

Here s identifies one strategy evaluation record; R_s and B_s in Table 1 denote these same two recorded returns. This alpha is a period return difference; it is neither annualized nor a regression estimate. No new Sharpe, drawdown, or gain-to-pain calculation is introduced.

For evaluation s , Table 1 defines the recorded metrics. Risk level r selects three thresholds from Table 3, while reward level w independently selects the alpha floor. A candidate fits only if

$$\begin{aligned} S_s &\geq S_{\min}(r), & D_s &\leq D_{\max}(r), \\ P_s &\geq P_{\min}(r), & \alpha_s &\geq \alpha_{\min}(w). \end{aligned} \quad (4)$$

All four comparisons must hold. A higher alpha cannot compensate for a failed drawdown limit. Missing or invalid measurements fail selection; the gain-to-pain no-loss sentinel is excluded. Every cell requires at least Sharpe 1.0 and alpha 5 pp.

Table 3. Independent Risk and Reward Requirements

Each risk level combines with every reward level, giving $5 \times 5 = 25$ requirement cells.


Risk level	Sharpe $\geq$	Drawdown $\leq$	Gain-to-pain $\geq$
1	2.00	5%	2.00
2	1.75	10%	1.75
3	1.50	15%	1.50
4	1.25	20%	1.25
5	1.00	30%	1.00

Reward level	1	2	3	4	5
Alpha $\geq$ (pp)	5.0	7.5	10.0	15.0	20.0

To plot a strategy, select the lowest risk row where it passes all three risk measurements, then the highest reward row whose alpha floor it passes. A strategy can qualify for several user cells while having one plotted location. Section 2.3 applies this placement rule to three recorded strategies.

Figure 2 applies the risk-3 / reward-3 requirements. The user occupies one cell; the nine-cell region covers matching strategy placements at risk levels 1-3 and reward levels 3-5.

User Profile and Strategy Requirements

**Figure 2. User Profile and Strategy Requirements**

The left blue border marks the selected profile. The right border encloses all nine compatible requirement cells. “You” identifies the same profile in both views.

2.3 Example Strategies

Matched strategies A (inside bar), B (moving-average cross), and C (balance-of-power cross) are three recorded BTC evaluations. Figure 3 places them in the same requirements matrix; Table 4 lists the supplied measurements used to check each match.




Figure 3. Matched Strategy Placement

A, B, and C occupy cells (2, 3), (2, 4), and (3, 5).

All three strategies pass the user’s four thresholds. They need not occupy the user’s exact cell: a strategy with a smaller historical drawdown and greater outperformance can also fit. The blue outline encloses the same nine-cell region as Figure 2. The matrix condenses four measurements into two coordinates; final selection still checks each original measurement.

Figure 3 and Table 4 show three BTC candidates from the stored Strategy Engine snapshot. All use one-hour data from February 17 through August 17, 2020. Buy-and-hold BTC returned 20.1084% in that window.

The example is a demonstration of the matching rule against recorded results; it does not establish present-market eligibility or deployment readiness.

For A, the recorded return is 34.4086%, so $\alpha_A = 34.4086 - 20.1084 = 14.3002$ pp. Its four tests are $3.0538 \geq 1.5$, $5.416\% \leq 15\%$, $6.2795 \geq 1.5$, and $14.3002 \geq 10$ pp. The other two records pass the same tests. The linked source evaluations are [A, run 722](#), [B, run 426135](#), and [C, run 210893](#). Their full records are retained with the figure source.

3 Application

The historical Strategy Engine snapshot tests compliance with the declared rules. Onboarding applies the saved profile with additional current-market eligibility checks before offering a personal paper-trading copy.

3.1 Answer-Space Coverage

There are three goal choices, three familiarity choices, three ambition choices, and four decline choices including unsure. The total number of answer combinations is

$$\underbrace{N}_{\text{answer combinations}} = 3 \times 3 \times 3 \times 4 = 108. \quad (5)$$

The 81 combinations with a numerical decline answer reach all 25 cells. The remaining 27 assign risk 2 and retain their independently calculated reward level. All 108 combinations therefore produce usable requirements. Reachability does not depend on the default.

For goal scores 1, 2, and 3, the 5% decline cap gives risks 1, 1, 1; the 15% cap gives 2, 3, 3; and the 30% cap gives 3, 4, 5. Among the 81 numerical combinations, risk levels 1 and 3 each receive 27 combinations, while levels 2, 4, and 5 each receive 9. Adding the defaulted answers gives totals of 27, 36, 27, 9, and 9. The unequal counts follow from the caps and default; they are not an arithmetic discrepancy or an observed distribution of users.

3.2 Historical Evaluation and Paper Application

The evaluation snapshot contains 3,000 fetched rows across DOGE, BTC, ADA, XRP, SOL, and ETH, representing 2,834 identities defined by experiment, configuration, run, asset, and dates. Identical repeats are retained once; con-

flicting source records exclude 22 identities. Of the remaining 2,812 evaluations, 2,513 pass the global metric limits and 159 BTC evaluations satisfy the risk-3 / reward-3 example. These are historical evaluations, rather than counts of independent trading rules or present-market approvals.

Direct comparison of all four measurements agrees with strategy placement for every $2,812 \times 25 = 70,300$ record-cell comparison. The worked selection in Figure 3 illustrates this agreement. If the same user instead answers unsure for decline, the default changes risk to 2 while reward remains 3. A and B still qualify; C fails because its 12.88% drawdown exceeds the risk-2 ceiling of 10%. The default thus changes actual eligibility, rather than merely changing the displayed profile.

In onboarding, raw survey answers are saved to the signed-in profile and the requirements are derived from those answers. Current-market and evidence checks further restrict the candidate set before up to three examples are offered. A selected example can then be copied into personal paper trading with simulated funds. The historical analysis above verifies the selection calculation; it does not by itself establish current-market eligibility or successful continuous paper execution.

3.3 Limitations

Computational consistency does not establish preference validity. User studies should test interpretation, repeatability, the decline caps, and familiarity’s role in the reward score. Historical selection also needs consistent periods, fees, data quality, and out-of-sample evaluation. Multiple backtests can inflate apparent performance [5]. Current-regime eligibility and paper activation require separate application checks; registration alone does not establish continuous paper execution or daily performance.

4 Conclusion

Mangrove’s survey-to-strategy method applies constraint-based recommendation to digital-asset onboarding through an explicit, reproducible mapping. Each survey produces independent risk and reward requirements, including risk 2 when the decline answer is unsure. Candidates qualify only when all four supplied measurements pass. Enumeration establishes that all 108 answer combinations are usable and that numerical answers alone reach all 25 cells. The historical evaluation confirms agreement between placement and direct matching across 70,300 comparisons and demonstrates how a changed answer changes eligibility.

This contribution makes the selection policy inspectable and computationally testable. The analysis does not show that the survey improves preference measurement or investment outcomes. Validation should next test how users interpret the questions and default, whether their answers are repeatable, and whether selected strategies remain useful in out-of-sample periods. Those studies would determine whether this application offers an empirical improvement over alternative onboarding and selection methods.

References

- [1] Grable, J. and Lytton, R. H. (1999). Financial risk tolerance revisited: The development of a risk assessment instrument. *Financial Services Review*, 8(3), 163-181. doi:10.1016/S1057-0810(99)00041-4.
- [2] Sonning, L. (2024). Ordinal response scales: Psychometric grounding for design and analysis. *Research Methods in Applied Linguistics*, 3(3), 100156. doi:10.1016/j.rmal.2024.100156.
- [3] Felfernig, A. and Burke, R. (2008). Constraint-based recommender systems: Technologies and research issues. In *Proceedings of the 10th International Conference on Electronic Commerce*, pp. 17-26. doi:10.1145/1409540.1409544.
- [4] Sharpe, W. F. (1966). Mutual fund performance. *The Journal of Business*, 39(1), 119-138. doi:10.1086/294846.
- [5] Bailey, D. H., Borwein, J. M., Lopez de Prado, M., and Zhu, Q. J. (2017). The probability of backtest overfitting. *The Journal of Computational Finance*, 20(4), 39-69. doi:10.21314/JCF.2016.322.