

ANTEdge — The Full Journey, One Deck

From a blind survey of 377 investors to a five-module platform, its system design and performance work, and how every piece of it was tested. A short overview; each companion deck goes deeper module by module.

2

How this is organized

FOUR PHASES, TWO SLIDES EACH

This deck moves through the whole project at a general level: why it was built, what was built, how it was made fast and reliable, and how it was proven to work. Anyone short on time can stop here; anyone who wants the full detail behind any phase can open the matching companion deck.

1

Product Research — why these five modules, not any others

2

Product Build — what the five modules actually do

3

System Design & Optimization — the architecture and the performance work on top of it

4

Testing — how every piece of it was checked before being trusted

3

Product Research — the evidence first

PHASE 1 OF 4 • 1 OF 2 • BEFORE ANY PRODUCT DECISION

Four academic findings on why unaided investors underperform came first, then a 10-question blind survey, no product or feature mentioned, sent to self-directed investors. 377 responded.

89.7%

don't systematically adjust for market conditions

86.7%

report real friction evaluating a company

75.6%

research each pick case by case, not by rule

64.0%

have no rule for what to do when a position falls

The problem: self-directed investors are trying to manage their own money with real discipline, but lack systematic, tested methods for stock selection, position sizing, risk response, market condition awareness, company valuation, and separating a genuinely robust idea from a lucky-looking one.

377 respondents • 10 questions • 8 competitors checked afterward, none covered all five domains

4

Product Research — five modules, five answers

PHASE 1 OF 4 • 2 OF 2 • WHY THESE FIVE, NOT ANY OTHERS

1

Macro Regime (AI/ML)

Answers the 89.7% who don't read the market environment systematically; no checked competitor names a specific regime or timing signal, AI-based or otherwise

2

Factor Strategy

Answers the 75.6% who pick stocks case by case; broad tools show scores, not a disclosed, repeatable method

3

Portfolio & Risk Management

Answers the 64.0% / 54.4% with no rule for loss response or real diversification, tying sizing to actual correlation

4

Long Term Valuation

Answers the 86.7% who report real friction evaluating a company; deep tools exist but assume a finance background

5

Short Horizon Trading

Answers the 67.7% who want robustness evidence, not a headline return; no checked tool discloses that kind of check

5

Product Build — the five modules

PHASE 2 OF 4 • 1 OF 2 • WHAT EACH ONE ACTUALLY DOES

Factor Strategy

Seven factor lenses; five tested against six market regimes, not just averaged. Pick ratios, build a shortlist, backtest against SPY since 2009.

Turns case-by-case picking into one repeatable, tested process.

AT A GLANCE

7 factor lenses**5 of 7 regime-tested**

Portfolio & Risk Management

Three basket diagnostics (volatility, correlation, return-quality spread), then five allocation methods: equal, inverse-vol, optimizer, HRP, regime-adaptive.

Turns basket risk into a chosen allocation, not a default.

AT A GLANCE

5 allocation methods**3 basket diagnostics**

Macro Regime (AI/ML)

An AI/machine-learning model (a Gaussian Hidden Markov Model) classifies the market into one of six regimes live, from thirteen real signals across stocks, bonds, and credit. Overlays on every other module.

The one signal every other module can lean on.

AT A GLANCE

AI/ML model, live**6 regimes, 13 signals**

Long Term Valuation

Independent five-step estimate of what a business is worth: ROIC, WACC, Forecast, DCF, Comps, each flagging what data was missing or estimated.

Separates price from what the business is actually worth.

AT A GLANCE

5-step process**ROIC → WACC → DCF**

Short Horizon Trading

Ten mechanical trading rules tested on real prices, checked against six regimes to see when they actually work, not just whether.

Proof an edge holds up, not just a good-looking backtest.

AT A GLANCE

10 strategies**Tested across 6 regimes**

6

Product Build — how it was learned to build well

PHASE 2 OF 4 • 2 OF 2 • SCOPE FIRST, COURSEWORK AFTER

Every module started as a product decision made first, from the research in Phase 1. The coursework and certifications below were taken afterward, specifically to learn how to build that decision well, not the other way around.

1

Factor Strategy

Quant investing & feature engineering for trading, plus a Georgetown regression course for combining multiple ratios into one ranking

2

Portfolio & Risk Management

Hierarchical Risk Parity and quantitative portfolio management courses, plus Georgetown Corporate Finance for cost of capital and CAPM

3

Macro Regime (AI/ML)

The nearest formal course stopped at basic clustering; the AI/machine-learning model actually used, a Hidden Markov Model, was learned independently, starting on YouTube

4

Long Term Valuation

Georgetown's Valuation, Corporate Finance, Financial Modeling, and Applied Value Investing courses, each mapped to one step of the five

5

Short Horizon Trading


One course per strategy category needed: algorithmic, day trading, event driven, and volatility, plus a shared backtesting-engine course

Two courses sit underneath all five: Linear Algebra & Calculus, and Georgetown Managerial Statistics.

System Design — from five islands to one platform

PHASE 3 OF 4 • 1 OF 2 • ARCHITECTURE, BEFORE AND AFTER


Before — four isolated modules, no API links



* whole-module figures are each module's own average across its measured endpoints (see that module's own before/after pages)

BEFORE • five isolated modules, no shared layer

After — one gateway, shared layer, live API links



* whole-module figures are each module's own average across its measured endpoints (see that module's own before/after pages)

AFTER • one gateway, shared layer, live links

Fourteen separately-started servers with no shared cache or link between them became one gateway in front of all five modules, a shared backend layer, and live AI/ML regime and scoring APIs in place of stale local copies and UI-only links.

8 System Design — four modules, measured end to end

PHASE 3 OF 4 • 2 OF 2 • BUILT ON TOP OF THE SHARED LAYER

Module	Before	After	Stages of change
Factor Strategy	1,107ms cold avg	~8.6ms cold & warm	6: instant lookup, parse once, vectorize, cache fixed history, persist the index, warm everything
Portfolio Management	1,362ms cold avg 1,236ms warm avg	1,005ms cold avg 898ms warm avg	4: warm at startup, parse once, reuse the cluster tree (trade-off disclosed), give the optimizer the exact gradient
Short Horizon Trading	~1,802ms warm avg	~235ms warm avg	4: zero-copy index, vectorize the trade sheet, parse once, warm everything
Long Term Valuation	~23,059ms cold avg	~38ms cold avg	4: collapse a full scan to a snapshot, chase down a second hidden file, fix “warmed” vs “loaded”, zero-copy index & warm everything

The single biggest win: Long Term Valuation's comps page, 28,973ms → 31.31ms cold (about 925x). The one deliberate trade-off — Portfolio Management's HRP cluster reuse — was measured and disclosed, not hidden.

9 Testing — module by module, then across all of them

PHASE 4 OF 4 • 1 OF 2 • UNIT TESTING, THEN INTEGRATION TESTING

Every module was checked in isolation first, then checked again against every other module it now actually calls — real bugs caught both times, not just theoretical risk.

UNIT TESTING • MODULE BY MODULE

- 1 HMM Macro Regime**
An 18-day “Systemic Crisis” streak was actually an unadjusted stock split, not a real crisis — rebuilt from a clean source.
- 2 Factor Strategy — short legs**
A short position was never actually short, no sign flip existed — found in Skewness, confirmed in Momentum too.
- 3 Factor Strategy — annualized return**
The shared formula averaged instead of compounding — fixed platform-wide, all 129 tests re-verified.
- 4 Portfolio Management**
The real worst-ever drawdown was mislabeled by 38 points, a dollar-vs-percent bug shared by all five methods.
- 5 Short Horizon Trading**
The SPY+VIXY hedge was quietly 1.5x levered, never divided by the real capital deployed.
- 6 Long Term Valuation**
36.5% of the ticker universe had no usable ratio table — standardized fallbacks built for all 2,034 tickers.

INTEGRATION TESTING • ACROSS MODULES

- 1 Three modules, one stale regime**
Each kept its own local copy of the market regime; a reproduced cold-start outage caught a real disagreement live. All three rewired to one live source, confirmed to agree after.
- 2 Six broken links a live boot test found**
The gateway's rewrite rules were logic-tested but never booted live — walking all 14 routed sections found six broken links a code read alone missed, all fixed and re-verified hop by hop.
- 3 A dividend mismatch, caught on the first loader**
Migrating SPY's price source silently dropped dividend adjustment; caught before it could propagate into any of the six real downstream consumers.

9 headline bugs like these (6 unit, 3 integration), plus 19 more found and fixed along the way — the full list lives in [Testing_Unit_and_Integration.pdf](#).

10

Testing — real numbers behind the two passes

PHASE 4 OF 4 • 2 OF 2 • WHAT IT ACTUALLY CAUGHT

28

real bugs found and fixed — 6 unit-test headlines, 3 integration-test headlines, 19 more logged

36.5%

of the ticker universe, 742 of 2,034, had no usable ratio table before Long Term Valuation's fix

6 / 6

broken gateway links found by an actual live boot test, not a code read alone

1.5x

a hedge strategy was quietly levered by, undetected until it was actually checked


Nothing here was assumed correct in advance. Every one of these bugs, from an unadjusted stock split mislabeled as a market crisis to a hedge quietly levered 1.5x, was traced to its real root cause and reproduced directly, not patched blind and hoped-fixed.

That same discipline, real behavior checked before a fix, real output verified after, real trade-offs disclosed, is what makes every number in this deck one that can be trusted rather than just asserted.

11 The Timeline — How This Came Together

FROM DOMAIN KNOWLEDGE TO BETA, STAGE BY STAGE

One continuous path: the coursework and domain knowledge that came first, then the build, then the system design and optimization pass, then testing and documentation — and what comes next.



Sixteen months building the foundation, four months building the product, two months on system design, two more on testing and documentation — and beta testing next.

One platform, five modules, four phases of real work

If you want to go deeper on any phase, each one has its own full deck:

Product_Story_Research_Survey_Competitive_Analysis

The research, the survey, the competitive check, and the case for five modules

Product_Walkthrough_Non_Technical / Technical

Each of the five modules, in plain language and in build detail

Architecture_and_Optimization

The architecture evolution and every performance win, module by module, in full

Testing_Unit_and_Integration

Every real bug caught testing each module alone, then testing them together — nine headline stories, told in full

HMM_Macro_Regime_Deep_Dive

The AI/ML-based Macro Regime model's full story, from motivation to six states, live

Error_Correction

The six headline bugs found and fixed, traced to their real root cause

This deck is the short version. Every companion deck above is the long one.