# What Moddy Delivers

Moddy runs over 300 independent predictive models across the NFL, NBA and MLB, and turns their output into data your customers can use: predictions, picks, performance records, ranked boards and written analysis. Each layer builds on the one before it. Moddy builds and runs all of it. Data partners get it under a custom agreement that includes the rights to offer Moddy data to their own customers.

**SIX LAYERS, ONE PLATFORM**

| | | |
|---|---|---|
| 1 | **Models** | Over 300 independent models, one market each |
| 2 | **Predictions** | Projection, probability and a bet or no-bet call |
| 3 | **Picks** | A model's latest prediction on a market plus its full history |
| 4 | **Performance** | Live graded records: accuracy, calibration, ROI, CLV and more |
| 5 | **Boards** | Ranked picks across a chosen set of models |
| 6 | **Briefings** | The Read, The Rundown and The Parlay |

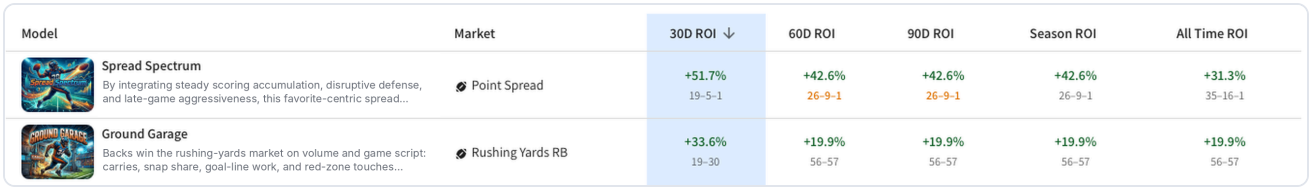
NFL · NBA · MLB | Game lines and player props across 31 markets | Updated every prediction run

## 1 Models

Everything starts with the models. Moddy has built over 300 of them, each a machine-learning model focused on one market, such as an NFL spread, an NBA total or an MLB pitcher's strikeouts.

Each model is its own independent view. Two models on the same market may weigh different factors and reach different answers, and that variety is the point: it lets your customers see where models agree, where they split, and which ones have earned trust.

The models are maintained and run by Moddy. Your team never has to build, train or operate them.



| Model | Market | 30D ROI ↓ | 60D ROI | 90D ROI | Season ROI | All Time ROI |
|---|---|---|---|---|---|---|
| Spread Spectrum<br>By integrating steady scoring accumulation, disruptive defense, and late-game aggressiveness, this favorite-centric spread... | Point Spread | +51.7%<br>19–5–1 | +42.6%<br>26–9–1 | +42.6%<br>26–9–1 | +42.6%<br>26–9–1 | +31.3%<br>35–16–1 |
| Ground Garage<br>Backs win the rushing-yards market on volume and game script: carries, snap share, goal-line work, and red-zone touches... | Rushing Yards RB | +33.6%<br>19–30 | +19.9%<br>56–57 | +19.9%<br>56–57 | +19.9%<br>56–57 | +19.9%<br>56–57 |

Moddy-built models in the business app, with return and record across each performance window

## 2 Predictions

Moddy re-runs every model throughout the day; each pass is a prediction run. In every run, each model re-reads every upcoming game in its market using the latest information: stats, schedule, lineups, injuries and the current market price. For each market on each game, it produces a prediction that says both what the model expects and whether it's worth betting.

Each prediction includes:

- **A projected number**, where the market has one: the projected total, margin or player stat, such as 268 passing yards or a 3-point margin
- **A probability**: the model's chance that its side hits
- **The line and price** it was judged against: the consensus line and price, offered by at least half of the US-regulated sportsbooks Moddy tracks, rather than a single book
- **Expected value**: the model's edge at that price, in profit per unit staked
- **A bet or no-bet recommendation**, with the side and a strength of normal or strong

> **The recommendation is built into the model.** Each model's thresholds for when an edge is worth betting are set during its training and backtesting, not by hand, so every model knows when to pass as well as when to bet.

Predictions stop at game time, so nothing a model says can use information from the game itself.

EXAMPLE · BASED ON A REAL PREDICTION, ABBREVIATED

```
{
  "batchId": "2026-11-08T17:22:26.003Z",      // the prediction run
  "model": { "name": "Air Traffic Control" },
  "domainId": "NFL",
  "marketId": "ocm.passing_yards",
  "entity": { "away": "Los Angeles Rams", "home": "Washington Commanders" },
  "topic": "Jayden Daniels",
  "operatorOdds": "-110.0",     // consensus market price
  "point": 249.5,               // the line
  "prediction": 268.4,          // the model's projected passing yards
  "probability": 0.589,         // chance the over hits
  "expectedValue": 0.125,
  "betAction": "bet_over",
  "betStrength": 2.6            // 1-2 normal, 2+ strong
}
```

268.4 projected yards against a 249.5 line: a 59% chance on the over at -110, a 12.5% edge.

## 3 Picks

A pick is a model's full record on one market in one game: its latest prediction plus every prediction it made before that. Every model has a pick on every market and game it covers, and Moddy serves all of them. A pick opens when the model first recommends a bet.

Because models re-predict right up to game time, a pick shows how the model's view moved as news broke and prices changed: when it first recommended a bet, whether it strengthened or backed off, and where it stood at game time. For an open pick, that final prediction is the one that gets graded. A pick that never opened isn't graded as a bet, but its projection still counts toward the model's accuracy.

**EXAMPLE · BASED ON A REAL PICK, ABBREVIATED**


```
{
  "id": "pick.rpy3GMO4DZAjrIaVpcWWX",
  "marketId": "ocm.passing_yards",
  "entity": { "away": "Los Angeles Rams", "home": "Washington Commanders" },
  "strengthBand": "strong",
  "snapshots": {
    "first": {
      "batchId": "2026-11-04T11:18:38.831Z", "point": 245.5, "operatorOdds": "-114.0",
      "prediction": 258.0, "probability": 0.565, "expectedValue": 0.061, "betStrength": 1.3
    },
    "timeline": [
      { "batchId": "2026-11-07T18:22:08.036Z", "point": 249.5, "operatorOdds": "-112.0",
        "prediction": 266.9, "probability": 0.585, "expectedValue": 0.107, "betStrength": 2.3 }
    ],
    "latest": {
      "batchId": "2026-11-08T17:22:26.003Z", "point": 249.5, "operatorOdds": "-110.0",
      "prediction": 268.4, "probability": 0.589, "expectedValue": 0.125, "betStrength": 2.6
    }
  },
  "totalPredictionCount": 101,
  "predictionChangeCount": 12,
  "entityResult": { "home": 13, "away": 20 },
  "topicResult": 274,
  "result": "won",
  "roi": 90.9
}
```


Opened as a normal bet on the over at 245.5. As the projection climbed to 268 yards, it strengthened to a strong bet and held through kickoff. Daniels threw for 274; the over won.

## 4 Model performance

Every open pick is graded on live results, never backtests, so your customers never have to take a model on faith.

When a game ends, each open pick is graded at the price it had at game time, and voided the way a sportsbook would void it (postponed games, players who don't play). Each model then carries a full record:

| Measure | What it tells you |
|---|---|
| **Accuracy** | How close its projected numbers come to actual results, measured on every prediction, whether or not its pick opened |
| **Calibration** | Whether its probabilities can be taken at face value: when it says 60%, does that side win about 60% of the time |
| **Record** | Wins, losses and pushes |
| **ROI** | Profit as a percentage of the amount staked |
| **Closing line value** | Whether the market moved toward the model's picks between the time it made them and the close, a leading sign of real edge that shows up well before win-loss records settle |
| **Skill vs. market** | Whether its probabilities beat the market's own implied probabilities |
| **Delivered vs. claimed** | Whether its actual returns live up to the edge it claimed when it made its picks |

Every measure is available across several windows, and each window is split by bet strength:

Last 30 days · Last 60 days · Last 90 days · Season to date · All time · × Normal · × Strong

That shows not just whether a model is good, but whether it's good right now and whether its strong picks really are stronger. Performance updates daily. It's what tells your customers which models to trust, and for what.

## 5 Boards

A board turns many models into one product: a single ranked list of the best picks across upcoming games.

You choose which models feed a board. That could be every NBA model, only player-prop models, or a hand-picked group. After every prediction run, Moddy's ranking algorithm gathers those models' live picks, weighs each one by how its model has performed at that strength, removes duplicates, and orders them best first. Ranking is set per board. There is one ranking algorithm today, and Moddy will build another if a partner needs it.

Every board run is kept as it stood, and each board has its own scorecard, so your customers can see what the board said at any point and how it turned out.

EXAMPLE · ILLUSTRATIVE, ABBREVIATED

```
{
  "runId": "2026-10-05T20:20:04.526Z",
  "picks": [
    { "rank": 1, "pickId": "pick.xRk1Cpy-LUMrBE8qyXnmj", "bet": "Falcons ML (+102)",
      "rankScore": 18.8, "strengthBand": "strong", "agreementCount": 2,
      "bandRoi": 21.6, "bandCount": 11, "bandWindow": "season" },
    { "rank": 2, "bet": "Bears/Raiders Under 44.5 (-108)", "rankScore": 15.2,
      "strengthBand": "strong", "agreementCount": 3 }
    // ... every live pick from the board's models, best first
  ]
}
```

## 6 Briefings: The Read, The Rundown and The Parlay

On top of each board, an AI analyst writes The Read and The Rundown after every run, plus The Parlay on boards where it's turned on. It works from the board's top picks and their models, and checks current news such as injuries, lineups and weather.

| Briefing | What it is |
|---|---|
| **The Read** | A one-sentence headline and a short paragraph on what kind of board this is today and why. For example, a board leaning on unders because of cold weather across several parks, plus You Should Know: board-wide watch-outs such as picks that are really the same bet, picks that work against each other, and edges fading as lines move |
| **The Rundown** | The three to five picks most worth attention, each with the case for it and its biggest risk named |
| **The Parlay** | One ticket, usually three legs, built from the board, with why the legs belong together, a safer alternative ticket, and an honest view of its chances against what it pays. Optional, switched on per board |

They're written as model-driven analysis, not guarantees, and are ready to publish as is. Below, from real board runs, are a Read, one Rundown featured pick and a Parlay.

### The Read

EXAMPLE · FROM A REAL BOARD RUN, SHORTENED


```
{
  "runId": "2026-10-04T07:30:55.402Z",
  "summary": "A split-league board mixing NFL Week 4 props with MLB postseason picks,
             leaning heavily toward unders and prop bets rather than sides or totals.",
  "detail": "The NFL side is defined by a sharp lean on quarterback rushing lines ...
            On the MLB side, the picks arrive in postseason context, with hit unders
            in two playoff games where lineup depth drops and pitcher quality rises ..."
}
```


### The Rundown · one featured pick

EXAMPLE · FROM A REAL BOARD RUN, SHORTENED


```
{
  "pickId": "pick.xRk1Cpy-LUMrBE8qyXnmj",
  "rankScore": 18.8,
  "headline": "Falcons moneyline (+102)",
  "summary": "A second model independently agrees on Atlanta tonight, backing the Falcons
             as road underdogs against a Saints team missing its lead running back.",
  "detail": "The model gives Atlanta a 73.8% win probability against a price implying
            roughly 49.5%: that gap is the case. Travis Etienne Jr. has been ruled out ...",
  "riskSummary": "The edge has meaningfully decayed as this opener at +127
                 has been bet down to +102.",
  "riskDetail": "The expected value has dropped from 61.40% at open to 49.20% now ..."
}
```

### The Parlay

EXAMPLE · FROM A REAL BOARD RUN, SHORTENED


```
{
  "intro": "Three legs, two leagues, two game days ...",
  "legs": [
    { "pickId": "pick.1ZwRGJzWBY7sXU8wE1GSD",
      "summary": "Colson Montgomery ... projects roughly 0.3 hits, well below the 0.5 line ...",
      "risk": "Lineup isn't confirmed yet; if Montgomery sits, this leg voids ..." },
    { "pickId": "pick.M3HetcVefDuWkuebbupw9",
      "summary": "Bhayshul Tuten's receiving-yards over ... six independent models ...",
      "risk": "..." },
    { "pickId": "pick.eNKJeGwfvDS3LC7QOc0vJ",
      "summary": "Matthew Stafford's rushing-yards over 0.5 ... betStrength has climbed ...",
      "risk": "..." }
  ],
  "thesis": "These three legs are structurally independent ...",
  "alternative": {
    "legPickIds": ["pick.1ZwRGJzWBY7sXU8wE1GSD", "pick.M3HetcVefDuWkuebbupw9"],
    "rationale": "Dropping the Stafford rushing leg ..."
  },
  "assessment": "... even well-constructed parlays lose far more often than they win ...",
  "pricing": {
    "combinedAmerican": "+784", "combinedDecimal": 8.843, "impliedProbability": 0.1131
  },
  "result": { "outcome": "hit", "profitUnits": 7.84 }
}
```


## What you can sell

Sell one product or all three, as an add-on to the odds and stats feeds you already sell. Your customers get projections, bet calls and a graded record they can check. You don't need a data science team, model training or a grading pipeline.

**Prediction feed**

Projection, probability, consensus odds, expected value and a bet call for every game and market, with each pick's full history

**Performance data**

Each model's live graded record, across windows from the last 30 days to all time

**Boards and briefings**

Ranked picks from the models you choose, with The Read, The Rundown and The Parlay written on top

## Getting it to your customers

You decide what to sell and in what format, and Moddy shapes the data to match, mapped to your own game and player IDs during setup, so it drops into your existing pipelines.

| What you get | |
|---|---|
| **Coverage** | NFL, NBA and MLB: game lines and player props across 31 markets. Other sports can be scoped as a funded project |
| **Freshness** | Updated every prediction run, around the clock, with graded results and performance the next day |
| **History** | Data starts with the current season, so everything reflects today's modeling |
| **Delivery** | REST API, plus webhooks |
| **Operations** | Moddy runs everything behind the data: the models, grading, ranking and analysis, plus monitoring, data quality checks and incident response |
| **Rights** | A custom agreement that includes the rights to offer Moddy data to your own customers |
| **Territory** | Available outside the US, with the partner responsible for local laws. Consensus odds come from US-regulated sportsbooks; non-US pricing is scoped in the pilot |

## How a partnership works

From first call to launch, with a 30-day pilot before you sell.

1

**Scoping call**

Tell us what you want to sell: sports, markets, layers and format.

2

**Custom agreement**

One agreement that includes the rights to offer Moddy data to your own customers.

3

**Setup and 30-day pilot**

Moddy maps the feed to your game and player IDs, then runs a 30-day pilot with a test group, which your team grades for itself.

4

**Decide and launch**

Continue to a one-year agreement, or stop. Once live, you sell it and Moddy keeps running the models, grading, analysis and monitoring.

Pricing is quoted after the scoping call, based on what you plan to sell.

**You define what to sell. Moddy handles the models, the data and the operations behind it.** **moddy.ai**

*Moddy's predictions and briefings are model-driven analysis, not betting advice or a guarantee of results. For adults 21+ only.*