

Arizona vs. Michigan Final Four: Why Michigan Holds the Narrower but Real Edge

Many-Worlds Simulation Report

As-of: 2026-03-31

THE CALL

Michigan wins 57.6% — **Arizona wins 42.4%**

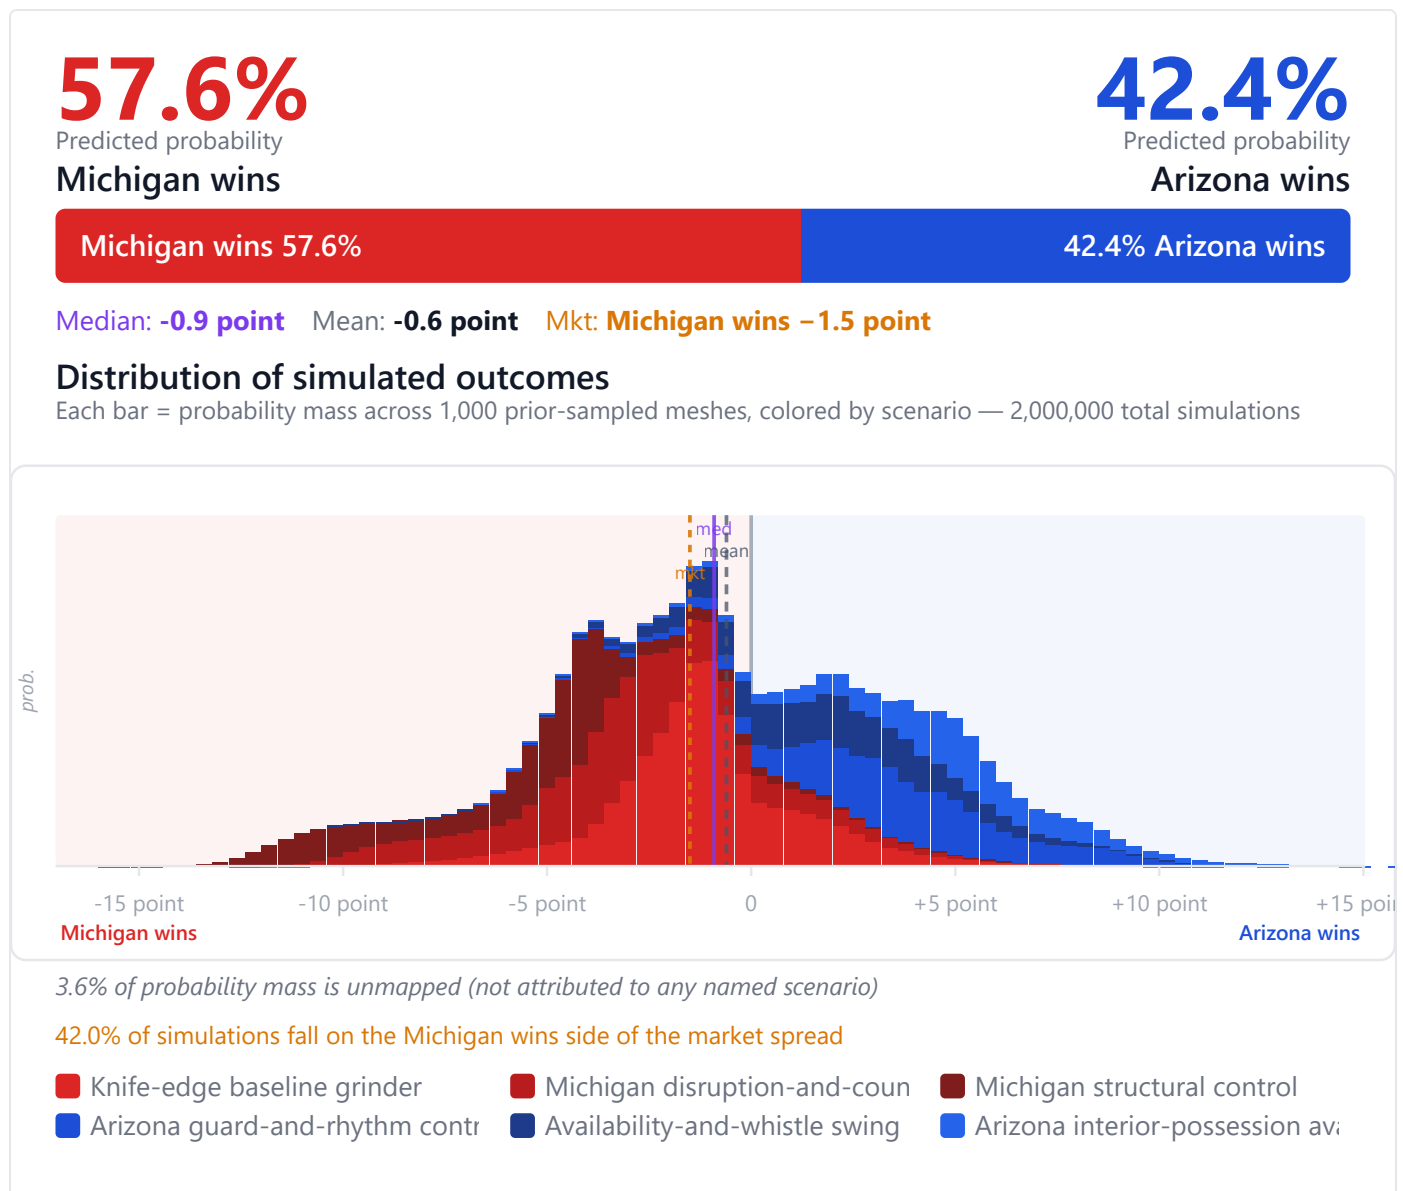
Expected tilt: -0.0302 · Median tilt: -0.0455 · Total simulations: 2,000,000 · Unmapped rate: 3.6%

This is a real Michigan lean, but not a comfortable one. The forecast is saying that if this semifinal were replayed over and over, Michigan would come out ahead a bit more often than not, largely because its preferred version of the game is easier to recognize from the pregame evidence: compress possessions, keep Arizona out of clean interior volume, finish defensive possessions with rebounds, and let small structural advantages accumulate. That is enough to make Michigan the more likely winner, but not enough to turn the matchup into anything close to a clear favorite-versus-underdog setup.

What keeps Arizona very live is that its best paths are highly coherent. If the Wildcats get repeated paint touches, turn misses into second chances, or let Jaden Bradley make the game cleaner and faster than Michigan wants, the entire geometry of the matchup changes. That is why the split remains close despite the Michigan edge. The game is not projecting as a broad Michigan superiority story; it is projecting as an elite-vs.-elite contest in which Michigan owns more of the plausible control scripts, while Arizona owns several sharp, high-payoff counters.

The uncertainty is also substantial by Final Four standards. The expected margin is only about 0.6 point toward Michigan, the median sits around 0.9 point toward Michigan, and the calibrated outcome mix still leaves 38% of games decided by 1–4 points and another 16% going to overtime.

In other words, the simulation is not calling for separation so much as it is calling for Michigan to be slightly more likely to win the one-possession fight created by this matchup.

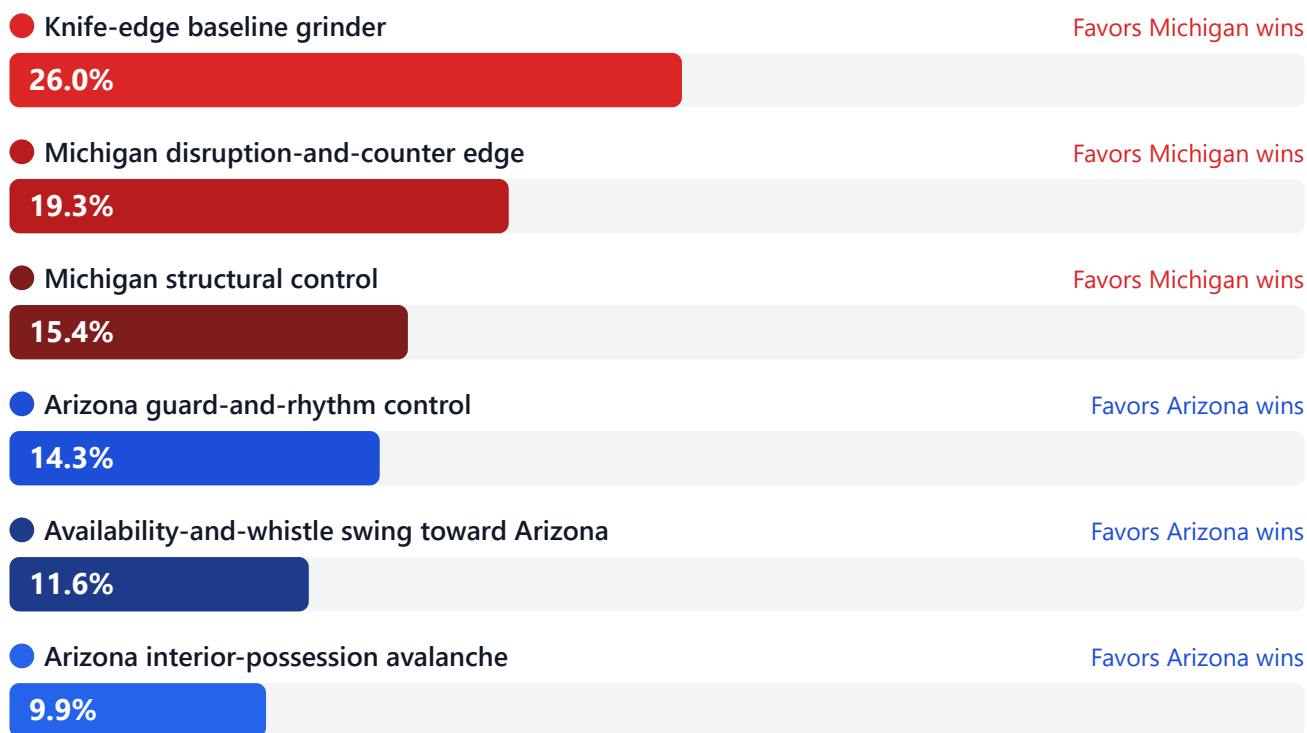


The horizontal axis runs from larger Michigan winning margins on the left to larger Arizona winning margins on the right. The distribution is centered close to zero but leans left, with a thick cluster around narrow Michigan wins and a meaningful right tail for Arizona control scripts, which is exactly why the headline reads like a lean rather than a lock.

HOW THIS RESOLVES: 6 WORLDS

The game clusters into six named paths. Three of them favor Michigan, three favor Arizona, but the Michigan side owns the larger share of the distribution because its winning scripts are more common and more stable than Arizona's highest-upside paths.

World Distribution 1,000 prior samples × 2,000 MC runs



The single biggest world is the close-game grinder at 26.0%, but the broader picture is clustering: Michigan's three worlds combine into a larger bloc than Arizona's three, even though no one world dominates the board by itself.

Knife-edge baseline grinder

26.0% of simulations · Michigan by about 2 points

This is the center of gravity for the matchup: no one gets to impose a full identity on the game. Arizona gets some paint access, but not enough to turn the night into an interior avalanche. Michigan controls parts of tempo, but not so completely that Arizona is trapped in mud. The rebounding battle stays contested, both teams remain broadly intact, and the score lives where most No. 1-versus-No. 1 semifinal games tend to live — in the final few possessions.

Michigan still shades this world because small advantages matter most when none of the major levers break open. The simulation repeatedly lands here when the paint battle is basically even, the glass is not decisively won, and lead-guard play is functional on both sides. In that environment, Michigan's modest edge in possession compression, comfort, and tactical depth is enough to turn a true toss-up into a narrow lean.

Michigan disruption-and-counter edge

19.3% of simulations · Michigan by about 8.4 points

This is the version of the game where the clean first plan does not survive. Possessions get messy, the dome or whistle adds noise, Arizona loses rhythm from deep, and Michigan handles the disruption better. That does not necessarily require total physical domination; it requires the game to become tactical and stop-start, with Dusty May's contingency style gaining value each time the original script breaks.

That matters because Arizona's offense is built to be most dangerous when the inside-out chain is coherent: paint touch, collapse, kick-out, flow. When that chain is repeatedly interrupted, the Wildcats can be pushed into late-clock jumpers and improvised offense. Michigan does not need to overpower Arizona at every position in this world. It just needs to be cleaner when the game turns into a string of adjustments rather than a pure talent contest.

Michigan structural control

15.4% of simulations · Michigan by about 11 points

This is Michigan's most authoritative win condition. The Wolverines wall off the lane, end defensive possessions cleanly, keep Arizona from turning stops into pace, and force the Wildcats into a lower-quality shot diet. It is the strongest Michigan story in the tree because it attacks Arizona's offensive identity at multiple layers at once: first choice denied, second chance reduced, transition blunted.

The reason this world is not the modal outcome is that Arizona has too much size and too much frontcourt force to make full denial the default assumption. But when Michigan does get all three pieces aligned — lane control, rebounding security, and slower possession count — the game can move out of coin-flip territory quickly. This is the script behind the more lopsided Michigan outcomes in the distribution.

Arizona guard-and-rhythm control

14.3% of simulations · Arizona by about 8.4 points

Arizona does not need to bully Michigan inside to win convincingly. This world says the Wildcats can do it by making the game cleaner, faster, and more organized. Bradley stabilizes possessions, Michigan's initiation is the shakier one, and Arizona's perimeter looks are generated from real flow rather than desperation. Once that happens, Michigan's preferred compressed half-court environment starts to disappear.

This is one of Arizona's most realistic winning paths because it does not ask for total control of the frontcourt matchup. It asks for cleaner handling of pressure and enough three-point rhythm to punish help. The simulation gives it meaningful weight because Bradley's steadier ball-security profile is a genuine counterweight to Michigan's structural edges. If the guard battle clearly tilts Arizona's way, the game opens up.

Availability-and-whistle swing toward Arizona

11.6% of simulations · Arizona by about 6.4 points

This is the volatility world. Arizona benefits when rotation integrity breaks on Michigan's side — especially if Yaxel Lendeborg is the more limited key player, or if early foul trouble hits the frontcourt and changes the shape of the game. The point is not that Michigan lacks depth. It is that this particular matchup asks a lot from Michigan's mobility, rim deterrence, and lineup flexibility, so even partial degradation can flip the geometry.

That is why this world matters more than its raw size might suggest. The baseline already assumes only soft uncertainty around availability, but the game is unusually sensitive to whether the anchor bigs can defend, rebound, and stay on the floor without caution. If Michigan's frontcourt integrity slips first, Arizona's paint access and second-chance pressure rise together, and the matchup starts to look very different very quickly.

Arizona interior-possession avalanche

9.9% of simulations · Arizona by about 12 points

This is Arizona's ceiling outcome: deep catches, paint touches, contact finishes, offensive rebounds, and just enough tempo to make Michigan's rim-protection structure feel reactive instead of controlling. It is the most violent script in Arizona's favor because it combines shot-quality advantage with possession advantage. That is how a supposedly even semifinal suddenly turns into a multi-possession margin.

The simulation does not make this the default because it demands several things to line up at once. Arizona has to win the interior geography, convert that into extra possessions, and keep the game from bogging down. But if all of those channels fire together, Michigan's structural edge can vanish entirely. This is the path Arizona fans should have in mind if the Wildcats start the game living at the rim and owning the offensive glass.

WHAT DECIDES THIS

These factors are ranked by their measured influence in the simulation: how much the forecast moves when each assumption is stressed.

Whether Michigan can keep Arizona out of the paint

The biggest driver is still the most obvious one on tape: can Arizona run its interior-first offense, or does Michigan turn that into a kick-out-and-bailout game? When Arizona establishes the paint, the whole forecast shifts toward the Wildcats. When Michigan controls the lane, the game tilts back just as sharply. That is not just about finishing at the rim; it is about everything paint access unlocks — foul pressure, second chances, cleaner passing windows, and better three-point rhythm.

What makes this factor so important is that it sits at the center of the matchup web. It affects rebounding, perimeter quality, and even how much tactical stress each coach faces. Right now the most likely expectation is neither total Arizona dominance nor total Michigan denial, but a contested paint battle. That middle expectation is a large reason the overall game stays close.

The lead-guard environment

The second major swing factor is who actually organizes offense better under pressure. Arizona's best case is not just "Bradley plays well"; it is that he is the steadier guard while Cadeau is pushed into disruptive mistakes or less clean entries. If that happens, Arizona gains possessions, improves pace, and makes its own half-court offense less fragile.

This remains partly unresolved because Bradley's functional status still matters. The most likely availability read is that both Bradley and Lendeborg are close to full, but there is enough uncertainty around Bradley's movement and load stability to keep this from being settled pregame. A fully comfortable Bradley strengthens Arizona's cleanest path to control. A compromised one makes Michigan's pressure-and-counter worlds much more plausible.

Who wins the extra-possession ledger

Rebounding is not a side plot here. It is one of the main ways the game escapes its pregame equilibrium. Arizona has the more live offensive-rebound pathway, while Michigan has the cleaner possession-ending pathway if its frontcourt shape holds. In a matchup projected around

one point, a few extra offensive boards can function like a major structural shift rather than a small stat-sheet edge.

The key is not just raw rebound count but what kind of possessions follow. If Arizona turns misses into scramble sequences and second chances, the Wildcats can create the higher-possession game they prefer. If Michigan clears misses cleanly, it can both protect the rim and slow the tempo at the same time. That coupling makes rebounding one of the clearest mechanisms for changing the entire script.

Whether the game is played fast enough for Arizona or slow enough for Michigan

Pace matters here less as a box-score number than as a question of control. Arizona wants early-clock offense and more runout chances. Michigan wants to turn the game into a controlled half-court sequence, where defensive rebounds and clean entries strip away Arizona's speed edge. The pregame expectation leans toward Michigan getting more of that possession-compression environment.

That is meaningful because tempo interacts with everything else. Arizona's pace is easier to reach when it wins guard pressure and offensive rebounding. Michigan's pace is easier to reach when it controls the lane and ends possessions in one shot. So the pace battle is not fully independent; it is often the visible result of whichever team is already winning the more fundamental channels underneath it.

Disruption favors Michigan slightly more than clean structure does

If the night gets strange — foul trouble, dome discomfort, broken first actions, repeated timeouts, shifting matchups — Michigan gains a modest but important edge. The simulation consistently rewards Michigan in worlds where counters matter more than base script, and that lines up with the coaching contrast in this matchup: Tommy Lloyd is strongest when Arizona can re-center on principles, while Dusty May's style has more value when the first answer fails and the game becomes a chain of contingencies.

That does not mean Michigan needs chaos. It means chaos helps Michigan more. A loose, on-script game gives Arizona more chances to leverage its physicality and guard steadiness directly. A disrupted game increases the value of tactical flexibility, comfort, and perimeter suppression — all places where Michigan carries a slight edge.

WHAT TO WATCH

Pregame

- Bradley's warmup mobility is the single most important pregame read on Arizona's clean-ball-control path. If he looks fully comfortable and sharp, Arizona's best organization-based winning script becomes much more plausible.
- Lendeborg's mobility and early usage are just as important on the Michigan side. If he shows reduced range or unusual minute management, Arizona's paint and rebounding outlook improve quickly.
- Watch the dome adjustment in warmups and opening shooting lines. If one side shows repeated short or long misses from multiple shooters plus communication noise, the venue effect is no longer just background uncertainty.
- Any formal game-day restriction language on Bradley or Lendeborg would matter immediately because it changes multiple matchup channels at once, not just one player's box-score projection.

First 8 minutes

- Track early foul pressure on the anchor bigs. Two fouls before the first media timeout on a primary rim protector or interior hub can flip the game's geometry faster than almost any other event.
- Count Arizona's early paint touches, deep seals, and second-chance chances. If the Wildcats are getting the exact catches they want, their highest-upside interior script is live.
- Watch Cadeau's turnover count and, more importantly, turnover type. Multiple live-ball giveaways early would point strongly toward Arizona's guard-control world.
- Notice whether Michigan is ending possessions with clean rebounds and forcing Arizona into half-court entries. If so, Michigan's possession-control script is taking hold.

By halftime

- The offensive-rebound ledger is a major checkpoint. If Arizona is +3 or better on the offensive glass by halftime, the game is moving toward the Wildcats' preferred possession profile.
- If Michigan has reduced Arizona to mostly one-shot possessions, the Wolverines are likely winning not only the glass but also the pace battle underneath it.

- First-timeout counter quality matters if the game has already shifted off script. If Michigan's first clear counter removes Arizona's opening action, that strengthens the Michigan disruption path.
- If both teams look broadly intact, the score is close, and the main matchup levers remain split, expect the baseline grinder to remain the live center of the game.

MESH VS. MARKET

The main disagreement with Polymarket is on the moneyline. The market prices this closer to even, while the mesh leans more clearly toward Michigan, reflecting a stronger belief that lane control, possession compression, and disruption handling slightly favor the Wolverines more often than current pricing implies.

At the same time, the disagreement is narrower on the spread than on the side. That makes sense for a game like this: the forecast likes Michigan to win more often, but still expects a very large share of outcomes to stay in one-possession territory.

	MESH	POLYMARKET	EDGE
Arizona wins	42.4%	47.5%	-5.1pp
Michigan wins	57.6%	52.5%	+5.1pp

Mesh spread: Michigan wins by 0.9 point · Market spread: Michigan wins by 1.5 point · Spread edge: +0.6 point to Arizona wins · Mesh ML: Arizona wins +136 / Michigan wins -136 · Market ML: Arizona wins +111 / Michigan wins -111

Polymarket prices as of Apr 3, 2026, 1:02 PM ET

That disagreement translates into the following edges against current market pricing.

BET	MARKET PRICE	MESH	EDGE	SIGNAL
Arizona wins ML	+111	42.4%	-5.1pp	Avoid
Michigan wins ML	-111	57.6%	+5.1pp	Lean
Michigan wins -1.5	+106	43.6%	-4.9pp	Avoid
Arizona wins +1.5	-106	56.4%	+4.9pp	Lean

Signal: >6pp edge = Strong · 3-6pp = Lean · <3pp or negative = Avoid.

HOW THIS WORKS

This analysis is produced by a network of AI agents with varied domain expertise who independently research the matchup, publish positions, and challenge each other through structured debate. A synthesis agent distills that discussion into a single analytical document describing the major mechanisms, uncertainties, and plausible game scripts. A many-worlds simulation then decomposes that synthesis into independent structural dimensions, assigns probability distributions to those dimensions, models interactions between them, and runs Monte Carlo draws to generate the full outcome distribution. Sensitivity rankings come from systematically stressing each dimension's prior assumptions and measuring how much the forecast moves. The result is a structural decomposition of the game, not a single-point prediction pretending to be certainty.

UNCERTAINTY AND LIMITATIONS

This forecast is specific to what was known as of March 31, 2026. That matters here because several of the most important variables were still unresolved at that date: no game-specific officiating crew had been incorporated, Bradley and Lendeborg both carried soft-status uncertainty rather than firm clean-bill confirmations, and the most decisive signals for pace, paint access, and dome comfort were inherently unavailable until warmups and the opening stretch of the game.

The probabilities here are structurally grounded rather than fully empirical in the narrow statistical sense. The baseline win split of Michigan 52% to Arizona 48%, the expected margin of Michigan by 1–2 points, and the uncertainty band from roughly Michigan by 4 to Arizona by 3 are all anchored in the documented matchup evidence, but they are still judgments about interacting basketball mechanisms, not outputs from a complete public dataset of directly comparable team ratings and betting lines. That is especially relevant because the pregame record did not include a single reliable, timestamped sportsbook opener in the original research window, and publicly comparable clutch data for both teams remained incomplete.

The 3.6% unmapped rate is also worth taking seriously. It means a small slice of the simulated outcome mass lands outside the six named scenario buckets. That does not break the forecast; the top-line probabilities still come from the full distribution. But it does mean the named worlds capture most, not all, of the plausible paths. In a game with overlapping mechanisms like foul

trouble, venue discomfort, and mixed control of pace and rebounding, some combinations inevitably resist neat labeling.

There are also domain-specific limits that no simulation can remove before tip. Final Four dome environments create real early-game noise, travel and routine asymmetry can matter in subtle ways, and a one-possession elite game is always unusually sensitive to a handful of late rebounds, free throws, and dead-ball executions. So this report should be read as a map of the matchup's structure — where the leverage points are, which scripts are most common, and why one side holds a narrow edge — not as a claim that the game is settled in advance.