

How To Play Boros: A Case Study of Three Red-White Limited Deck Archetypes

Zev Goldhaber-Gordon

November 2025

1 Introduction

Have you ever forlornly set down your pile of Sultai cards, and thought to yourself: I wish I could draft Boros? Well, this guide has some tips and tricks for using cards with red, white, and red-white color identities from packs of MTG cards! We will center our discussion around three deck archetypes which I tend to see pop up again and again. While these deck archetypes are most prevalently Boros, they also inform aggressive decks in every color combination, as well as how to play against aggressive decks.

The through-line of these three decks and of good aggro play in general, is that **with clever deck construction and gameplay, you can force your opponent to misuse their resources in various ways**. While playing aggro has a lower skill-floor than most other archetypes, as it is easy enough to turn creatures sideways until your opponent is dead, what really raises aggro's skill ceiling is that once your opponent's resources have been taxed (i.e. you voted for a small game, as discussed in Sam Black's article on the concept¹), gaining your own marginal resources becomes far more important. When games hinge on single lifepoints, squeezing every bit of advantage out of your cards becomes more difficult to do when autopiloting. Playing on the margins in Boros involves knowing exactly how much of your opponent's life each of your creatures is worth, and making as many bluff attacks as possible without sacrificing too many resources.

This article assumes some familiarity with limited theory and a wide variety of recent limited sets. We'll first discuss how to accrue these advantages with disciplined and calculated gameplay, then discuss deck construction composition and card roles, before continuing on to the three deck archetypes.

¹Elvish Visionary VS Elderfang Disciple: The Nature Of Card Advantage (Sam Black)

2 General Boros Skills

2.1 Bluffing

Consider the following scenario. You and your opponent both kept 7 card hands, with you on the play. Turn 1, you play a Kavaron Harrier, then after your opponent's island-pass, turn 2 you play a Knight Luminary, making a 1/1 token. Your opponent plays a Zookeeper Mechan and passes back. It is now your turn 3, and you have no combat tricks or spells to cast in hand. What do you do? (See Fig. 1).



Figure 1: Turn 3 On the Play

You have a few choices here. The easiest decision to make, especially if you are playing on autopilot, is to just pass. You don't want to throw away your Kavaron Harrier for nothing, and you have no way to protect it in hand, so why attack? One level higher is recognizing that your opponent likely does not want their mana dork to die (bolt the bird's corollary: attack into the bird), so they'll probably let both your creatures through if you attack. But what about a stronger opponent? Now you need to think about what spells you're telegraphing. You have eight options for your attacks:

1. Nothing
2. Just the 1/1 token
3. The Kavaron Harrier, without making a 2/2 robot
4. The Kavaron Harrier, making a 2/2 robot and leaving up a white mana
5. The Kavaron Harrier, making a 2/2 robot and leaving up a red mana

6. The previous three options, but also attacking with the 1/1 token

Of these, you can immediately eliminate attacking with just the 1/1 token, as this does not telegraph anything that attacking with the Kavaron Harrier wouldn't already. You can also likely eliminate attacking with nothing, the naive choice. While you might have a card in hand that incentivizes playing slowly, allowing your opponent's 2-drop mana dork to stonewall two cards worth of your early game value while simultaneously ramping is probably not a winning line. Unfortunately, a strong opponent will be aware of this, and know that an attack doesn't necessarily represent a spell.

So, what spells can you represent? If you leave a white mana up, you are representing Focus Fire and Squire's Lightblade, as well as Reroute Systems which isn't really a consideration since it doesn't do anything to progress our gamestate. If you leave a red mana up, you represent a Full Bore, and if you pay for Kavaron Harrier on attack, a post-combat Plasma Bolt to kill the Zookeeper Mechan after the 2/2 robot dies. If you leave all your mana up, you add some further range, including Zealous Display, Rig for War, or a post-combat Cut Propulsion to kill the Zookeeper Mechan if it blocks the Kavaron Harrier. Crucially, these >1 mana cards are much worse cards than the 1 mana options you could have access to, meaning your deck is less likely to contain them, and your opponent isn't as likely to even think of them. (Cut Propulsion is a good card in a vacuum, but is a disastrous enough line for you that your opponent would be happy if you went for it.) If they see you not pay for Kavaron Harrier's attack trigger, they're likely to think you have a post-combat play you want to develop, rather than a Rig for War. Since the name of the game is to convince your opponent you have a combat trick while pushing as much damage as possible, you can exclude any option which doesn't involve making a robot token.

You can also exclude leaving up a red mana, not because it telegraphs less threatening lines (although it does), but because of how they are incentivized to respond to those lines. The common Plasma Bolt is more likely to be in your hand than the uncommon Full Bore, and the way to play around Plasma Bolt is to block the Kavaron Harrier, securing a 2-for-1 instead of the 1-for-1 and 4-5 damage they would take if you can post-combat Plasma Bolt with void.

Since you have nothing to play, that leaves you with a single white mana. Should you attack with the 1/1 token? If you do not, you are telegraphing a very specific threat, a copy of Squire's Lightblade. Even though Focus Fire is the more powerful and thus more common card, you would attack with your 1/1 token if you had it, so your opponent will immediately notice the strange attack and think of Squire's Lightblade. How they respond depends on their hand, and their read: if they don't particularly care about their Zookeeper Mechan surviving because they have a 3 drop, they could choose to block your 2/2 robot, forcing you to either equip your token and let it fall off if you want to win the combat, or waste the 2 mana you spent on the robot. If they doubt you have Squire's Lightblade, they could even choose to block the Kavaron Harrier and call your bluff, but if you did have the equipment, it would then stay equipped,

opening up better future attacks.

On the other hand, if you attack with the 1/1 token, while it might seem like you are expanding your range to include Focus Fire, you are also somewhat contracting Squire's Lightblade out of your range, since the naive attack with it in hand is just the 2 power creatures. This shifts the calculus a bit, making the risk-reward more favorable to block, as seen in Fig. 2. Your opponent may also not notice the safer available attack for Squire's Lightblade or think that you are bluffing not having it and split the difference, blocking your 1/1, which is also an unfavorable result. If you have a read on your opponent especially wanting to keep their Zookeeper Mechan alive, such as having cast multiple Kav Landseekers in a previous game, or if you have previously shown Focus Fire, you could still go for the full attack, but I think I would tentatively land on the specific threat of just attacking with your Kavaron Harrier and paying for it, leaving a white mana up. A different Boros player might be more fearless and attack with all three, deciding that the gamestate is dire enough that you need to take a greater risk for a single extra point of damage. There isn't a single correct answer, since if you always play the same way, your bluffs won't be effective.

With 1/1 attacking:	You have a trick	You don't have a trick
They block Harrier	They lose Mechan, take 3	You lose Harrier
They don't block Harrier	They take 5	They take 5

Without 1/1 attacking:	You have a trick	You don't have a trick
They block Harrier	They lose Mechan, take 2, equip Lightblade	You lose Harrier
They don't block Harrier	They take 4	They take 4

Figure 2: If you're telegraphing Squire's Lightblade instead of Focus Fire, blocking becomes marginally more painful compared to letting everything through. Here assuming a free equip is more painful than taking an extra damage, since it represents 1 damage next turn, when you have another revealed play of Knight Luminary.

Now, this decision making was all for only two creatures and one blocker, on turn 3, without even considering too deeply what your opponent could use to read whether you are bluffing. Imagine having to figure this out for a more complicated board state later in the game. When an attacker says "Math is for Blockers," it is not because there is no math available for the attacker — it is because that math is completely untenable in any reasonable time!² For every

²For anyone except Zevin Faust and some hall of famers.

arrangement of blocks a blocker has to consider, an attacker has to consider those blocks for each possible configuration of attacks, including the spells they are telegraphing, and the information their opponent has revealed about their hand. While blocker math scales quickly as $\mathcal{O}(mn)$, where m and n are the number of attackers and blockers respectively, attacker math scales as the insane $\mathcal{O}(2^m mn) \approx \mathcal{O}(2^m n)$. Although the larger the boardstate is, the more attacker options you can immediately exclude, the possibility space is just too large to think through everything.

Furthermore, as you make your decisions, a competent opponent can use the time you spend thinking to judge what you have in hand. If you rapidly untap, attack with Kavaron Harrier and summon a 2/2 robot leaving up a single white mana, that much more effectively sells a Squire's Lightblade than if you took your time thinking through all these considerations. You can think of your plays by imagining a rectangle, in which one side represents the quality of your decision, and the other side represents how quickly you made that decision (See Fig. 3). A good play should maximize the area of that rectangle, and the shape with the largest volume compared to the sum of its sides is a square: that is, you should think until your thinking starts to get diminishing returns, and no more. Playing too quickly will give up costly mistakes, and playing too slowly also takes time off the clock for both players.




Figure 3: Tradeoff between speed of decision making and efficacy of decision.

As an aggro player, you (and your opponent) are entitled to play a bit more slowly than in most matches, since your games will naturally be over

more quickly and involve less decisions that each become commensurately more important. Still, going too far in the tank is actively stealing time from your opponent, and if your slow play consistently leads to draws, it's probably worth speeding up.

This is a lot of pressure when you're deciding which creatures to turn sideways! Fortunately, there's a fix. Your opponent's turn is free time. If you're constantly thinking on their turn, not only will you more quickly make decisions on your turn without giving up timing tells, you also deprive them of time to think on your turn, making any removal or combat tricks they choose to hold up all the more obvious. It's particularly useful to count out how much mana you'll have available, and decide on your highest value usage of that mana, assuming no relevant topdeck — calculating combat math ahead of time is a lot harder, since your opponent can play extra blockers which shake things up. As with anything, moderation is important, since you don't want to think too hard about your attacks and miss a detail about your opponent's land sequencing or a facial tell. Regardless, the more you can frontload your combat math, the smoother your turns are going to look.

2.2 Mulliganing

Another crucial aspect to maximizing your resources in Boros aggro play is to have tight mulligans. Mulligans are more important in Boros than in any other deck because as a game speeds up and your deck has access to less card draw, your starting hand becomes a greater percentage of the total cards you'll see during the game. Mulliganing in Boros means balancing the need for an aggressive start and at least three lands on curve with the difficulty of recovering from down resources without any catch-up mechanisms. We'll attempt to quantitatively progress our analysis beyond vibes-based mulliganing.

The first set of numbers is inspired by an old blog post from 17lands³ on the impact of mulliganing on winrate, as a calibration for exactly how punishing a mulligan is. Here is a modernization of the post's primary chart, updated to all recorded play booster standard set games.

Play/Draw	# Mulligans	Win Rate	% Games	# Games
Either	Any	50.00%	100.0%	1071086
Play	0	56.38%	39.42%	422211
Play	1	42.15%	9.22%	98796
Play	2	25.27%	1.26%	13529
Draw	0	50.05%	39.48%	422887
Draw	1	37.83%	9.24%	98964
Draw	2	24.66%	1.19%	12728

This data includes both the player and the opponent from each log, so it excludes any bias from 17lands users being stronger than your average player.

³Impact of the London Mulligan (Robert Conroy)

It also only uses game 1 data, as in games 2 and 3, the player who won the previous game is on the draw, and thus the player on the draw is more likely to win in a vacuum. In fact, if you include games 2 and 3, being on the draw has a 50.18% winrate!

On the play, a mulligan from 7 to 6 means a 14.23% hit⁴ from 56.38% to 42.15%. In addition, with 9.22% of games tracked on 6 and 1.26% on 5, there's a 12.02% chance of having to mull from 6 to 5, for a further hit of 16.88% down to 25.27%. Putting it all together, mulliganing your 7 represents a total of a

$$0.5638 - [(0.1202 \times 0.2527) + ((1 - 0.1202) \times 0.4215)] = 0.1626 \quad (1)$$

or 16.26% hit, taking us to an approximately **40.12% expected winrate**. Similarly, on the draw, mulliganing represents a 13.72% hit, taking us to an **expected 36.33%**. This already tells us something: that while mulliganing a marginal hand to six and being forced to go all the way down to five feels disproportionately bad (“why didn’t I just keep the seven?”), the possibility shouldn’t statistically influence your mulliganing decisions too much, as it’s only about 2% off your expected winrate in total.

Now, let’s combine these numbers with an important tool for a competitive player who wants to find the exact correct decision every time: the Hypergeometric Calculator.⁵ This widget allows you to find the probability of drawing specific cards in specific sample sizes under specific distributions, like “I need to hit at least 1 of my 15 remaining lands from my 33 cards in deck in the first 2 draws,” which has a 71% chance of occurring. When deciding whether to keep or mulligan hands, I’m always thinking back to this result and similar ones, as well as doing rudimentary calculations in my head, and comparing them to the 17lands mulligan results.

One implication of this type of combinatorial analysis is that the fewer lands in your hand, the better it is to be on the draw, as each extra draw towards a land becomes more important. An oft-unrecognized corollary follows from this reasoning:

1. Being on the play leads to somewhere between a 50% and 54% winrate, depending on the format, averaging to 52.88%.
2. Being on the draw has a higher winrate compared to play winrate when low on lands than the average draw winrate compared to the average play winrate.
3. Therefore, being on the draw must have a somewhat lower winrate when high on lands⁶, so if your hand has a lot of lands, it’s disproportionately better to be on the play.

⁴Here using absolute percentage decreases instead of relative, since we care about total winrate.

⁵Aetherhub Hypergeometric Calculator, or the more advanced Card Combination Calculator for particularly intractable problems.

⁶You may say that play-draw bias is not necessarily linear with land count, and it could be bimodal, with draw winrate peaks at both high and low land counts, and a trough at mid land counts. To that I say: shut up nerd.

4. Finally, just like you mulligan low-land hands on the play more, you should commensurately mulligan high-land hands on the draw more.

The simplest heuristic stemming from this assertion is that 5-land hands are rarely keepable on the draw, but can sometimes be keepable on the play. When looking at a 5 land hand, this is not intuitively obvious, since it feels like on the draw you have more time to draw action. Unfortunately, every land you draw with a large number of lands in hand is less likely to be played than the last, whereas when drawing nonlands in a hand with few lands, you'll still potentially play every nonland in your hand. This makes each draw lower value, favoring the potential tempo from being on the play.

Let's apply this math to some practical examples. Consider the following hand on the play (see Fig. 4): Mountain, Rogue's Passage, Axcgard Cavalry, Resolute Reinforcements, Inspiring Paladin, Felidar Savior, and Angel of Finality.



Figure 4: Passage gone rogue. To keep or mull is plains as day.

This is not a great hand, and you don't need numbers to tell you that. While you do have a valuable castable 2-drop, you don't have your third land drop, and there are four cards in hand you can't cast. But if you do topdeck a Plains, this hand is perfectly fine, so how likely is that? With eight Plains in the deck (for a total of 17 with eight Mountains and the unfortunate Rogue's Passage), our handy dandy Hypergeometric Calculator tells us that we have a 43.2% chance of drawing one of eight cards in two draws from our 33 card deck. If we were on the draw, that would go up to 57.8% in three draws. If we take this as an average seven card hand if we do hit a Plains and claim that we lose on the spot if we don't, this hand would have a $0.5638 \times 0.432 = 24.4\%$ winrate, lower than the 40.12% expected off of one mulligan.

But that’s just a floor: this hand will sometimes win games in which it misses that third Plains for a turn, which we can evaluate by using the Card Combination Calculator. There are six castable cards that we could topdeck, at 21.6% of draws (not including draws with both Plains and the other castable), and three possible red 3-drops alongside a Mountain, for a 4.0% chance of getting both. You can also win games in which neither of those happen, just by the opponent also stumbling early: with a 25.8% chance of topdecking a Plains on the third turn after missing on the first two, it’s at least worth considering. These options aren’t created the same: topdecking a 3-drop and a land is a lot worse than a Plains, and relying on your opponent to stumble is worse than either. But they do add to our win probability.

The hand also isn’t necessarily average if you do hit a plains. I would guess it is a good 10% better than average in that case, or $56.38\% \times 1.1 = 62.02\%$ winrate. Getting a land and a 3-drop seems 60% as good as your average seven, 33.83%, and getting a 2-drop or missing a land drop until turn four is perhaps 30% as good, 16.91%. These numbers are completely made up, but should serve as reasonably starting points. If we add everything together (neglecting the possibility of drawing both a 2-drop and a Plains) we get:

$$0.432 \times 0.6202 + 0.216 \times 0.1691 + 0.04 \times 0.3383 + (1 - 0.432) \times 0.258 \times 0.1691 = 0.3428 \quad (2)$$

a 34.28% winrate, still lower than the 40.12% mulligan winrate but close enough that you could see keeping being worth it if you’re a bit higher on the hand in the draw-a-Plains case! Herein lies the issue: as soon as we introduced variables corresponding to the quality of each possible draw, we reintroduced subjectivity to our calculation, but sanitized it through a numerical lens as being “more objective.”⁷ For this reason, I prefer to stick with the absolute floor we calculated earlier, 24.4%, and admit to myself that how much I value this hand over its floor winrate is vibes-based. Still, going through this whole exercise can tell us three things:

- By calculating the odds of drawing the cards we need and comparing those possibilities to an average seven, we can compare apples to apples (sevens to an average seven), instead of trying to compare a seven to an average six, which is trickier. These calculations could inform how we think about similar hands in future when we’re deciding how much to raise a hand above its floor: in this case, the full calculation had an approximately 40% higher relative value than the floor, which gives us a good order of magnitude estimate for how much to care about the cases in which we whiff on lands, vibes-wise. Pay attention to any cards you’re unlikely to play, though, as a 6-drop in a two lander or a fifth land may both be practically

⁷This is also a problem with projected metagame analyses in constructed: instead of directly voicing your preconceived notions about relative deck strength and meta spread, you construct the metagame that would manifest if everyone thought like you. This metagame will then, of course, back up your preconceived notions.

useless, making your hand an effective six already. That would make it worth comparing oranges to oranges (your effective six to an average six) instead.

- Mulliganing is painful. While a 16% hit in winrate may not seem too bad, when you add up all the possibilities for good cards you could draw, it takes a pretty bad hand (like this one) to be worth mulliganing. Limited isn't like constructed in which the nature of your cards matters much more than their quantity; a limited player can't afford to be picky, because every card in hand is important. Similarly, you can win a lot more games in limited after stumbling, because no deck has enough power to cleanly end the game, so you see a lot more scrappy games coming down to weird board states. This also makes every card more important, although Boros is the deck that can least afford to stumble as stated above, so this point is less relevant to the article.
- Finally, we've seen here how punishing it is to be kept off a color, far more punishing than missing land drops. If this Rogue's Passage was a Plains, we'd have a full 71.0% chance of drawing a third land on time (and another 2-drop if we do miss that third land drop for a turn), making this a snap keep. It's just much easier to have to draw any land than it is to draw a specific land. As is, one-color hands can get choked on their second color of mana even once you topdeck a land of that color, and a lot of the time will simply lose without hitting their second color at all. A hand with neither a turn two play nor a second land color is usually a mull.

That's exactly what happened to this hand, which has claimed a dubious honor: **it is the only hand to ever lose against an opponent who mulliganed to two** in this entire dataset. You can see from the game log⁸ that while our Boros hero makes several deeply questionable attacks, they were not favored in this game even with good play given their draws. They never hit their white source, stranding six cards in their hand, and effectively mulliganing one more than their opponent.

Another, less discussed mold of hand is the type that is just somewhat meh. It's a lot easier to think about how many draws you have to draw a land than it is to think about the character of threats and answers you need to draw to make a 3-4 land hand functional. Consider this hand on the play (See Fig 5):

Plains x3, Mountain, Cunning Maneuver, Airbending Lesson, and Kyoshi Warriors.

This hand is a snap keep in any slower deck, if you replaced the Cunning Maneuver with a more defensive spell, since those decks don't require an early play, and are fine using Airbending Lesson defensively. In Boros, however, waiting for turn four to make your first proactive play can be a death sentence.

⁸<https://www.17lands.com/history/ead5dab1a6164525930c9e288e1f8a74/2/1/3>



Figure 5: Makes me want to fall asleep.

My instinct, before doing any math, is that this hand doesn't quite get there. This deck has eight 2-drops and four 3-drops, so we can calculate the chances of drawing those on curve like before: 24.2% to draw a 2-drop in our first draw, and 60.2% to get either a 2-drop or a 3-drop by turn three. It also has a single 1-drop, with a further 3.0% to draw it turn two, but by turn three holding up Airbending Lesson is likely more impactful than casting the 1-drop.

So, what's this hand's absolute floor percentage? For discussing when to draw lands, it's easy to define a floor as "having missed your land draw by turn N," but when you can draw various nonland cards to save the hand, that's harder to do. In the event that we miss for three turns, Airbending Lesson on three, and Kyoshi Warriors on four, I think we can call this sequence a loss when calculating a floor. If you topdeck a 2-drop on two, we can consider that a functional hand. But what if you get a 3-drop on turn three? What about a 2-drop on turn three? An absolute floor is a lot less useful with such a gradient of possible hand textures: surely this hand doesn't have a $0.242 \times 0.5638 = 13.6\%$ winrate. Qualitative statements like "a Boros deck must play a creature by turn 3 in order to have the tempo to win" can also be dubious choices to fall back on and convert into numbers. This whole article is about how to grind out value advantages as an aggro decks, and as emphasized repeatedly, limited games often go off the rails from theory (excepting Section 3.1).

I think the best way to bridge the gap between the absolute floor of this hand and its actual winrate is to consider what you could topdeck turns three to five if you don't hit a creature turn two. In this particular deck, with eight 2-drops, could things get better over time if you do try to use Kyoshi Warriors as your primary threat with combat tricks and removal? Can we productively use the

extra card advantage from Cunning Maneuver, or will it and Airbending Lesson just draw you the 2-drops you wished you had on turn two? With 21 subpar draws after turn three (eight 2-drops and thirteen lands), we'll hit one two-thirds of the time. In the hand as initially presented to me, Cunning Maneuver was actually Lost Days. While the deck was still aggro and expected to be the beatdown, it was Azorius aggro with several pieces of card advantage. These questions differentiate the Boros version of the hand from the Azorius one. Since there was a real possibility for that deck to claw back the resource game even if it did miss on its eight 2-drops, I felt it was worth keeping, while with that possibility diminished, there are too many lose-cases to keep.

This might seem less precise than what we did with the two land hand, and that's true. In general, it's much easier to quantitatively analyze low possibility-space hands, where the spells you cast are likely to be mostly the ones you're already holding. Just as deciding attacks optimally in a sufficiently large board state is structurally impossible, you can't consider every potential draw and numerically break them down within the time constraints of a Magic game, so being comfortable with the vibes side of mulliganing is just as important as knowing your land odds. That starts with some common probabilistic calculations, but it mostly relies on the questions you ask yourself in the gaps between certainties. Everyone has a slightly different mulligan instinct, so the more you interrogate your own assumptions, the better your draws will be on average.⁹

2.3 Card Choice and Deck Construction

This section will largely be a compilation of some pretty widely known heuristics for Boros deck construction, along with a few trends I've noticed. So, what cards do your Boros decks want?

To start with, 2 drops. Any Boros shell wants the ability to proactively influence the board, and the lowest mana cost in limited is 1, but most aggro decks don't have too many 1-drops. While 1-drops are nice to have, they aren't printed in sufficient quantities to build a functional limited curve with, and they tend to be significantly poorer topdecks than 2-drops. One point in 1-drops' favor is also a reason to not run too many of them. Part of their power is that you can generally consider them to cost 0-mana in sufficiently small quantities, as you'd otherwise waste your mana on the first turn of the game. If you fill your deck with enough 1-drops and tap-lands, though, you'll need to cast some later in the game, which often ends up being turn 4+ if you have other creatures to fill in your curve before then. At that point, your 1-drops don't pull their weight anymore.

So in most limited formats, your curve starts on 2. While removal famously gets picked much more highly in strong pods, I personally prioritize 2-drops even higher in aggro. In a slower deck, they represent a road block against aggro decks even when understatted, but in an aggro deck, 2-drops are your lifeblood. If you

⁹Originally I phrased this as "the less nongames you will have," but it can be a leak to prioritize safe hands over hands which could potentially fall apart but have a better EV.

can't hit six 2-drops, you'll have serious consistency issues and need to mulligan more often, since your strategy depends on putting your opponent on the back foot immediately. As discussed in Section 2.2, I am always looking for castable 2-drops in my Boros hands, and if there isn't one present, the rest of the cards need to pick up the slack. I would consider almost any aggro deck with four or less 2-drops to be disastrous regardless of what the rest of the curve looks like. Most of my aggro decks will have six to seven 2-drops, and I am always thinking of what I need to fill the slot out. You can play anywhere from two to six 3-drops, and up to four 4-drops, although each 4-drop becomes harder to support the more 2-drops you have, as they can be replaced by a pair of 2-drops or a 2-drop and a combat trick on turn four. Above four, your mileage may vary: some aggro decks can play the odd 6-drop, but you basically cannot support a 7-drop in a typical aggro shell.

For your removal suite, large removal is better than small removal, because small removal is more replaceable. When your opponent plays a 2-drop into your 2-drop, you could either trade your creatures off and develop a 3-drop, or use a cheap combat trick to run over their creature if they block and play another 2-drop. Worst case scenario, their 2-drop neutralizes your 2-drop until a turn or a few turns later, but your 3-drop still probably has attacks. When your opponent plays a larger creature on turn four or five, though, it can hold back your entire board until you find a way to deal with it, or accept sacrificing a creature with each attack, so it's important to have hard removal ready when that large creature arrives. Overall I'd like 3-5 large removal spells, 0-2 small removal spells, and 0-3 combat tricks in my aggro decks.

Filling out your deck are lands. A pretty sophisticated piece of 17lands data analysis convinced me that more decks should be running 16 lands, aggro decks especially.¹⁰ While it makes sense to run less lands in decks which top out at four mana spells, the difficulty lies in the bare minimum number of sources you need to reliably cast each of your colors. Ideally you would want nine sources of each of your primary colors, meaning a 17 land manabase with eight of each basic type, and a dual land of both colors. The eight-eight manabase on 16 lands is doable, but it certainly doesn't feel good for either color, meaning you'll mulligan more and lose win percentage. Mulligans are punishing in limited in general but particularly painful in a value-light aggro deck. You can play multiple dual lands, but that takes away from your playable picks, and you need an extra playable pick with 24 nonlands. Taplands also interrupt your curve at some point, which feels bad, but is sometimes necessary. One strategy I've used in some sets recently is to lean heavily into one color, running the other as something closer to a splash than a primary color, and then run a 9-7 or 10-6 manabase. 6 to 7 lands is perfectly fine for casting your splash, and you get plenty of lands for your main color.

For your final cards, you have a bit of leeway. There are some cards that play quite well if you only have a few of them, like card draw spells: if you draw

¹⁰16 is the new 17: Mathematical analysis of 17lands data (TimLewisMTG). The analysis uses Best of 3 data, so does not incorporate the Best of 1 hand smoother, which would presumably make running 16 lands even better.

multiple, they can clog up your hand and delay your aggro starts, but with only one in the deck, you can treat it as a top end threat. Bombs are separate from the curve, since the way you sequence them are much more gamestate contingent than your typical curve fillers. Synergy pieces, game-enders like Act of Treason and Cosmotronic Wave, and extra 2-drops also find their way into your deck sometimes.

In the last few sets, there has been one more card that has snuck its way into my typical Boros deck: the evasive equipment creature. Living Weapon and its derivatives (Job Select, For Mirrodin!, and all the unkeyworded equipments that enter with a creature token) have been revolutionary for Boros in the sets they appear, granting early aggressive threats along with a value outlet later. Creature equipments that grant flying in particular have consistently overperformed. Auxiliary Boosters, Dragoon's Lance, Mandibular Kite, Hexgold Hoverwings, and Dragonwing Glider have all been powerful, even above most of them being quite initially pushed. Historically, Boros decks could not run flying equipments for creature-count reasons, as aggro decks need 15-18 creatures to guarantee their draws get off the ground, but in creature-dense hands, flying equipments can be quite powerful. I distinctly remember Raven Wings in particular overperforming in Kaldheim limited, despite being actively weak on rate. Flying equipments allow you to present threats repeatedly, turn after turn, and your opponent must answer each of those threats at a typically more expensive rate in cards and mana than it took to set them up. Whether opponents trade real fliers with smaller creatures or spend a card to remove a token, flying equipments steadily generate advantage and make any stumble more punishing, as a three or four power flier can kill remarkably quickly after an aggressive start. When Wizards started printing creature equipments, their weakness of not being either removal or creatures was filled in. I usually would like to have a flying equipment creature in my deck, set permitting; up to three at balanced rates or infinite Dragoon's Lance-tier copies.

Putting it all together, here's the skeleton of your platonic ideal Boros deck:

- 2 1-drops
- 7 2-drops
- 5 3-drops
- 2 4-drops
- 1 5-drop
- 4 large removal spells
- 1 small removal spell
- 2 combat tricks
- 16 lands

- A flying equipment replaces one of those N-drops

As a note, the platonic ideal Boros deck is not the most common Boros deck, as most will have one or less 1-drops, a few more 4-drops, less large removal spells and more small ones, etc. Don't sabotage your card quality to fit this mold, and deviate from this skeleton as necessary, including going up to 17 lands often. The hip bone's connected to the...bone saw.

I won't derive all of these numbers from first principles, but I will do so for the most fundamental, the removal count and the 2-drop count. To start with, from that historical play booster 17lands dataset earlier, Boros games have taken 8.80 turns in wins and 9.29 turns in losses on average, so about 9 turns. In general, we want to draw approximately two and a half interactive spells per game, and at least one. On the play (where each card is more relevant and gains a power multiplier from the additional agency gained), we should expect to see eight draws in those nine turns, for a total of 15 cards. With seven interactive spells, we have a 83.2% chance of drawing at least two interactive spells at that depth, and only a 22.4% chance of drawing at least four, where four is the number that you'd probably prefer to have a creature over the fourth one. The distribution of large removal, small removal, and combat tricks is a product both of what you can expect to see in a draft, as well as having some diversity of responses to be prepared for more varied boardstates, while minimizing the probability of drawing multiple combat tricks to 13.5% by only running two.

With the number of interaction spells locked in, we can be a lot more specific about our 2-drop count. We have two constraints on how many 2-drops we want to draw:

1. We would like a 2-drop in our opening hand, as hands with 2-drops are considerably stronger than those without. Here we specifically want one in the first seven cards instead of the first eight or nine (the number of cards we see before casting our two drop), because we want knowledge of that 2-drop when deciding our mulligans.
2. We want to draw less than four 2-drops over the course of the game. Why four? The third 2-drop is still reasonable to fill in our curve, as Boros will often peak at four lands, letting us double spell our 2-drops. At the fourth one, our draws start to lose impact, and we'd prefer a higher curve card. We also only need to consider the first 14 out of 15 cards we see, since the 15th card would be a creature anyway (as we already locked in our interaction spell count, we're only shifting around our creature curve). A topdecked creature is largely less relevant on the game's final turn as the beatdown player, since it has no opportunity to attack.

So what are the odds of drawing a hand with a 2-drop in the opener but three or less in the top 14, what we'll call a "golden hand"? To calculate that, we need something stronger than either Hypergeometric Calculator: some code. Let's call the probability of satisfying the first condition $P(A)$, and the second condition $P(\text{Not } B)$, B being the probability of drawing four or more 2-drops

in the first 14. The probability of getting a draw which satisfies both conditions is:

$$P(C) = P(A \text{ and Not } B) = P(A) - P(A \cap B) \quad (3)$$

where $P(A \cap B)$ is the probability of BOTH A and B occurring. In human language, this means that out of all hands which draw a 2-drop in their opening seven, we exclude those in which we then also draw three more 2-drops in the next seven, to get the probability of our golden hand. In Fig. 6 we have a graph of what this looks like:

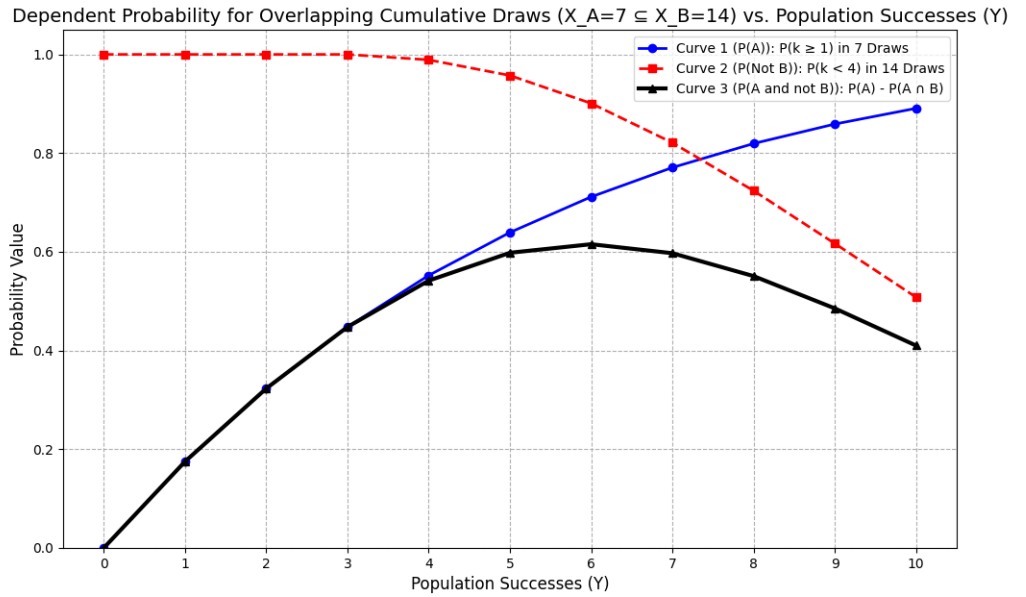


Figure 6: The probability of a golden hand (black), alongside the probability of each of its two components (red and blue respectively).

The golden hand is most likely at six 2-drops in the deck, which is the gold-standard number to run. Since condition A is more important than Not B , seven is also acceptable, with the occasional eight. Above that, you're hurting your late game draws too much for marginal mulliganing benefit. Below that, you risk having to mulligan too often, or accept more common slow draws.

I won't pretend each of the numbers in the Boros skeleton are so rigorously grounded. A lot is based in what has felt good in the past, but humans are remarkably bad at approximating probability distributions, so there could be some edge to be gained by repeating versions of this analysis for different mana values. One point of confidence for this approximate curve is that it's basically

a skewed Gaussian Distribution; if I was advising you to play one 4-drop and four 5-drops, that might be reason for concern. There are a few 5-drops that I'd occasionally encourage you to play above this skeleton's quantities, though...

2.4 Pseudobombs

The final concept to discuss before we jump into some specific deck archetypes is how aggressive decks utilize bombs. In general, a bomb is a card that will win the game if left unanswered, but that isn't a specific enough definition. A Storm Crow provides a 20 turn clock if unanswered, but nobody would unironically call it a bomb¹¹. I will impose two more constraints on bombs: they must meaningfully outperform the vanilla test, and their win-clock if unaddressed must be measured in less than three or so turns. Some bombs will be faster than that, and others will be slower, but the slower ones will generate an insurmountable lead in that time.

In aggro decks, bombs serve a very different role than in slower decks, because your win-clock is already measured in less than three turns if your primary strategy isn't addressed. Slow decks use bombs to overpower their opponents' fair Magic in the long game, or stabilize defensively in the shorter game. For aggro decks, a bomb serves to saturate opponents' removal. Removal is already taxed against aggressive decks, so opponents are less likely to have the answer to a bomb on turn 4 or 5, and if they do, that's mana and a piece of removal not directed towards staying alive.

Not every aggro deck opens a bomb, and that's fine, because bombs are a garnish on aggro's gameplan, whereas slower decks that see more of their decks are more likely to need bombs to close grindy matchups. But what if we could Moneyball a bomb? For those unfamiliar, Michael Lewis' *Moneyball* tells the story of how the Oakland A's, having freshly lost their star player Jason Giambi (among others), are able to recreate him in the aggregate by signing several less brilliant players who together could do an impression of his skills. Nothing in the role bombs play in aggro requires them to outperform the vanilla test, it's just nice to have.

In Murders at Karlov Manor, Boros had access to several more premier bombs than usual: Aurelia's Vindicator, Agrus Kos, Spirit of Justice, Aurelia, the Law Above, Wojek Investigator and more, to say nothing of aggressive non-Boros bombs like Ezrim, Agency Chief. The core of those aggro decks was primarily their efficient combat tricks that replaced themselves, like Auspicious Arrival, The Chase is On, and On the Job, so forking your opponent between answering a 2 for 1 combat trick and a game-winning bomb was a powerful strategy. But not every Boros seat opened a bomb, and to keep the gameplay of my Boros decks as smoothly similar to each other as possible, I would plug that hole with what I call a "Pseudobomb."

Pseudobombs are much weaker cards that demand an answer when played by

¹¹For further taxonomy of bombs, see "Making The Most Of Your Mythics".



Figure 7: “Are you the strongest because you’re a bomb? Or are you a bomb because you’re the strongest?”

an aggressive deck, in MKM’s case Haazda Vigilante and Rubblebelt Braggart. While these cards barely pass the vanilla test, they provide pressure, and can stretch removal thin just like a bomb can. In Avatar: The Last Airbender the corresponding Pseudobombs are Master Piandao, Combustion Man, Cruel Administrator, and maybe Beetle-Headed Merchants¹². Notice a few things all these Pseudobombs have in common:

1. They all have more than 3 toughness. This means that, while they’ll usually trade down in mana to removal, they won’t get killed by a Lightning Strike, so your opponent probably has to tap at least 3-4 mana to kill them before they...
2. Attack! Since commons and uncommons are less likely to have upkeep or endstep triggers than rares, the easiest way to design a Pseudobomb with persistent value is to give them an attack trigger. This means your

¹²This portion of the article is being written before TLA’s release, so the Merchants may not get there. Fortunately, “being good” is not a requirement for Pseudobomb status.

opponent has to spend their mana and removal spell before you pass to blockers, freeing you up to more safely use your combat tricks.

3. They all cost 5 mana. While cheaper cards are usually better in aggro decks, for the Pseudobomb's niche, it's nice that they don't conflict with your uniquely packed turn 4, during which you can cast a pair of 2-drops or a 2-drop and a (usually) 2 mana combat trick. By turn 5, it's more likely you've gotten all the 2-drops out of your hand and you won't be sacrificing the double spell to cast your Pseudobomb.

The biggest issue with Pseudobombs is that they tend to not leave any value behind when removed, unlike a lot of the most powerful bombs, so you want to be a bit more careful about how to play them out, to improve the odds you'll either get value from them attacking, or from the resulting combat trick or free attack when they do get removed. Once again, you can use your role in the game (pressuring life total) to force your opponent to misuse their resources (removal) and maximize your own (situational 5 drops that need to attack and/or vulnerable combat tricks).

With these general Boros considerations out of the way, let's talk about three Boros decks that execute that strategy in particularly interesting ways.

3 Three Decks

3.1 The Single Quadrant Deck

Quadrant Theory, developed by Brian Wong and popularized by Marshall Sutcliffe, is a foundational card evaluation heuristic for modern limited play. It states that a card can perform differently in four game states, or "quadrants:" Opening/Developing, Parity, Winning, and Losing. One common trap new limited players fall into is to only evaluate cards in a specific quadrant, most commonly Winning. In the Single Quadrant Deck, we intentionally fall into this trap.

This deck is born of necessity. Imagine a scenario in which your draft is going poorly, and you've found your way into an open Boros lane (as this entire article presupposes - don't draft Boros if it's closed!), but power just hasn't been opened at your table. While you are wheeling plenty of C tier creatures and getting last-pick 23rd playables, your most powerful cards are uninspiring, and most importantly, your removal suite is barren. This is a common scenario in high power draft pods, as removal gets snapped up aggressively, and it can feel difficult to escape. Especially if you open off-color bombs in the third pack which you have no way to remove, it can be tempting to hate-draft them, since you can't imagine your mediocre deck beating them, and you already have plenty of playables. Unfortunately, early hate-drafting is usually a mistake without very solid information, as you're not guaranteed to encounter that bomb if you pass

it, and you have to pass up a valuable early pick.¹³ Fortunately, you're playing Boros, and Boros has access to the best removal: player removal.

When you notice you're falling into such a seat, you should immediately reframe your perception. From now on, you are always in the Winning quadrant¹⁴, and every pick should serve to move you further into it. In exchange, you accept that if your deck ever falls behind, you will lose (unlike both decks discussed later, which have a lot more fight to them from behind). From now on, you will ignore any games in which you fall behind, and focus on maximizing your winrate in games in which you start ahead, as well as maximizing the raw number of games in which you start ahead. From this mindset, several details change in my card evaluation.

1. 2 drops become more important. Above, I mentioned that I always want at least six 2-drops in a Boros deck. In the Single-Quadrant Deck, you should try to play 7 to 9. This deck cannot win without being in the Winning quadrant by turn four, so you will not keep many hands without 2-drops, and it is more important to have one in your opening hand: eight 2-drops gives you an 81.9% chance of having one. Mulliganing is quite punishing for this deck, as it is easy to not draw the pieces you need, and as soon as you stumble you lose, so maximizing your range of keepable hands is also extra important.
2. Fixing lands, especially untapped, become more important, and you also gain more leeway to pick them up early. As stated before, to be desperate enough to draft this deck, you must not have too many powerful individual cards, or it would be unnecessary to rely so heavily on a single gameplan. Therefore, the Boros lane must be heavily open, or else you would have no bombs AND no cards, and be in a different color combination entirely. Between an open lane and a willingness to pick up 4-5 combat tricks, this deck should have access to 28+ playables at the end of the draft. Through that lens, it's less costly to use an early pick on fixing, since those extra 5 playables do nothing for you unless there's a specific card you need to sideboard out. Furthermore, since this deck has so little room for error without finding yourself in the Losing quadrant, stumbling on mana is disastrous. With so many 2-drops, you're also more likely than the average aggro deck to cut a land, which makes dual lands even more important for getting 9 sources in at least one color, as discussed in Section 2.3. Fixing lands aren't all upside, though. You are limited to one or sometimes two fixing lands in a deck, as missing your 2-drop to a tapped land is unacceptable, and not getting a 1-drop down until turn 3

¹³When piloting a slower deck with some amount of fixing, speculative hate drafts which you might yourself splash are an entirely different story, even in pack 3. Now, you're both denying an opponent a threat and potentially powering up your deck. Just don't sunk-cost yourself into playing the card if the fixing fails to materialize.

¹⁴Or Opening/Developing, but you should also try to shorten that quadrant as much as possible.

can be problematic in some sets.¹⁵

3. Most importantly, since you don't have enough removal, combat tricks become your removal. Every real game begins with you being ahead, meaning your opponent always has to play from behind. There are two ways for them to do so: playing blockers, and removing your creatures. In both of these cases, they are heavily incentivized to tap out on their turn (since holding up removal leads the risk of getting blown out by a protection spell and taking a full turn of damage for free). This makes your combat tricks much safer to cast, already increasing their value.



Figure 8: Run Amok's text in a Single-Quadrant deck. Sure, their blocker soaks up some of the damage, but it feels this good.

Your opponent is also incentivized to block aggressively — playing a bomb or utility creature and then not blocking with it is equivalent to having spent their turn doing nothing, if you can kill them in the next few turn cycles!¹⁶ Good players will try to play around combat tricks, but eventually, the best players know to sometimes “make you have it,” as tanking every attack means dying. It is your job to always have it, in multiples if possible; getting hit by the third combat trick can be very salt-inducing. I've won plenty of games in which my opponent let through a creature on low health, deciding that there was no way I had another combat trick. Combat tricks which grant either trample, to continue to push damage, or a persistent effect, like a flashed in aura,

¹⁵The Single-Quadrant deck does not agree that 1-drops cost zero.

¹⁶No, I did not use ChatGPT to write this. Em-Dashes are just great.

gain particular extra value, as do cheap tricks which allow you to trade up in mana while developing your board with another 2-drop - if your 1 mana combat trick kills an opponent's 3-drop, you've just contracted their available resources substantially. One repeated example of these effects are 2 mana +3/+3 and trample tricks — Blitzball Shot, Fanatical Strength, Staggering Size, and the identical Run Amok, since you're only casting these during combat anyway. Another is a permanent +1/+0 and temporary first strike: the aforementioned Squire's Lightblade, Coral Sword, Military Discipline, somewhat Zealot's Conviction. Notice that all of these are low picks which you can frequently wheel, but will often play better than pick 1-3s like Banishing Light effects in these decks.

On the other hand, some cards lose value in this archetype.

1. High mana value creatures become less effective. Where in a regular aggro deck, a 4-drop allows you to follow up on earlier pressure and close games, in the Single Quadrant deck, you want to be casting a combat trick on turn 4, often alongside another 2 drop. The same applies to 5-drops and higher (goodbye Pseudobombs); you really want to trim down your curve as much as possible. This has the added benefit of making it easier to go down to 16 or even occasionally 15 lands than in most decks.
2. Grindy or card advantage cards lose value. In general, I am a fan of running a few card advantage pieces in aggressive decks, to maintain a balanced draw. If the rest of your deck is cheap, efficient pieces, you can afford to spend some mana to draw a few more of them. If you only have one or two of them, there's little cost to draw extras, as discussed in Section 2.3. Here, even that goes out the window. The only card advantage allowed is from creatures that generate card advantage as they attack or combat tricks that draw on resolution.
3. You have less reach — the ability to close games in which your opponent is on low health. This one sounds a bit weird, because this deck's whole purpose is to close games quickly, but the problem lies in what cards tend to provide reach. The three easiest ways to close a game are to give one of your creatures evasion and swing with it, direct a burn spell face, or to cast several removal or tap spells in a turn and swing around your opponent's remaining creatures. You can also play haste creatures, but that tends to be more telegraphed. Unfortunately, these all become rarer in the Single Quadrant deck.

Evasive threats, especially persistent evasive threats as mentioned in Section 2.3, are a bit slower than this deck can leverage well, since using a turn to develop one and having it get removed can be enough to shut you out when your topdecks are all low-quality. Burn spells double as removal, and you definitionally are low on removal, so you'll more often have to direct what burn you have at creatures. This also shuts you out of removing two creatures for a big swing. By using combat tricks instead of removal,

you gain the ability to stack several combat tricks on one unblocked creature, which becomes your primary way to generate reach, but it's a lot more vulnerable than the other options, so you have to be careful if your opponent has mana up. This compounds with trample tricks being particularly effective as discussed above, but double strike tricks are also quite good in conjunction with other tricks.

The Single Quadrant deck makes the best use of the concept that **aggro is card advantage**. If your opponent dies before they can cast 3 cards in their hand, they have just mulliganed 3 times, and even your piddly 2 drops can compete with that kind of value. A Vaultborn Tyrant will usually win the game against an aggro deck, but if your opponent never hits 7 mana, you've just Thoughtseized it for free. This is why I'm willing to make concessions which make my deck composition, curve, and answer distribution worse in the Single Quadrant deck: if every turn that passes is an extra opportunity for my opponent to stabilize, I'd rather take a more dangerous line that kills them faster than develop my 4-drop into their potentially stronger 4-drop. Some aggro decks can fight for value with their opponent's hand empty, though, like the other two deck archetypes below.

Perturbations of the generic Boros skeleton include:

- +2 2-drops
- +1 1-drop
- +3 combat tricks
- -1 5-drop
- -2 4-drops
- -3 large removal spells

Once again, these perturbations will vary to your draft.

3.2 The Board Wipe Aggro Deck

It is common knowledge that you should not play board wipes in aggro decks. This may be the first lesson that a lot of limited players are taught about how you should try to draft a deck rather than a pile of cards. I usually agree with the common knowledge, but in the last few years, there have been several formats with aggro decks that can happily play a board wipe. In modern limited, with the power level of commons increasing, aggro decks have largely seen their power shift into stronger combat tricks and more powerful individual-card grindy engines, becoming less glass-cannon (or less Single-Quadrant) on average.

There has also been a proliferation¹⁷ of board wipes which leave behind a creature. Some of these are pushed enough on rate to always be playable like

¹⁷More of that strange card design ... It's probably nothing.

Nibelheim Aflame, Unstable Glyphbridge, and Sunfall, but others are less so like Zero Point Ballad and Starfall Invocation. With good but not exemplary winrates, these latter board wipes would generally not be played in aggro decks, but with a bit of value mixed in to your aggro deck, they open up a unique playstyle.

Consider the following scenario. You have been curving out in a regular aggro draw, getting your opponent down to 6-10 life or so, but they've managed to stabilize with two midsized creatures that can eat your four attackers. You swing with all four creatures, getting your opponent to 3 life, but entirely sacrificing your ability to progress your aggression. Another creature wouldn't push through their blockers, and while a combat trick might, it would be vulnerable to removal. You pass the turn without playing either of your 2 cards left in hand, eyeing their life and sighing — perhaps you ham it up and mutter “Blooming Blast” under your breath — and pass the turn back. Your posture is slumped; you are relying on topdecks to carry you over the finish line. Given this information, what should your opponent do?

The correct answer is to develop, and develop as quickly as possible. The sooner they have two to three more creatures than you on board, the sooner they can start attacking, and close out the game before you can topdeck your reach, whether the implied burn spell or a flying threat. They've confirmed you don't currently have the tools in hand to swing the board, as combat tricks, removal spells, and haste creatures cannot be in your current range. As soon as they drop a few more creatures, you reveal Starfall Invocation, and reset the board with one creature left on your side.

From here, you've effectively restarted the game with a few changed parameters. Both players start the “new game” with 5+ lands in play, which should ostensibly favor your opponent, as their deck is more geared towards the late game. At the same time they've lost a lot of life, so your smaller creatures have to connect fewer times to win, effectively increasing their stats. By 2-for-1ing (or better) your opponent, as well as keeping a Veteran Guardmouse in your hand you would otherwise play out, you've somewhat equalized the value game, especially since their creatures were more expensive than yours. You also have the first creature on the board, ready to attack next turn, and your opponent has already played out their best immediate board threats, so there's a real chance you get free attacks in while your opponent is forced to play draw spells or other setup cards.¹⁸ The best part is, your opponent played exactly how they should have, meaning this strategy can work at the highest levels of play, because **good players know not to play around board wipes from aggro decks**. From any other deck, passing the turn without doing anything should immediately raise red flags, but against aggro, the margin a slow deck can gain is from punishing it when it stumbles.

I've enacted this play pattern many times, and there are even more reasons why it scales to your opponent's strength. In the next game, you can take the

¹⁸These scenarios have often looked strangely like early Blood Moon plays in constructed — one player attacking with a small creature, and their opponent passing back. Rinse and repeat. When your opponent is already low on life, it becomes even more threatening.



Figure 9: A false equivalency, but a tempting one. Starfall is actually far better than Karn, because who plays 7 mana cards in their aggro deck?

board wipe out of your deck entirely, but their very knowledge of its existence takes up valuable attention from them, and can make them slow their plays down to play around it, improving your actual aggro draws. Furthermore, grindy value pieces are the first resources to get traded away against aggro, since you don't need an overwhelming lategame advantage usually, only to survive long enough to stabilize. Opponents are more likely to trade off value creatures early, they will discard draw spells or mulligan them or side them out, and with every decision, they will try to make the less greedy play which lets them live another turn. Aggro makes your opponents use their resources less efficiently but more quickly, and board wipes can punish that.

What about regular board wipes with no upside? I'm far less willing to try this strategy with a typical Wrath of God, but on occasion I find a more midrange Boros deck which can support a version of this playstyle while occasionally grinding into the lategame. I would say it's rarely worth trying to warp your deck around a regular board wipe, but there is some margin to be gained in considering the possibility.

This strategy does have a few real and highly punishing weaknesses. It tends

to work much better against medium to slow midrange decks which are trying to kill you anyway than against true control decks, since if they have countermagic in hand the equation changes, and they can sit back for as long as they'd like to, which instantly loses the game for you. True control is rarer in modern limited than historically, with soup being a more popular late game option that can outgrind those slow decks, but you do come across a draw-go deck from time to time. Giving your opponent time to cast their draw spells that they happen to have in hand which would otherwise be useless can also be an issue if they draw enough material to stabilize. The biggest issue, however, is what kept board wipes out of limited aggro decks for years, which is why your opponents know not to play around them: they're a major quadrant hindrance. In either developing or winning quadrants, board wipes sit dead in your hand, effectively mulliganing your aggressive start, and ideally you'd want to avoid being in the losing quadrant where board wipes are designed to be used (see Section 3.1). In an archetype that can't afford to stumble, sabotaging your primary gameplan to enable a finnick, bluffy backup plan which can fall on its face is a pretty daunting proposition. Especially if players in your area know you play board wipes in your aggro decks.

Still, the higher level Magic you're playing, the more you need to squeeze every last drop of value out of every card you open, and the weirder your decks have to get. Many recent boros mechanics have had grindy options, from equipment, to valiant, to tap your untapped creatures, to get rewards for tapped creatures. Being aware of the board wipe deck helps you use these grindy cards effectively in an aggressive strategy. But there is an even more commonly effective way to do so, the final deck archetype we'll discuss.

Perturbations of the generic Boros skeleton include:

- +1 board wipe
- +1 5-drop
- -1 1-drop
- -1 combat trick

It's basically a Boros aggro deck with an extra comeback card, that happens to care a bit more about its lategame threats and topdecks.

3.3 The Equipment Deck

As mentioned in Section 2.3, creature-equipments allow for Boros decks which run many equipments without diluting their creature pool. In some sets like Phyrexia: All Will Be One and Final Fantasy, equipment decks are the primary Boros archetypes, and they have finally become stably functional there due to equipment creatures. Other sets include cycles or individual copies of equipment creatures, such as Kaldheim's Valkyrie's Sword cycle or Adventures in the Forgotten Realms' Goblin Morningstar. Cheap equipment creatures tend to

be costed around half a mana higher what they would be if they were just a creature, even at common, when comparing to common vanilla test metrics. For example, Barbed Batterfist would be worth 1.5-1.75 mana as a 2 mana 3/1. More expensive or lower rate equipment creatures rarely cost more than 1 more mana than they would as a creature, like Vulshok Splitter at 3 or Dwarven Hammer at 4, but these tend to make up for the mana disadvantage by providing a more impactful equipment later, effectively two cheaper equipments (albeit without the two tokens attached to those cheaper equipments). This means that replacing typical creatures with equipment creatures in your curve doesn't hamper your aggressive gameplan too much, but it does open up an entirely new gameplan lategame.

At some point, if the board has stabilized (ideally with your opponent at low life), you can start stacking your equipments on a single equipment-generated token, and attacking. Your opponent is forced to block with multiple creatures, trading several resources for a single token, giving you card advantage every time you do so at the expense of tempo and mana. By tailoring what you equip to your opponents' blockers, you can set up profitable attacks through even decently clogged boards, and as an aggro deck, your opponent's board can't be too clogged since they're forced to trade creatures off early. Having Glimmerlight in play allows you to decide down to the mana how many equipments you want to throw onto your attacker.

Here the same factors work in your favor as in Section 3.2, in that your opponent is incentivized by your aggro start to sacrifice value in order to stay alive early, but the equipment deck has other tools to buy back time during this board state:

- Lifelink equipments and creatures like Dissection Tools and Adelbert Steiner can grow your health total higher than any other playable limited mechanic, which makes it difficult to turn the corner against you and punish you for spending 5+ mana to suit up a single token. If your opponent tries to race your equipped creature, a single turn of holding your creature back can stonewall their aggression, followed by an attack and a re-equip to constantly keep your blocker up (if you don't already have an equipment granting vigilance).
- Evasive equipments (as mentioned in Section 2.3) and creatures which take equipments well like Swooping Lookout fight on a different axis from ground attacking creatures, but unlike in a typical Boros deck, you can grow your fliers far beyond your opponent's fliers, making removal their only answer. If the evasion comes from an equipment, not even removal is an answer for long.
- By using your turns on equip costs, you can build up actual plays in your hand, allowing you to cast several real spells in a turn if tempo swings too far in your opponent's favor. This is the same concept behind "screw beats flood,"¹⁹ in that you are artificially mana-screwing yourself by spending

¹⁹Debatably true.

your turns on overcosted equips, before re-emptying your hand a few turns down the line. Similarly, your opponent is artificially flooding themselves by trading off multiple cards worth of material for your fractions of a card. Since paying equip costs is generally less efficient than just casting those spells in the first place, this is only worth doing if reequipping tangibly opens attacks for the next two turns which playing out new threats would not, or if you suspect your opponent can stabilize defensively in the next turn.

There are a few ways your opponents could beat the late game equipment plan once it gets rolling, none easy while at low health themselves. They could cast removal on every creature you play after you equip them, but eventually they need to kill you or you can keep drawing creatures and equipping them, forcing them off kilter like in Section 3.2. They could also develop a more powerful value engine than your equipment stacking and kill you in the air while repeatedly chumping if you don't have an evasive equipment or lifelink creature. Crucially, no individual nonengine creature can stop the equipment-stacking for long, since with enough equipments, you can make your tokens the largest threats on the board!




Figure 10: An avengers level threat.

Since both ways for your opponent to outpace your clunky equipment stacking involve tempoing you out, small removal gains value in equipment aggro decks compared to a generic Boros deck, as it can be used either offensively in the early turns, or defensively later. The typical large creatures you would ordinarily

need large removal to get past are smaller than a souped-up token. Just like a regular Boros deck facing an opposing 2 drop in the early game, you can wait a turn and start attacking again. Combat tricks also lose value, as equipments are basically sorcery-speed combat tricks, although evasive combat tricks like Jump Scare gain a bit of value to sneak one large creature through in the end game if you don't have access to a flying equipment.

While equipment-creatures almost stack up against their early game creature counterparts, every bit of power dilution matters more for an early game creature expected to attack multiple times, so you do lose a meaningful amount of explosiveness in the early game by playing so many equipments. Aggro's strength is in depriving your opponent of the opportunity to use their cards properly, and when you take your foot off the gas by playing a slightly subpar creature or spending the turn re-equipping a Warrior's Sword for 5, you sap that strength rapidly. This is particularly pronounced in the aggro mirror, where it can be difficult to find time to reequip your dead equipment creatures, meaning you'll simply play the game down resources against your non-equipment-using opponent. If you can't stay in the Winning quadrant and force trades to reuse your equipment value, you'll rapidly lose board presence to generic aggro creatures. Fortunately, most sets with enough equipment creatures to build a deck around them don't have an alternative aggro deck in Boros, but this is a real issue when playing against non-Boros aggro decks, especially green-base ones which should already outstat your creatures.

Regardless, the sheer flexibility of having both an above average early game and the most inevitable late game in several formats is enough to turn the previously inconsistent Boros equipment archetype around entirely. This is especially true since you can use your control over the early game to transition the gamestate favorably into the late game by forcing tough early decisions from your opponent.

Perturbations of the generic Boros skeleton include:

- -2 combat tricks
- -1 large removal spell
- +3 small removal spells
- No other changes, your creatures are just equipments.
- Try to get 7 of them I guess?

4 Conclusion

These have been my musings on Boros and aggro in general in modern limited. I hope they can help any readers build more functional and punishing aggro decks in the future. I've repeated the largest theme of this article several times, that you can catch your opponents offguard and squander their resources with sufficiently threatening aggro play. An undercurrent to this theme, evident

with every involved piece of math analysis and deep dive, is that there's always more to think about in any situation you can find yourself in, and the best way to squeeze equity from that stone is to always be thinking about where that depth is. Distracted play is all too tempting, but if you can sleep well before tournaments²⁰, eat throughout the day, and maintain your posture, you'll find the focus you gain will reap major dividends, despite aggro's high floor. Remember that life is a resource, and your opponents have none!

Large swaths of this text came from long discussions over limited theory with other Seedcore members, particularly Liam Etelson and Neil Estrada. Thanks to everyone who read over this text and gave me suggestions, several of which I expanded into full sections.

A Appendix: A Boros story

In the 2024 Las Vegas Duskmourn Limited 100k, I reached the finals playing the following decks:

- Blue-White in Sealed, 7-1
- Blue splashing White and Red in Draft 1, 3-0
- White-Red in Draft 2, 2-1
- White-Red in Top 8 Draft, 2-1

In the days after the event, it struck me that I had only played Jeskai for the entire weekend. That I hadn't touched black was no coincidence, as it was definitely a step down from the other colors in DSK. While I wasn't actively avoiding it, if I had ended up in that lane once or more, it was less likely I would have had such a high variance event. But green was quite strong, including the powerful Spineseeker Centipede at common. I wasn't actively avoiding green, I merely felt I was "drafting the hard way" and taking what was opened.

It was my opponents who were favoring green over white, since there were only 2 playable white color combinations: white-red and white-blue, both of which had high winrates on Arena. According to Mengu and other players whose accounts on the top tables I listened to, the fear of getting forced into a barely functional white-black or white-green lane was too great to bite at the late white cards, so I was left with two powerful Boros seats. Since then, I've played a lot of Boros in internal limited testing, and I would say I have a preference for it (I have 2.5% higher winrate in Boros than my overall Premier Draft winrate all time), but despite this piece being all about how to play Boros, you should always be ready to draft whatever seat you open. The other players in the 100k could just as well have biased **towards** Boros, in which case I'd hopefully have picked up on the open green signals.

²⁰I'm not a great model for this one, but I try.

