Cursed Problems in Game Design

•••

Alex Jaffe | GDC 2019 | @blinkity

Hello reader! I'm annotating each slide with roughly what I said in the talk. **BUT:** if you'd like to know more, I've got a bunch of extra slides in the back, including discussion of further cursed problems. Check 'em out!

Welcome to GDC!

where we make good games gooder™

With talks like...

- How I balanced my shotgun in 90 days and you can too
- Quest markers that even filthy casuals can follow
- 8 hot loot table distributions for 2019

Welcome to GDC! Where we teach you how to make good games gooder.

Talks like these are actually essential. They break incredible ground on difficult concepts, and show us how to make a more perfect game, a better player experience.

I love these design problems. But in many ways they're like pearl diving. Search for the shiniest, best solution, from a sea of possible solutions.

I've found myself particularly interested in problems that are more like wandering the desert. Searching for months, even years for a solution, and possibly finding none at all.

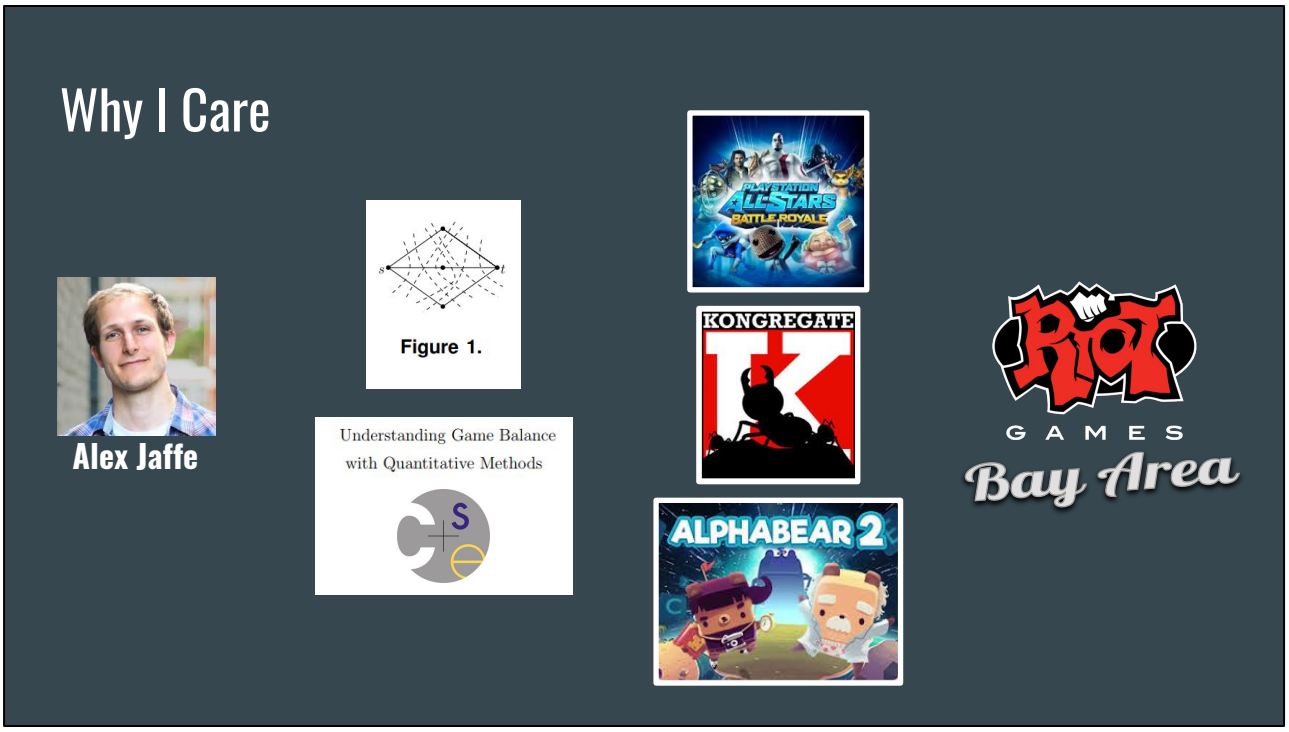

My interest in these problems comes from my own life.

For a few years, I was a mathematician. Mathematicians love hard problems. They work on a single problem for months, or even years, sometimes making no progress. You develop a kind of Stockholm Syndrome. And you start to realize that you're learning so much about the domain from the problem itself, even when no solution presents itself.

I carried this love of hard problems through my work in my PhD on computational game balancing, my design and data science work at Superbot, Kongregate, and Spry Fox, and now to my design work in R&D at Riot (in the Bay Area office). Throughout, I've wondered whether game designers as a group could spend more time thinking about the nature of our hard problems.

So I've thought a lot about the many seemingly disparate hard problems in game design, and tried to understand how they arise, what they have in common. And the answer is, I believe, a surprising amount. What they have in common is that they are "cursed". Don't worry, definitions and examples soon. But first an outline.

Outline

Part 1: Cursed Problems and Where to Find Them

- Defining Cursed Problems
- An Example Problem
- Identifying Cursed Problems

Part 2: Defense Against Cursed Problems

- A Game Model
- The Four Core Techniques
- Further Applications

In the first part of this talk, I'll try to build an understanding for what this class of hard problems really is, and how to recognize them. We'll see that these problems really are impossible - in a sense. And that's why this is the most important section of the talk. You can destroy a tremendous amount of team effort trying to solve unsolvable problems. So you want to see them coming a mile away. I'll define the category of problems, give you an example, then help you understand them better so you can identify them yourself.

In the second part, I'll tell you what you *can* do about cursed problems. It's not that you're just screwed - there's a lot you can do, and games throughout history have done so. But doing so requires making sacrifices to the core conceit of your game, however subtle. That's what part 2 is about. I present a model for what a game is. I'll use that model to describe the four techniques that pretty much every game throughout history has used to address cursed problems. Then I'll show additional applications of those techniques, so you can build intuition for how to use them yourself.

Outline

Part 1: Cursed Problems and Where to Find Them

- Defining Cursed Problems
- An Example Problem
- Identifying Cursed Problems

Part 2: Defense Against Cursed Problems

- A Game Model
- The Four Core Techniques
- Further Applications

On to the definition.

See Zak Mclendon's talk "Welcome to the Yard Sale" from Practice 2018

player promises

a game's essential experiences. why the player came to play.

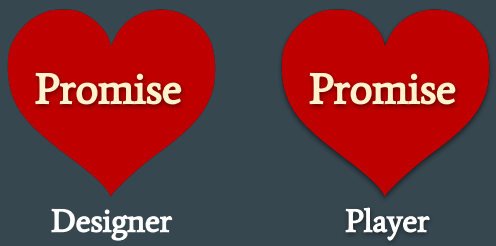

For the definition, we need a sub-definition. This is a useful concept, and you can learn more about it from Mclendon's talk. But for our purposes, an important note is that player promises often live in the heart of the designer. They're the essential aspects of the experience we want to create for the world. But oftentimes the promises can also live in a player's heart. Even if the designer never intended it, choices they make it in the core design of the game naturally "bring about" certain promises in many players. And we can't just break those promises haphazardly. We have to respect them and think carefully about them, even if we never meant to make them.

cursed problem

an unsolvable design problem, rooted in a conflict between core player promises.

No direct solution.

We win by giving up.

So now, a cursed problem. We'll come back to this definition repeatedly.

You can't just solve a cursed problem. But somehow we have to survive them, and we do. Like the Gordian knot, you win by giving up. It's an exercise in mitigation or compromise. This is a talk about what these problems really are, what defines, them, how to recognize them, and only then, how to think about solving them.

Outline

Part 1: Cursed Problems and Where to Find Them

- Defining Cursed Problems
- An Example Problem
- Identifying Cursed Problems

Part 2: Defense Against Cursed Problems

- A Game Model
- The Four Core Techniques
- Further Applications

We finally get an example!

Free-For-All Politics
A Cursed Classic

I’m going to walk you through an example cursed problem. It may seem a bit hypothetical at first, but it’s actually quite grounded in reality.

A cursed problem often
begins with a dream.

Imagine this dream. It’s the 90s. You love beat’em up games. Mashing on a bunch of enemies, kicking ass, crowd control. And you love Street Fighter, the head-to-head mind-reading, the deep systems, the skillful execution. And you think: what if we combined them? What if those enemies were other players, and I could harness my tactics and quick wits against several, simultaneously, kicking ass and doing crowd control? Like an old kung fu movie brawl. A brawler let’s say.

And you want to make it competitive - a journey of mastery. I know Final Fight isn't the most natural starting point here, but bear with me.

You make the game, and it's pretty fun. But pretty soon, something weird happens. You realize that the supposedly best player, the one who has great tactics and execution and prediction, isn't winning. The game has become something else. It's about politics.

See “Characteristics of Games”
by Elias, Garfield, and Gutschera

politics

competition through social negotiation, alliances, and manipulation.

For our purposes, here’s a definition of politics in games. You can learn a lot more about it from “Characteristics of Games”, a phenomenal and under-appreciated game design book.

Applied to our hypothetical brawler, we end up with a game *directed* first and foremost by this game of thrones chicanery. Players playing down to not get targeted, players ganging up on one another, making alliances, pleading not to be targeted, enacting revenge. You can't just kick ass at the fighting and expect to win - success is determined by your ability to politic. The game is not the game we intended to make. So what's happening here?

cursed problem

an unsolvable design problem, rooted in a conflict between core player promises.

Let’s look at the definition again.

Incompatible Player Promises

Free-for-all PvP games with high-skill gameplay

FFA Politics Problem

In a sense, this political focus is an inevitability! There are core player promises in conflict here, borne from the very premise of the game.

You can't make this game, at least with all this implicit promises associated.

Free-for-all PvP games with high-skill gameplay

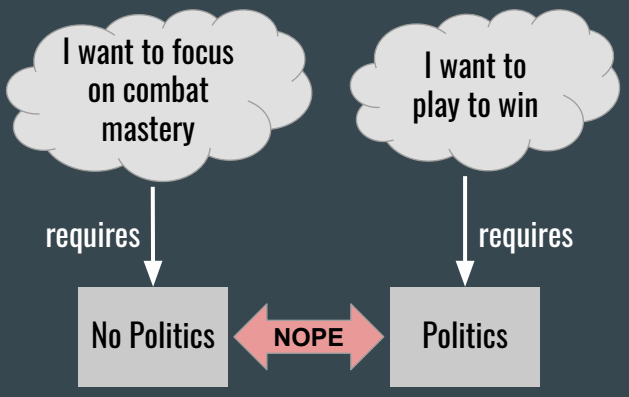

FFA Politics Problem

cursed

So it seems that the core promises of this hypothetical game are in conflict, because of what they imply. The problem is cursed. Which isn't to say there's nothing we can do! Only that doing something will require giving up some aspect (possibly implicit) of one of the promises.

This game is a fantasy.

So what now?

We let go of the dream. A little.

So what do we do? I told you this is a talk about giving something up. And you do have to give something up. But it doesn't have to make your game worse. Your game might be better as a consequence. But you have to let some of the dream go. And find a different, related dream, that doesn't put player promises into conflict.

You let go of the dream. A little. Okay, sometimes a lot.

But then you can find a design space for a game you CAN make.

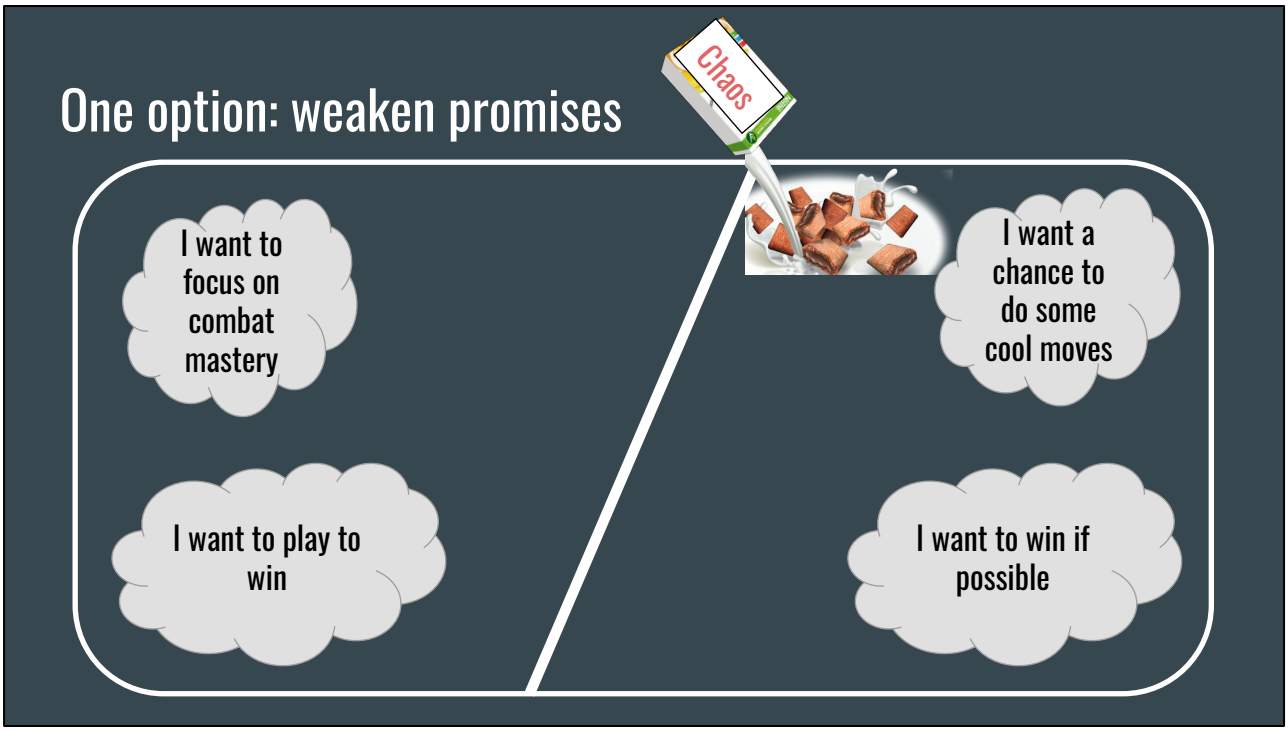

One thing you can do is soften both promises. Say, hey, you can play to win, but you can't control it that much. And you can focus on combat, but not all that much. You do this by inserting chaos. Changing the game in this way naturally changes the player's (and designer's) understanding of what the game is all about.

And as a consequence, you get something even better.

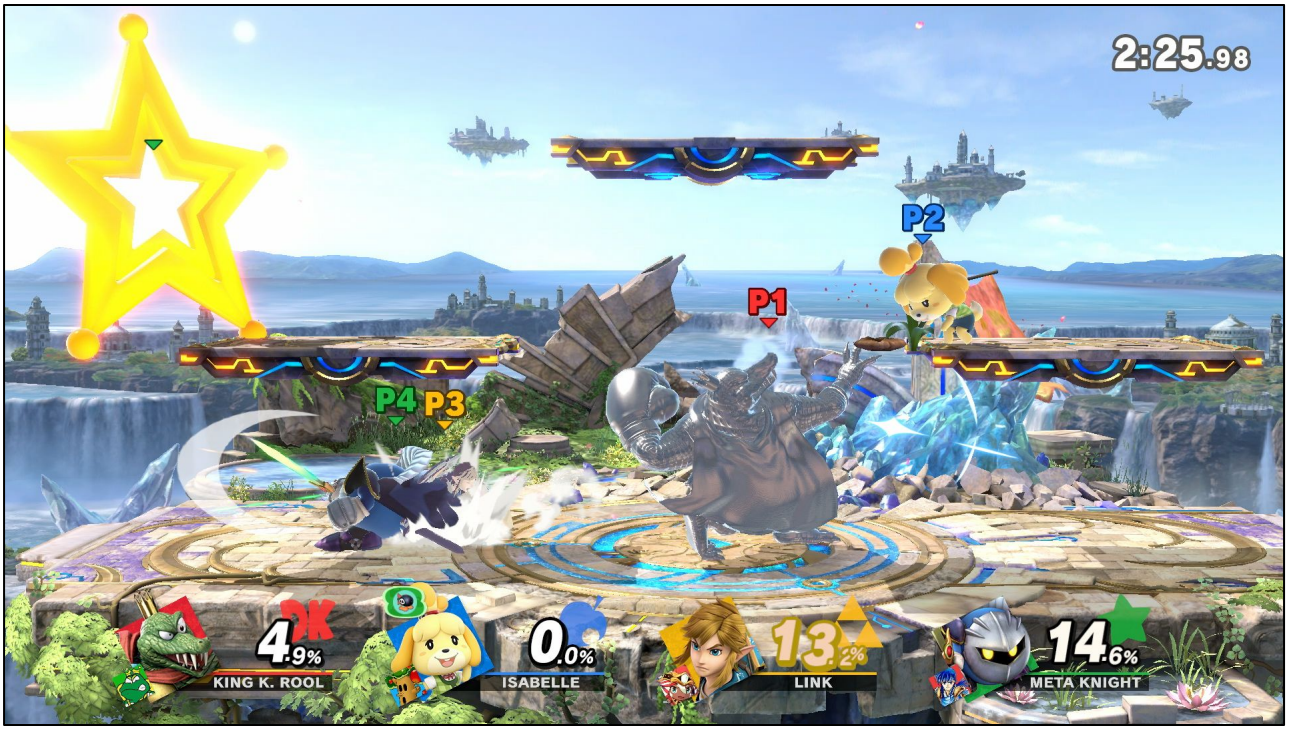

You get Super Smash Bros. One of the most beloved games of all time. And I'm not talking about the hyper-competitive stripped down 1v1 mode, though that's great too. I'm talking about the fun, low-stakes party FFA game, with stage hazards, random items, smash balls, high-variance attacks, etc.

It's not likely Nintendo was approaching the problem this way, seeing a cursed problem and finding their way through it. Their values probably led them the developers a little more directly to Smash. Many of the games we love are implicit or explicit answers to cursed problems - that's why they succeed. But that's good fortune for them.

Because if they'd pursued the fantasy I outlined initially, they would've found themselves cursed to wander the desert, trying to tackle a cursed problem. They'd hopefully find their way out eventually. But without the awareness that they were in a cursed problem, they'd spend a lot of time there. And worse yet, they'd be dragging a whole lot of very tired developers along with them.

You want to to recognize cursed problems before they really cost you.

And you want to understand your options to move forward.

So you really want to recognize these problems before they cost you. And if you're dead set on facing it, I want you to have the tools to more forward, despite the lack of a true solution.

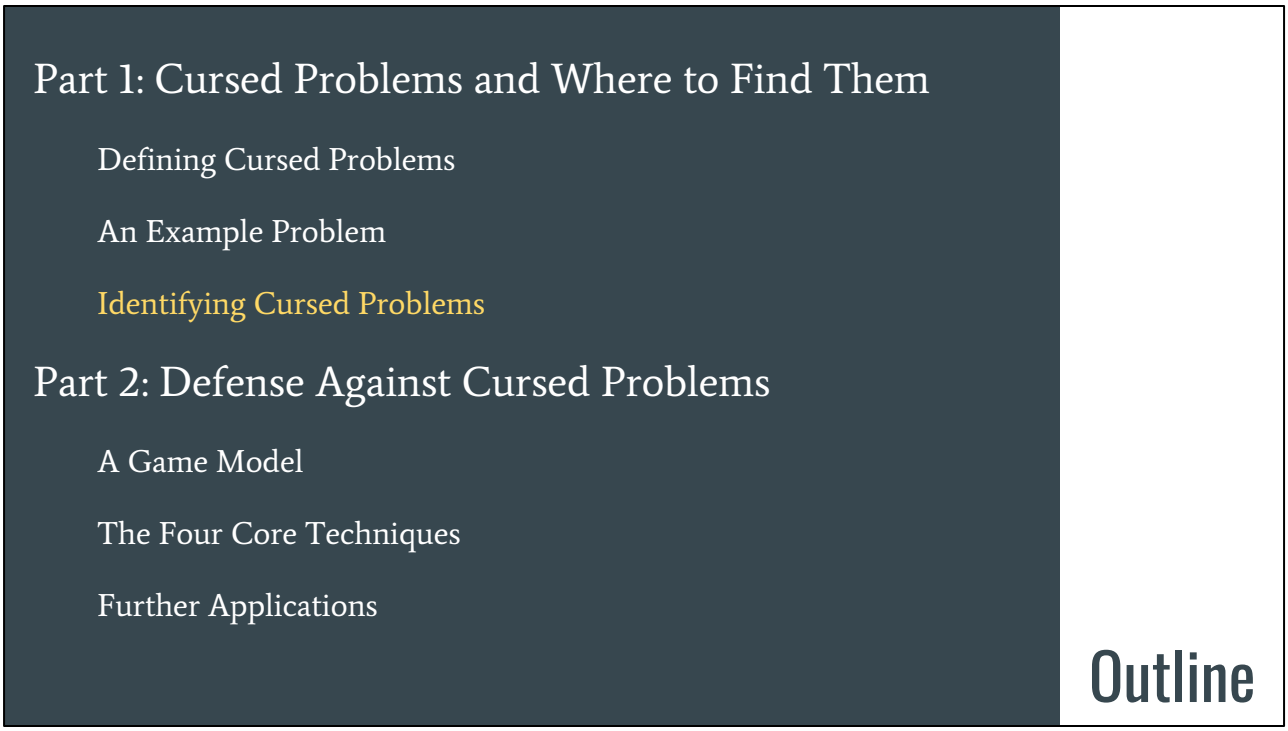

Now I'm going to take you through examples to help expose a general model, so you can build intuition for it, and look for cursed problems.

Note that I'll talk about some amazing games. Some of them are *answers* to cursed problems. Some still suffer from them, not necessarily hitting every aspect of their potential. This can be true even if the game is wildly successful! There's always room for a more perfect game.

And yes, this tells us that cursed problems don't necessarily kill games. Often they can survive just fine with them at their core. But sometimes the focus of the curse is something we really care about personally, sometimes it *is* severe enough to kill our game, and sometimes the biggest risk is that we'll spend too much effort trying to solve it.

Two More Cursed Problems

I’m going to show you some problems you’re hopefully familiar with. They’re very difficult to solve, they appear quite commonly, and there’s something somewhat inviolable about them. You may think devs have answered them in the past. But when they have, it’s precisely because they’ve given something up.

Turn-based cooperative games

Quarterbacking Problem

Like the FFA politics problem, it's common that problems pit some need against winning; against explicit goals set by the designer.

For instance, this happens in many turn-based co-op games, in the form of quarterbacking. Players have a good ol' time playing a co-op game like Pandemic, each doing their part in an Ocean's 11-like fantasy of cooperation. But over time, one player starts to make suggestions, the other players accept them (because they're good suggestions), and eventually one player is doing all the thinking. It's like a single-player game with many pairs of hands.

This is a natural tension between promises. We want a cooperative, harmonious experience, but we also want to win.

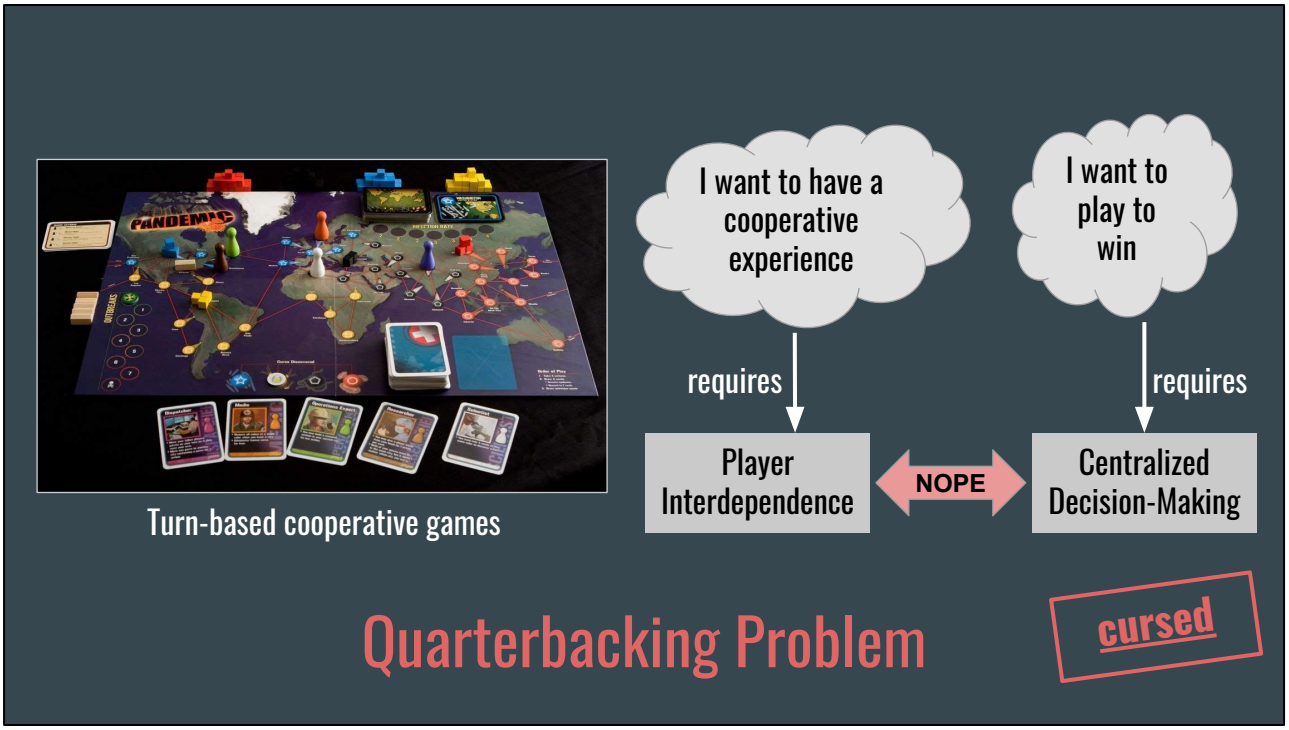

But true cooperation involves interdependence between players, not just being aligned. And playing to win in a typical turn-based co-op game is usually a matter of strategy. Unfortunately, strategy at this scale is usually best planned by a single centralized decision-maker. So you have a fundamental tension. Unless you give something up.

Evergreen competitive games

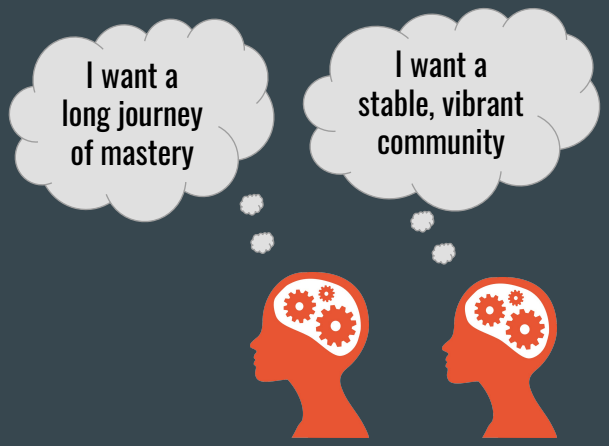

Skill Inflation Problem

When a competitive game comes out, it's a wide open field of play, experimentation, and learning. Over time, players start to get better. What's more, the ones who aren't getting better are more likely to drop out. By the time a game is mature, only killers remain. And it becomes very daunting for new players to enter. If you're lucky, there remain enough weak players to support matchmaking. But these players are often driven away, and even then, weak players are all too aware of the shadow you're living in, thanks to Twitch and smurf accounts.

The same can even apply to non-competitive games. This is a fundamental problem with long-term deep games. They promise me a long mastery path, and also a big engaged community. But the loss of the bottom end naturally causes the rest of it to slowly die out. It's hard to keep a game alive and thriving.

Evergreen competitive games

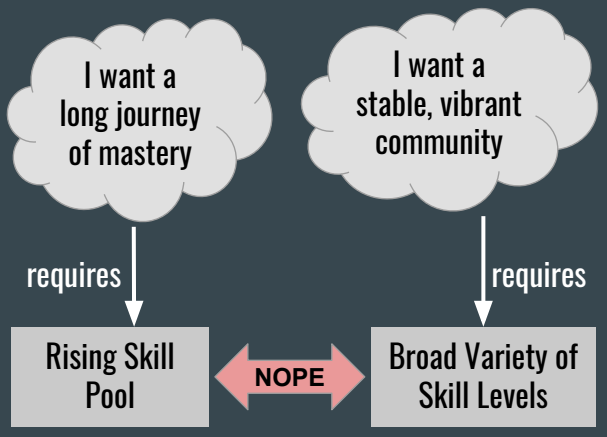

Skill Inflation Problem

cursed

Again this is due to a fundamental tension between promises. A given player, even a serious player, has violated promises. They want a long journey of mastery, which naturally leads to the skill pool rising. But they also want a stable, vibrant community, which is hard to pull off without a broad variety of skill levels. Fresh blood keeps the game fresh.

Developers certainly do deal with this problem, and not only through luck or money. But what do they have to give up to do it?

So what *isn't* a cursed problem?

Can we do this with FFA Politics?

Fix by tweaking

- Attacks
- Defense
- Movement
- Pace
- Map
- etc.

Example: Degenerate Core Gameplay Patterns

You may be wondering if every hard design problem is cursed. Nope. Consider something as simple as degenerate core gameplay patterns. Here's Hungrybox as Jigglypuff, in Smash Melee. (Now we're talking about the hypercompetitive 1v1 game, not the FFA party game.). Hungrybox's play style is often considered frustrating to play against. He's defensive, reactive, punishes a single mistake. He forces the opponent to play on his terms. It's *subjective* whether this is degenerate, but say you think it is. Are we in a similar situation? Is there anything we can do?

Sure there is! This playstyle is not a fundamental structural consequence of the core conceit of the game. It may be breaking a promise, but It doesn't take structural changes to fix. We can just take it out, by tweaking any number of things. Nor are violating a promise by taking it out. There's nothing fundamental about Smash that says you should be able to play such a reactive game.

Free-for-all PvP games with high-skill gameplay

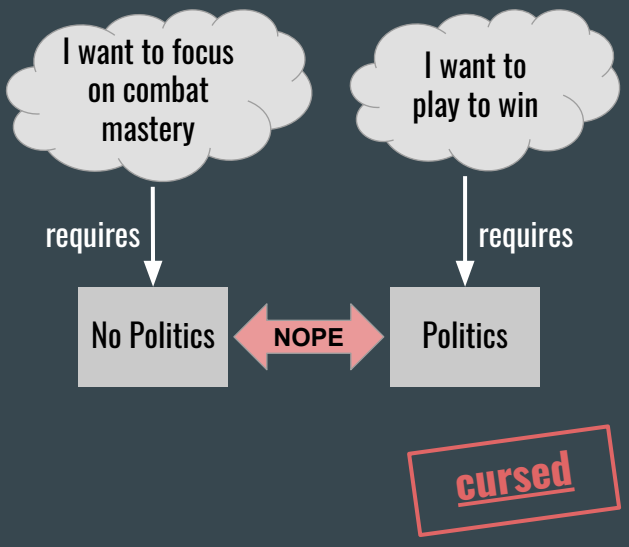

But if we go back to FFA Politics for comparison...Basically cursed problems are traps we make for ourselves. We make multiple promises, but we don't fully understand their implications, and how they conflict. There's no unilateral improvement of one aspect of the game that will make them not conflict.

Let’s play a game.

“Hard or Cursed?”

I’m going to show you three problems - real challenges encountered by big games. I want you to decide whether it’s a cursed problem or just a hard design problem. This isn’t easy - it takes a lot of thought! You have to identify two promises, and make a compelling argument that they’re inherently in conflict, rather than presenting some tension that easily supports a “have it both ways” solution. And of course, there’s some ambiguity; it depends on how you frame it. But I think reasonable people can *typiecally* agree, with training.

Exploration games with millions of worlds

Millions of worlds to explore

Vibrant, diverse, interactive ecosystems

How about this problem? No Man’s Sky promised a game of limitless exploration - out in space with your spaceship, uncovering beautiful, varied worlds, interacting with them. It launched to some serious disappointment Cursed or hard?

Im going to say it's just really hard. It was an ambitious promise, but not an impossible one.

Testament to this is that, a year later, the game released an expansion that made huge progress toward fulfilling player's hopes for the game.

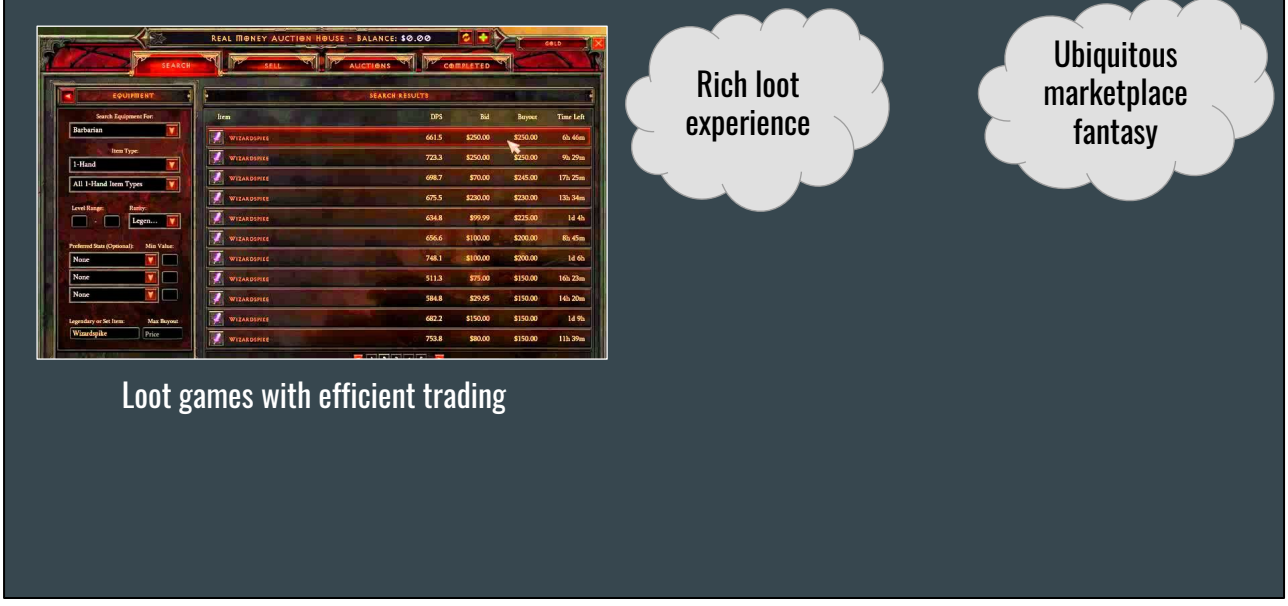

Diablo is a game all about the magic of exciting, surprising loop drops. But this fantasy naturally glows to encompass the fantasy of commerce. Players want to buy, sell, and trade the amazing things they find. So for Diablo III, they tried to explicitly make that possible for players, with a new player promise: that they could buy and sell found goods with millions of other players, through a shared auction house. It used real money, but that's not important for this story. The game launched to signficant player frustration, but was this a cursed problem, or just hard?

I'll say cursed. Arich loot experience requires a variety of loot drop experiences. It's about that randomized reward schedule, the feeling of anticipation leading to fulfillment, allowing you to fantasize about a given object then finally receive it. Conversely, the marketplace fantasy, at least the one Diablo III was pushing, involves buying and selling anything to anyone, with sheer efficiency. This means that all items become fungible for one another. Every loot drop becomes no more interesting than the amount of gold I can sell it for, since I can always buy that item on the store too.

As production director John Hight says, the auction house ""Undermines Diablo's core game play: kill monsters to get cool loot"". And indeed, they eventually removed it.

Note that WoW of course has an auction house, but key items tend to be bind-on-pickup, so the auction house's role is limited. The marketplace fantasy is severely diminsehd. And prior Diablos had trading, but it didn't have an efficient market, so finding an item in the wild still meant something.

Rewards at specific times and places

Need for personal safety and convenience

Always-on location-based games

Finally, always-on location-based games. These games promise a fantasy of a game that adds a magic layer to your life. Wherever you go, whatever you do, the game is there to offer its take on the world. On the other hand, almost all games implicitly promise to maintain a player's basic need for personal safety, convenience, even mindfulness (being actually aware of what they're doing). Pokemon Go launched to some horror stories of accidents, and people undergoing extreme inconvenience for the game. But it also just generally created a new vector by which phones can more proactively take us out of our daily lives. Is this tension fundamental, is it cursed? Or just hard to fix?

Always-on location-based games

Rewards at specific times and places

requires

Play anywhere and any time

NOPE

Need for personal safety and convenience

requires

Only play when appropriate

Life Disruption Problem

cursed

I say it's cursed. If I'm only playing when appropriate, if the game *isn't* interfering with my life, then I'm not playing the game on its own terms. And this fundamentally compromises the core fantasy of a location-based game that lies on top of your life.

Of course, there are actions you can take to mitigate this tension. But as always, they require sacrifice.

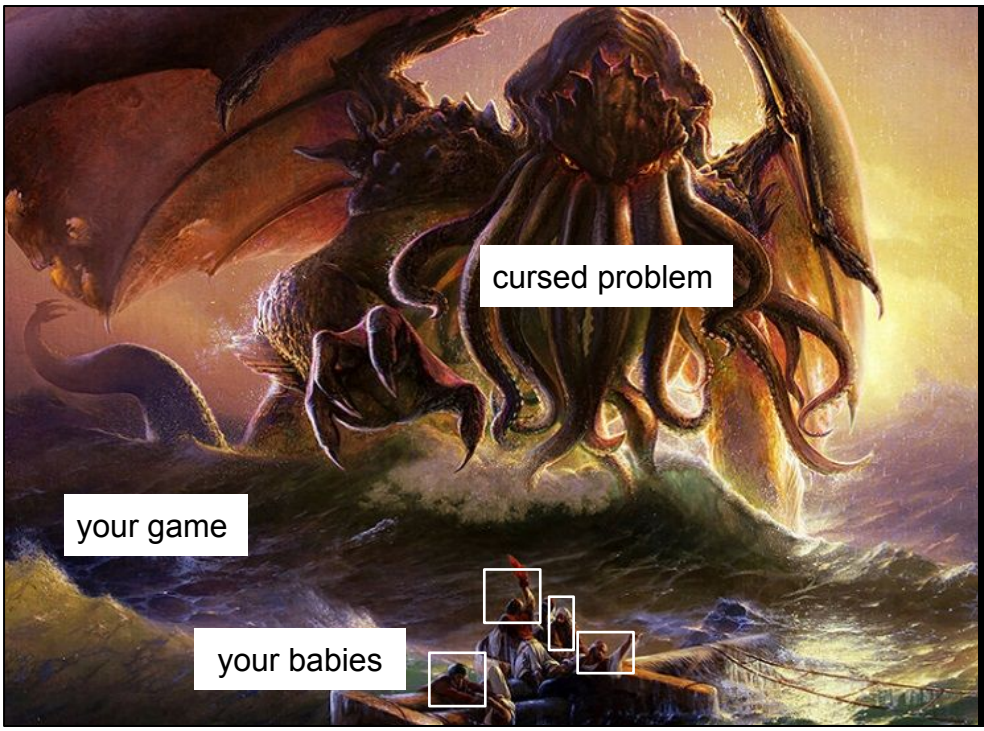

The Nature of Cursed Problems

Four Core Answers

And that's what it's all about. Like an eldritch horror, a cursed problem lives at the heart of your game, sleeping, waiting to strike. You don't get to just remove it through little tweaks. Just as you can't just pluck Cthulhu out of the ocean. He lives there; he is one with it.

Instead you placate it, by making sacrifices. This is what makes them cursed. They take sacrifices. As long as the core of the game remains the same, you'll have your problem. The question is how far you need to go in changing it.

Outline

Part 1: Cursed Problems and Where to Find Them

- Defining Cursed Problems
- An Example Problem
- Identifying Cursed Problems

Part 2: Defense Against Cursed Problems

- A Game Model
- The Four Core Techniques
- Further Applications

And that's where we are now. You've decided to face the cursed problem head-on, not just make a different game. But you can't solve it. Here's where we learn to compromise. To give up some of our dream, but hopefully keep enough to feel good about it.

We begin with a game model, on top of which we'll build the four techniques.

cursed problem

an unsolvable design problem, rooted in a conflict between core player promises.

Recall, here's the definition.

cursed problem

an unsolvable design problem, rooted in a conflict between promised experiences and objectives.

Common Special Case

For the rest of this talk, I'll focus on a special case of cursed problems, seen here.

experiences

moment-to-moment gameplay; what it is like to play.

objectives

goals, driving forces, what I want to get out of the game in the end.

Obviously these are some overloaded terms, so I'm going to tell you what I mean.

W

cursed problem

an unsolvable design problem, rooted in a conflict between promised experiences and objectives.

Common Special Case

We work with this special case because A) it's large - constitutes more than half of the cursed problems I've found. And B) it's easier to reason about. All the techniques I'll describe work in the general case, but they are a little less intuitive, harder to explain.

Free-for-all PvP games with high-skill gameplay

FFA Politics Problem

So in the case of FFA politics, we see that the first promise was an experience, and the second an objective.

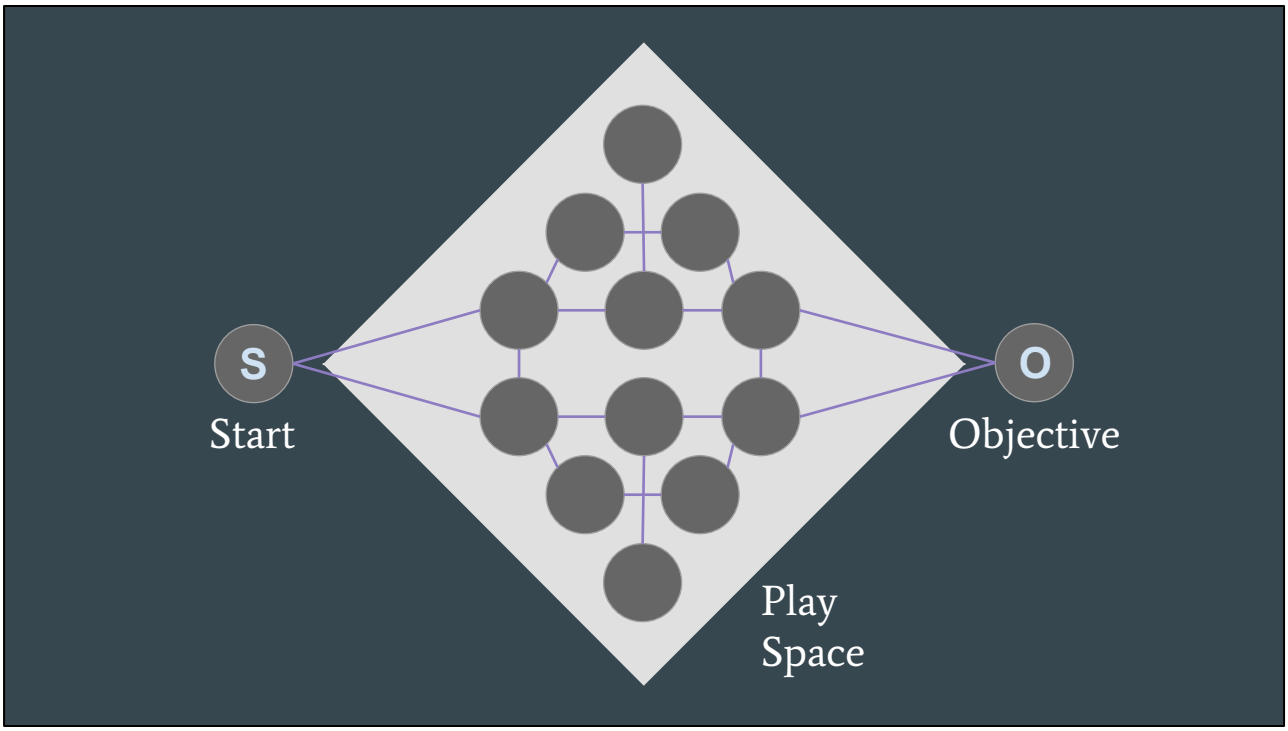

Alright then, this is your game. A play space, starting at S, in which players meander, moving toward their objective O. The objective can come from the player or the game.

It's a bit like a state machine, but the nodes here represent not just states of the game, but mental states as well. In other words, think of the nodes as experiences.

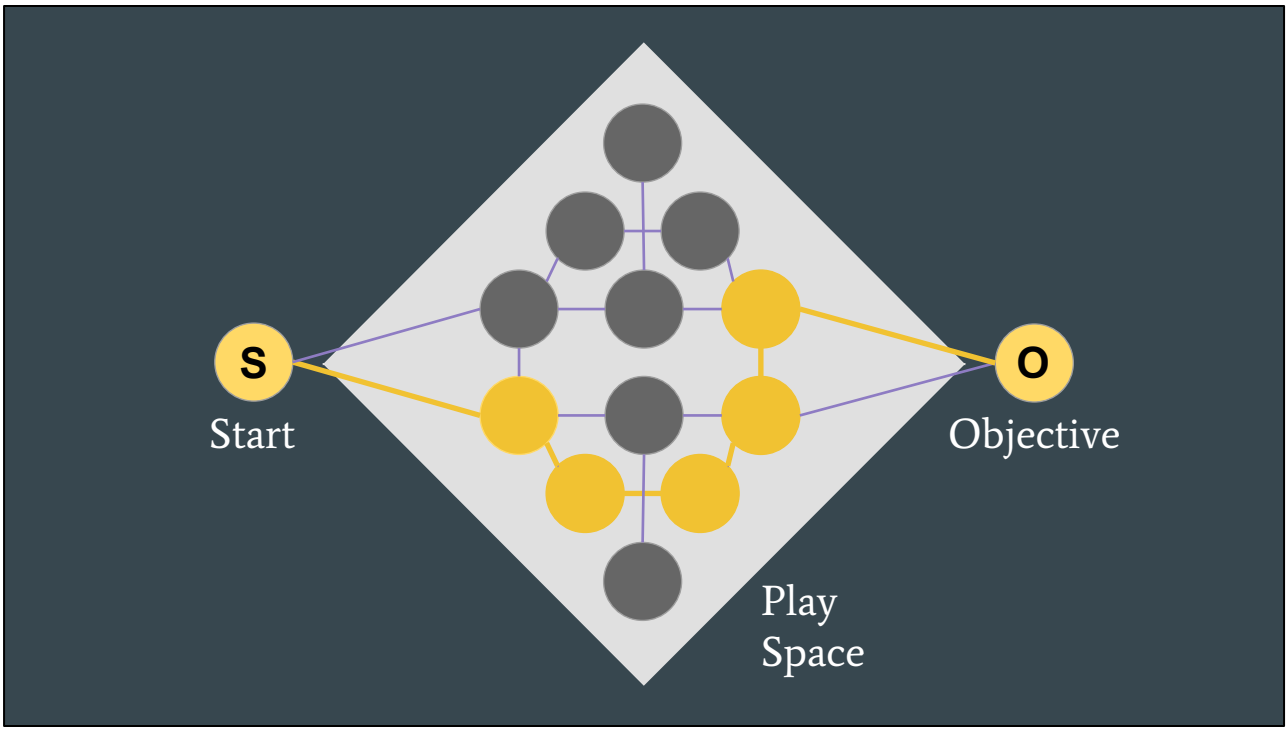

Players tend to take certain paths toward the objective, either because they are more optimal, or because something in the game draws them that direction. We'll call that the golden path - the path that players are drawn to, on account of the objective and their understanding of the game. There may be multiple golden paths, mind you.

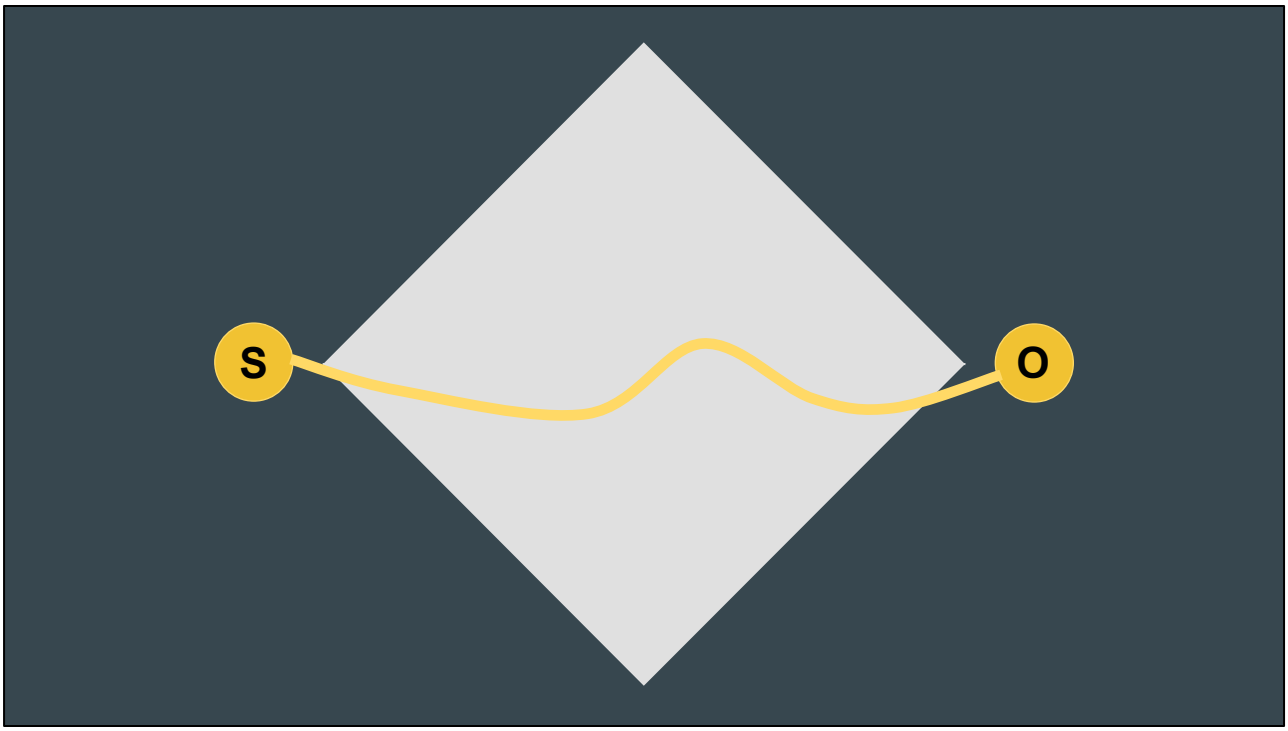

Let’s abstract this state space to make it a little easier to talk about. The diamond is the state space, the yellow line is the golden path. So. This is your game.

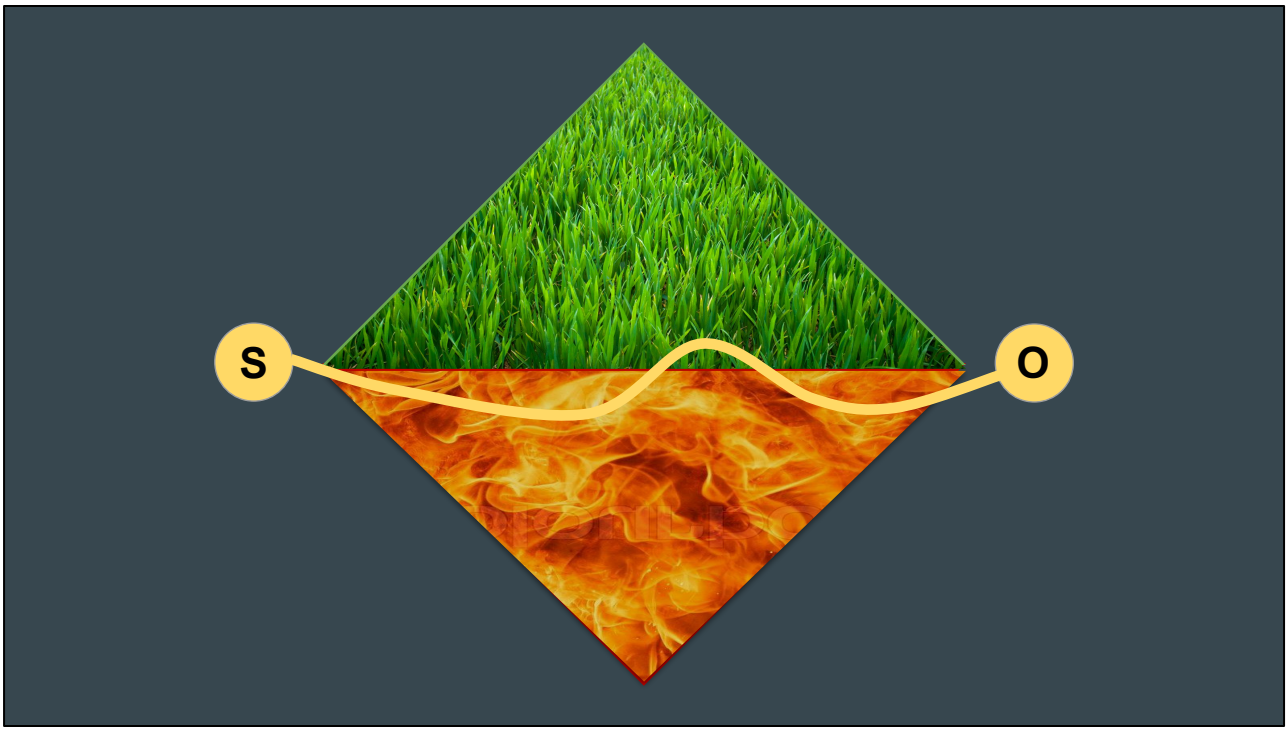

And this is your game on cursed problems. The grass are the experiences that fulfill the player promise. The fire are those that don’t. The golden path inherently takes the player straight through the fire, with occasional dips into the grass. Maybe I’m being a little dramatic. But it’s a helpful metaphor.

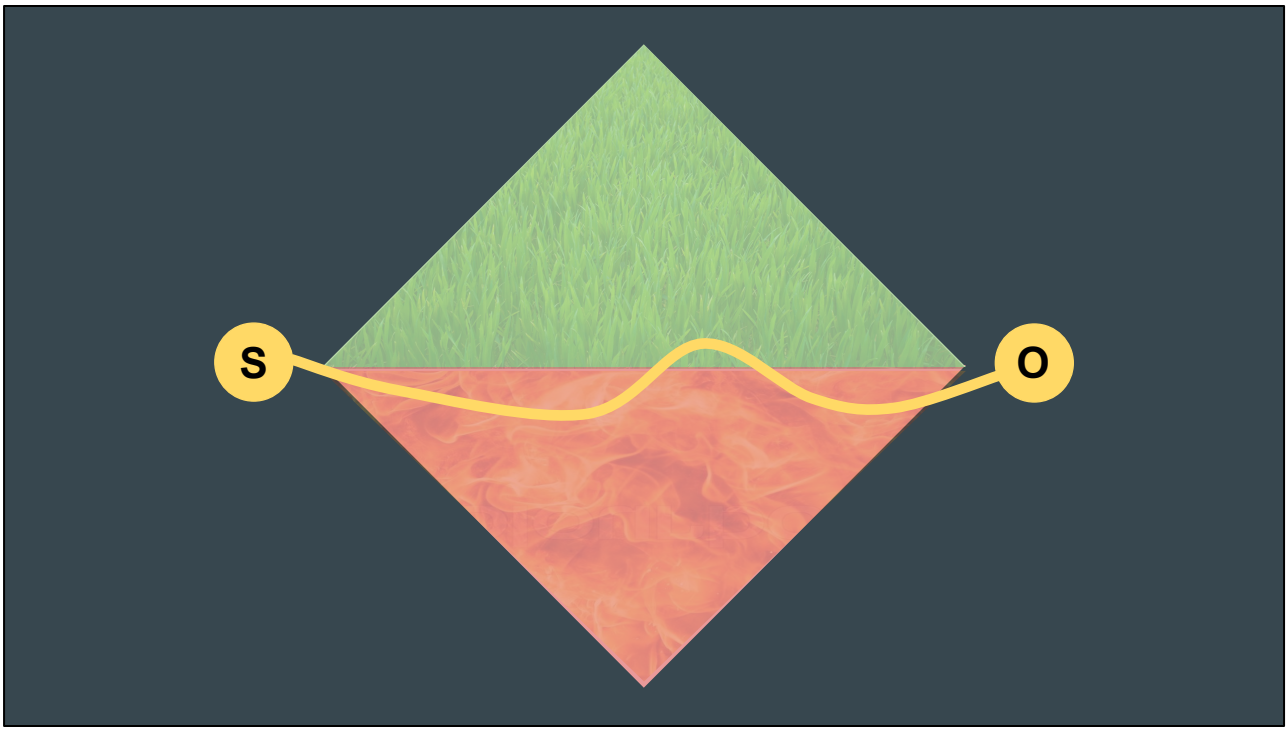

So, let’s remind ourselves of what’s going on here. The player believes they’ve been promised two things: to fulfill an objective O, and to have certain experiences, in the green. Chasing O naturally brings the player through the red zone, where they aren’t having the promised experiences. We’re letting, or even requiring, the player to act against their needs.

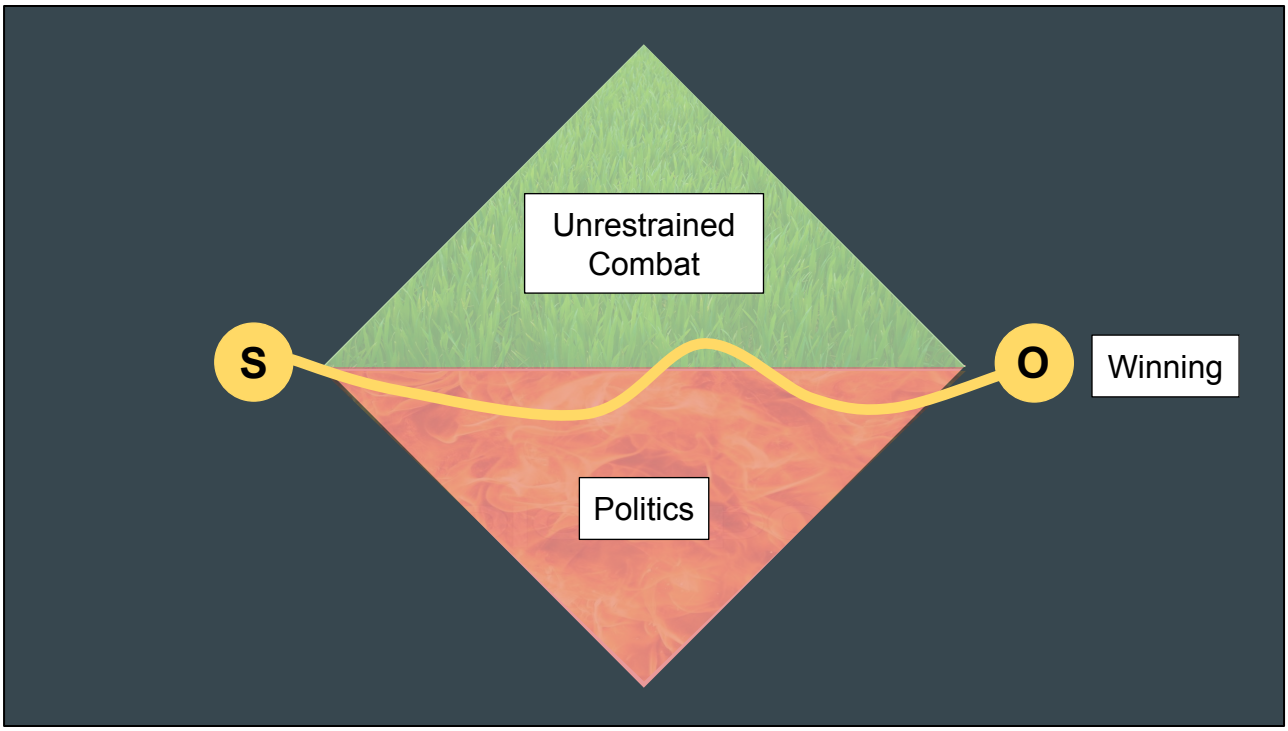

Let’s make this concrete for FFA Politics.

Answering a cursed problem means keeping players from spending time in undesirable states. For FFA Politics, this is the state in which players manipulate a specific opponent’s performance, rather than improving their own or reducing opponents’ overall. (Political games are great! It’s only a garbage fire when this fights your intended experience.)

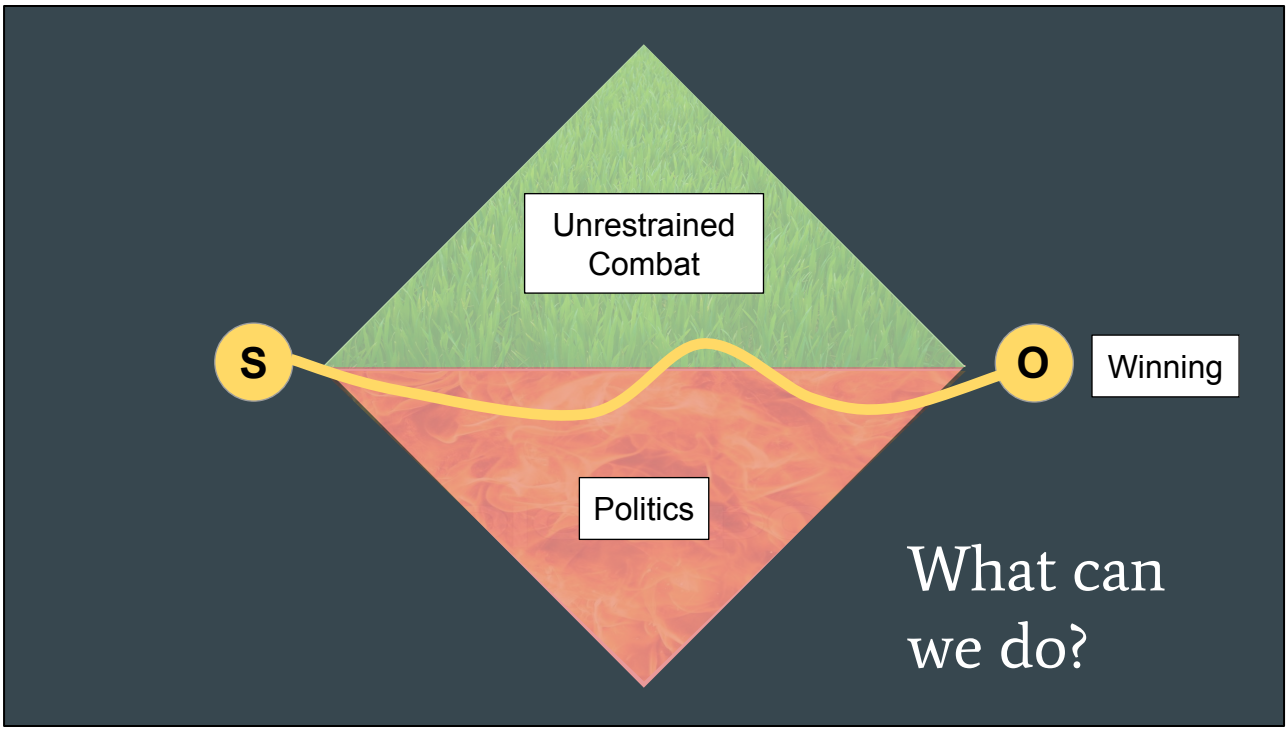

So now the question, in the abstract is this: what can we do to prevent players from taking these paths?

Outline

- Part 1: Cursed Problems and Where to Find Them
  - Defining Cursed Problems
  - An Example Problem
  - Identifying Cursed Problems
- Part 2: Defense Against Cursed Problems
  - A Game Model
  - The Four Core Techniques
  - Further Applications

Time to see the techniques.

Games have faced cursed problems, and made the needed sacrifices to address them.

They've used four techniques.

Good news.

I did a lot of breaking down the means by which cursed problems were addressed. And I was able to find four techniques at the core, that capture almost everything.

Games throughout history have used these, often unintentionally! Note that some of these techniques result in a vastly different game. It's up to you, the designer, to decide if that sacrifice is acceptable.

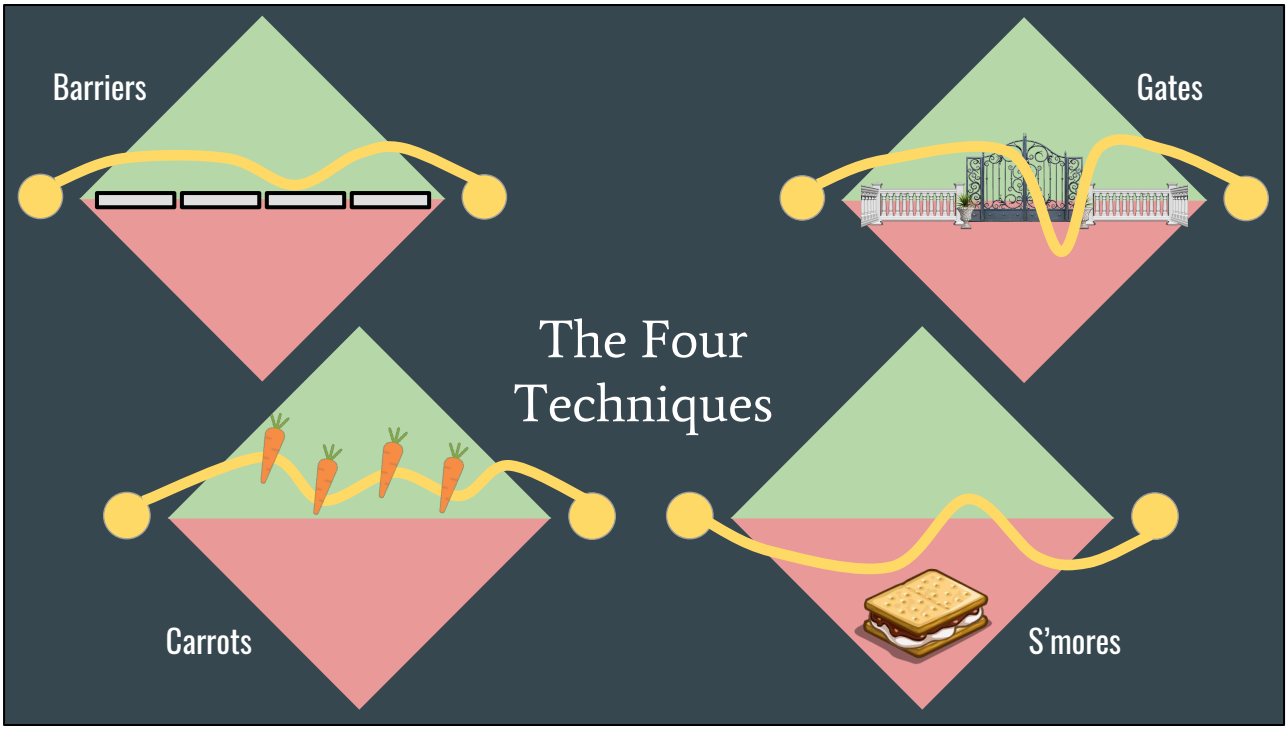

So here they are, the four techniques. We'll go through them one-by-one, using FFA Politics as an application.

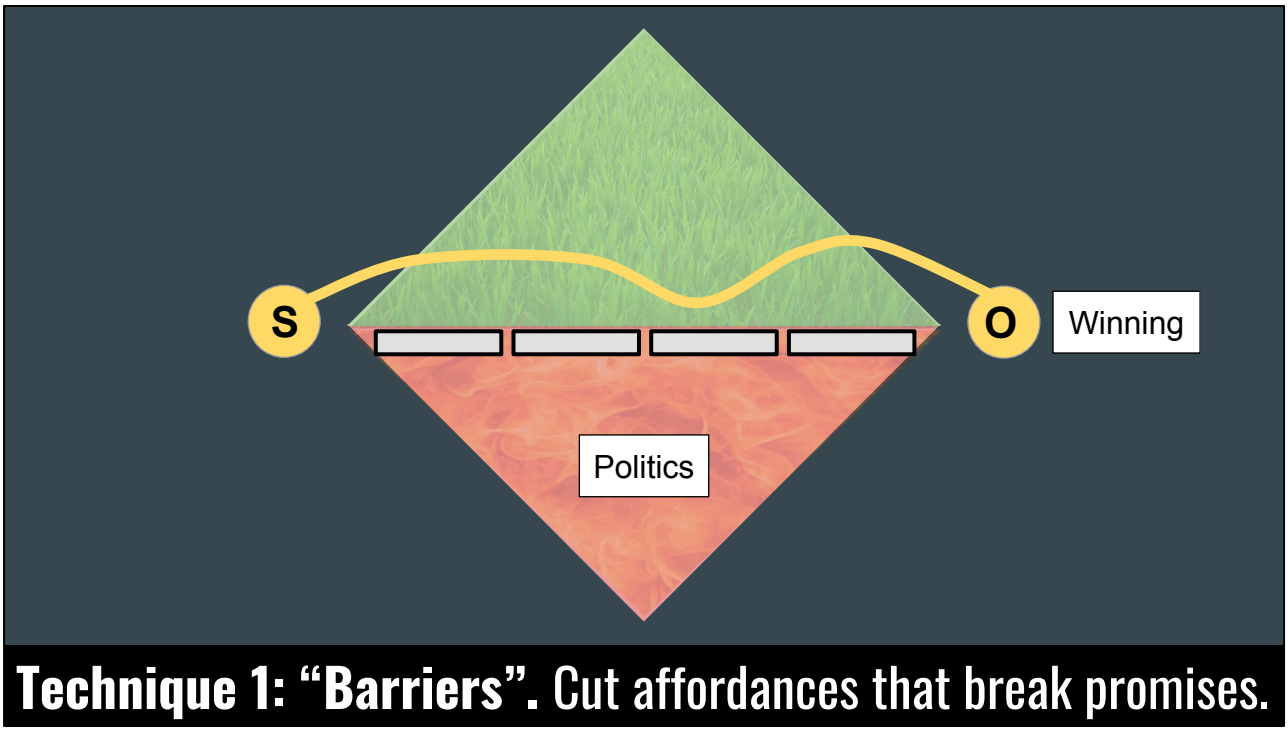

Technique 1 is to remove affordances that let players break promises. Literally make those actions impossible. So for example, how can we make politics impossible?

Free-For-All Politics - Technique 1: Barriers

Approach: Limit players' agency over one another.

Sacrifice: Some of the PvP fantasy of control.

One approach is to limit players' agency over one another.

In the extreme are foot races, in which there is almost no explicit mechanism to affect opponent performance. For a softer touch, we can look at Battle Royale games, in which player interaction is flexible, but so erratic that it's difficult to plan around. It would be essentially impossible to pick out a specific opponent, forge alliances, etc. The map is too large, there are too many players, lethality is high. In fact, Battle Royale is a great example of real progress on a cursed problem, through a new application of one of these techniques.

Note that we said every answer requires a sacrifice. And in this case we give up some of the core of the PvP fantasy - being able to control and dominate opponents.

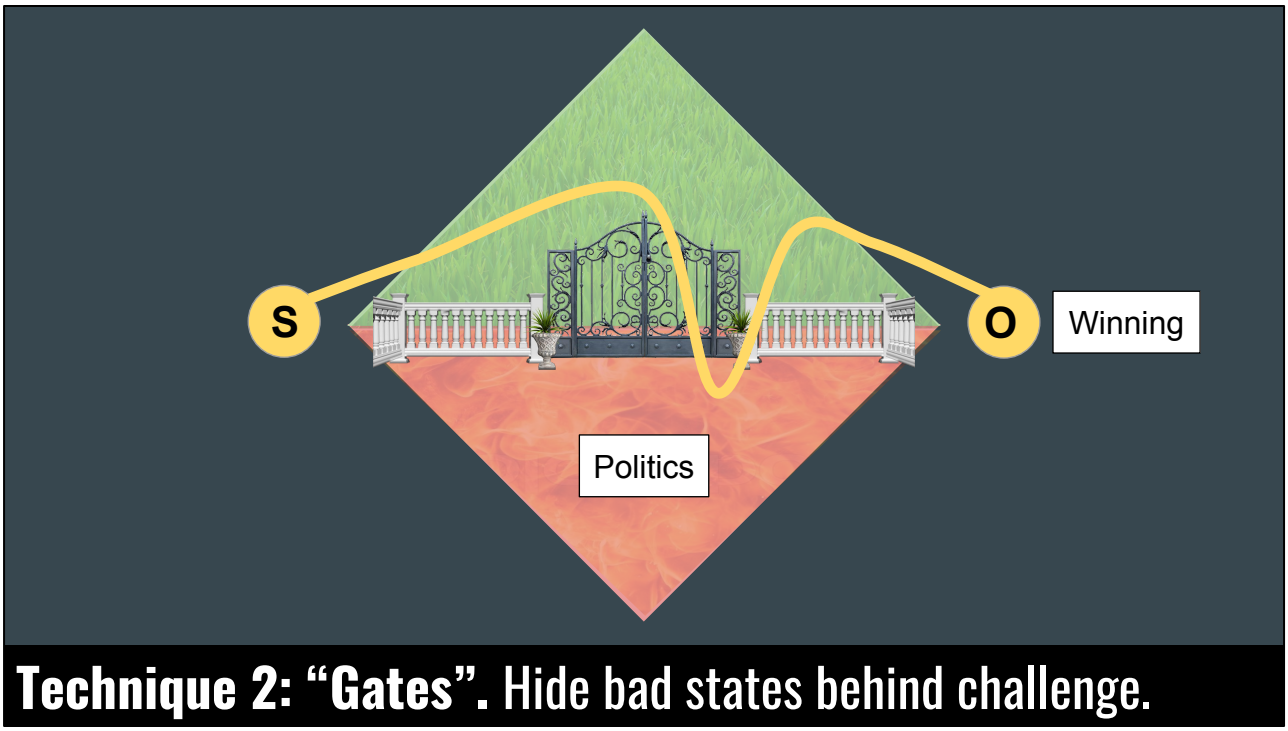

Technique 2: “Gates”. Hide bad states behind challenge.

This is a softer touch version of barriers. Rather than completely disabling the actions that lead to unfulfilled promises, instead simply make them challenging to perform or understand. This is less heavy-handed, and we can do this; games are all about voluntary challenge. Players might occasionally step over into the undesired territory, but not for long and not often.

Free-For-All Politics - Technique 2: Gates

Approach: Limit visibility of players' success.

Sacrifice: Some of the tension of buzzer beaters.

For instance, on FFA Politics, we can limit players' visibility of each other's success.

On Playstation All-Stars Battle Royale, we hid players' scores entirely. We also made them more difficult to track: a player's score was 2*kills - deaths. So in the heat of the battle, a player could plausibly believe that their opponents don't know how well they're doing, and this would give the player permission to play hard and focus on core combat.

The sacrifice here is the tension that can come from buzzer beaters - game endings can be somewhat anti-climactic. "Oh, that's what happened."

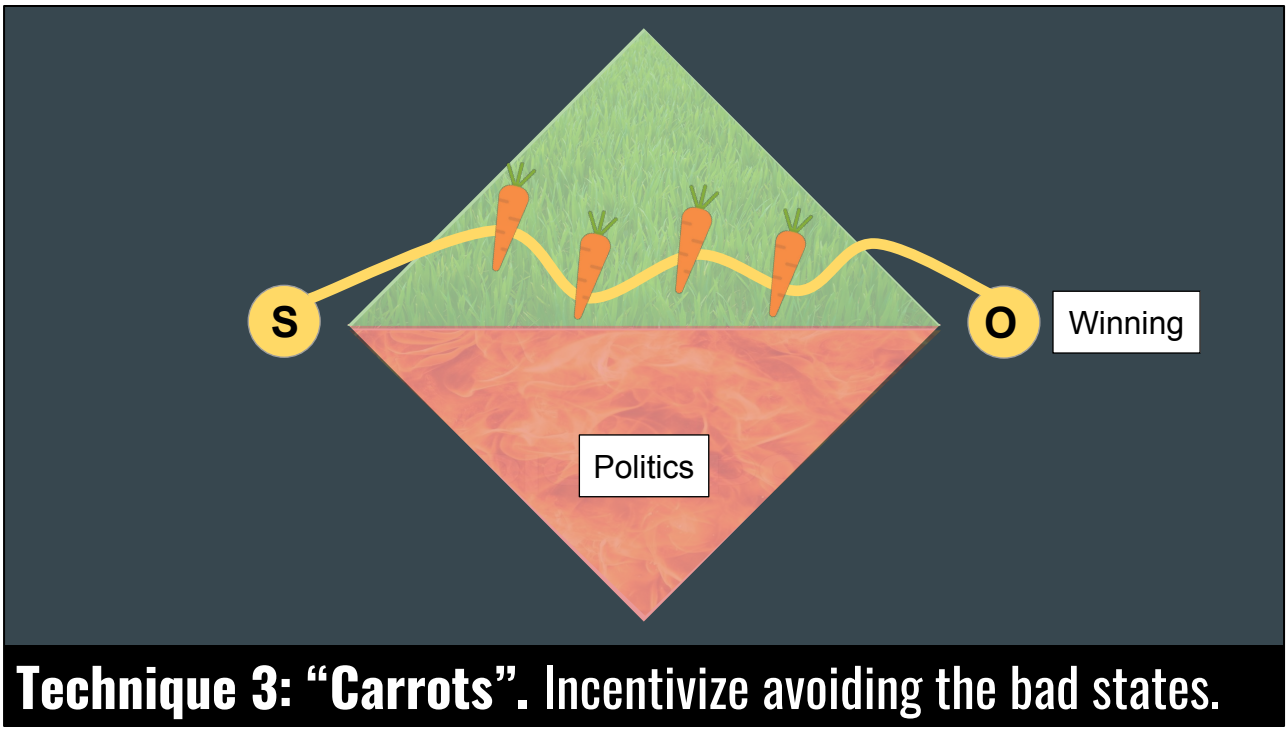

Technique 3: “Carrots”. Incentivize avoiding the bad states.

Games define or suggest user objectives, rather than being put to task for an existing user goal. This gives us freedom to simply shift the objective, or add new incentives. The idea is that players will then be naturally drawn through the “promised” land.

Free-For-All Politics - Technique 3: Carrots

Approach: Add meta-game effects.

Sacrifice: The magic circle of an individual game.

One way to disincentivize politics is to add meta-game effects.

For instance, Settlers of Catan tournaments are run in an interesting way. Players' progress through the tournament is not determined solely by wins/losses. Instead, they receive points for each game, proportional to their placement. 4 points for a win, 3 points for second place, etc. This can drastically change the incentives of a game. My time is better spent raising my own performance, rather than bringing an opponent down. This is a natural dampener on politics.

The sacrifice here is the magic circle of an individual game. A game in this context feels less like a self-contained entity, with clear winners and losers. Instead, it can start to feel like a meta-turn of the broader game that is the tournament.

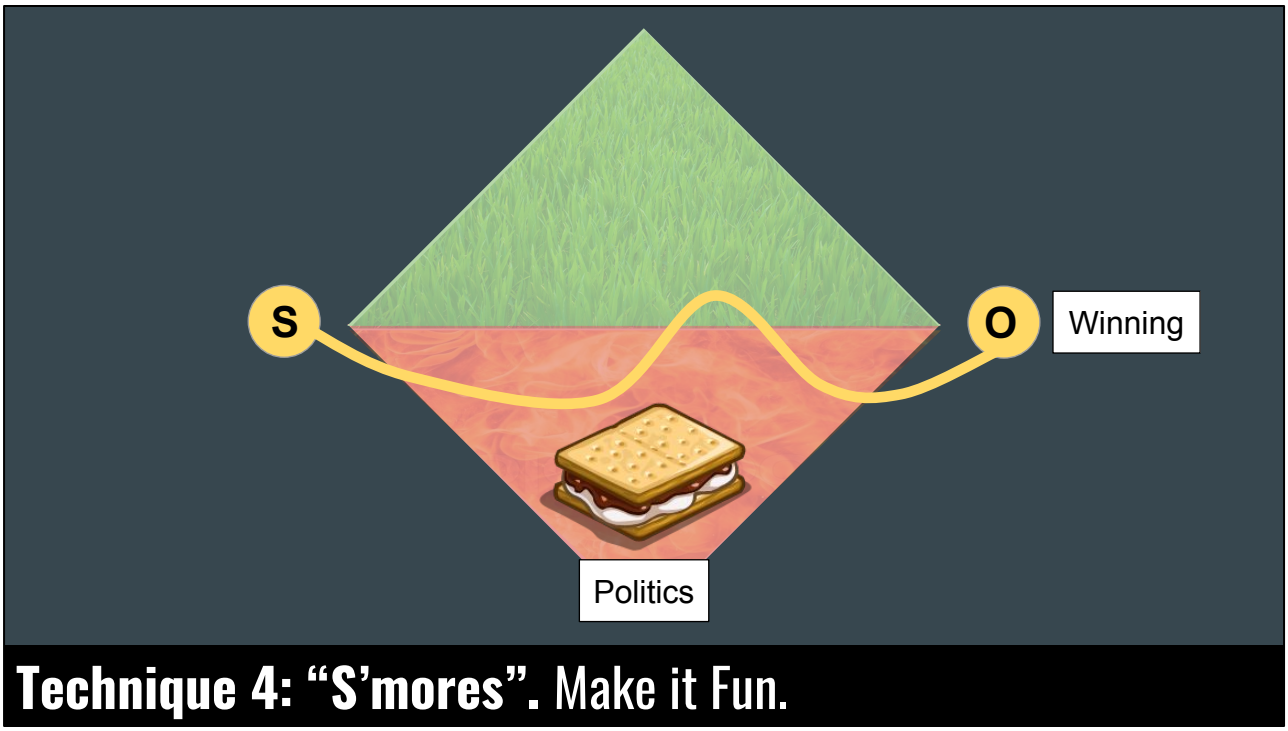

Technique 4: “S’mores”. Make it Fun.

The final technique is s’mores. If we really can’t keep players out of the fire, we can lean into it instead. Make it fun for them. At the risk of abusing this metaphor… give them marshmallows. Let them walk straight toward the fire and show them how to enjoy themselves there. By leaning into a design problem, it can be repurposed as enjoyable gameplay. Sometimes it can become an entire game. In that sense, this move can be the most aggressive. It can mean inventing a whole new game around your game, or pivoting to a different game altogether, that hopefully still captures the essence of what you care about.

Free-For-All Politics - Technique 4: S'mores

Approach: Give players tools for deep political play, e.g. secrecy.

Sacrifice: Emphasis on moment-to-moment action.

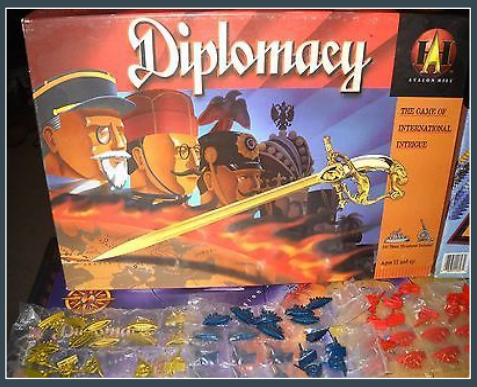

Can we lean into politics, building a game around it? Well sure, particularly if we give the player tools to make that gameplay deeper - e.g. secrecy.

This is what Diplomacy is. It says "political machinations can be fun, for the right crowd. So let's build the game around it."

The sacrifice is the emphasis on moment-to-moment action. The individual piece movements are de-emphasized over a more neutral game like Risk.

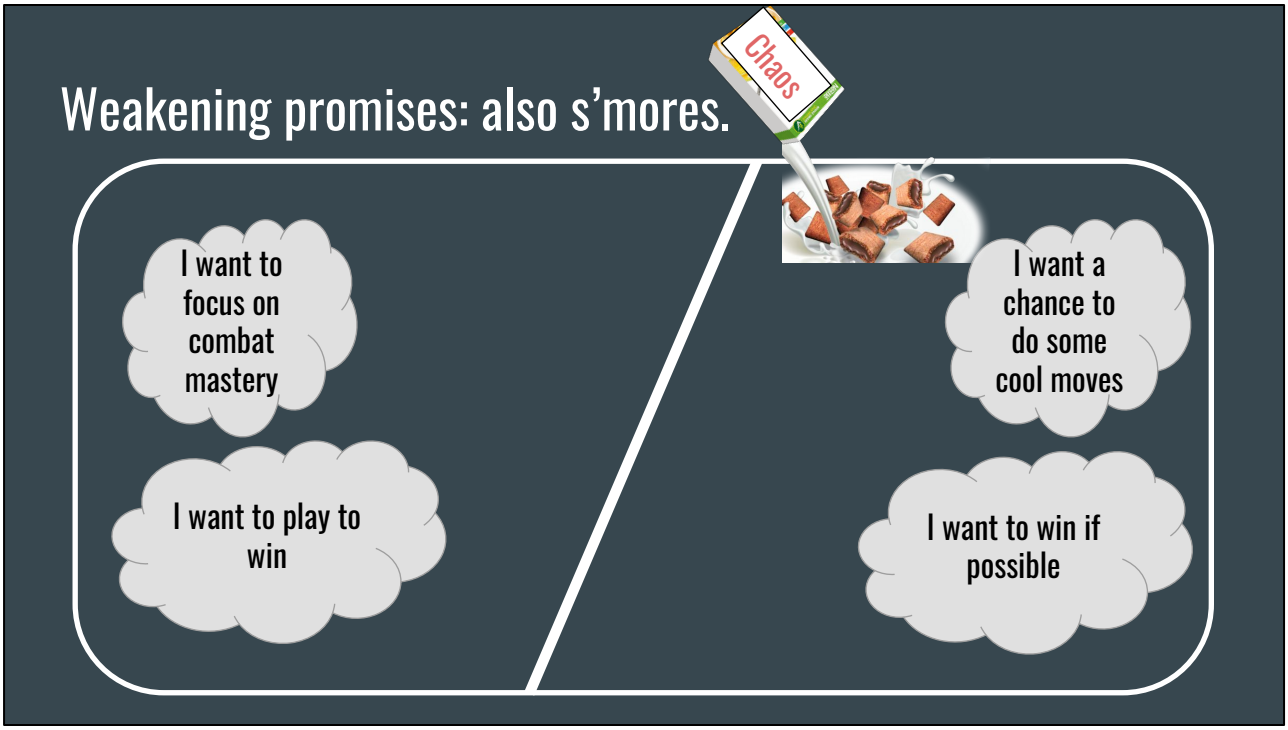

Note that we’ve already seen another example of s’mores: the weakening of promises that led us to Smash. We said: hey the promises of this game aren’t fulfillable, so let’s make that work for us.

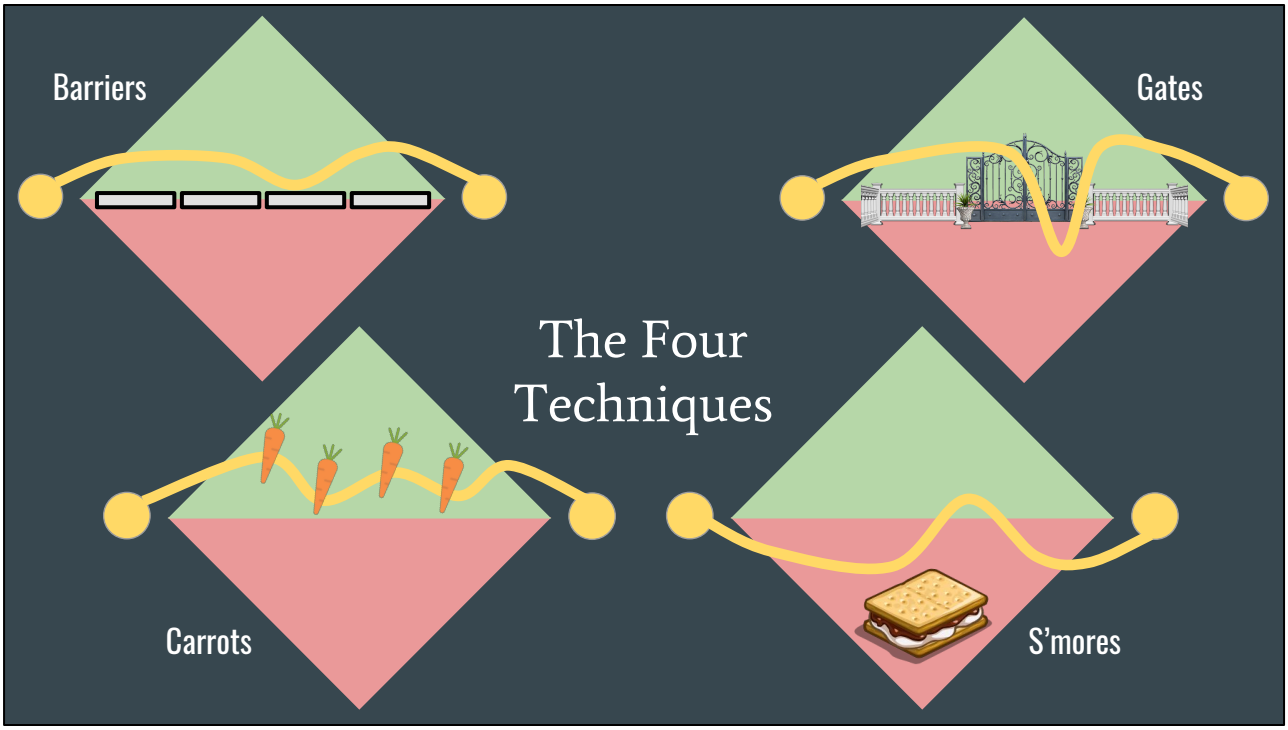

So that's it! The four techniques. Note that these are general design techniques that can work for any problem. The reason they're relevant here is because these are arguably the *only* techniques that we can use for cursed problems. The key is that we are assuming the problem is fundamentally baked into the game. We can't just twist the rules until it ceases to be there. We work around it.

This is a tiny design framework.

It won't tell you how to do the work.

But it can give structure to a difficult creative process.

Part 1: Cursed Problems and Where to Find Them

- Defining Cursed Problems
- An Example Problem
- Identifying Cursed Problems

Part 2: Defense Against Cursed Problems

- A Game Model
- The Four Core Techniques
- Further Applications

Outline

Now let's see one more example of each of the techniques, so you can see how they generalize. We'll apply two each to a couple more cursed problems.

	Barriers	Gates	Carrots	S’mores
FFA Politics				

Here’s what we’ve seen so far.

	Barriers	Gates	Carrots	S’mores
Coop Abuse	●	●		

And here’s what we’ll see next.

This is a sadly prevalent problem in any high-stakes co-op game. Players are frequently horrible to one another. And it's not obvious how this is a conflict between promises, so let's see them. On the one hand, players want to win. The game signaled for that is what it is about. The other hand players join a cooperative game, at least in part, due to a desire for social belonging.

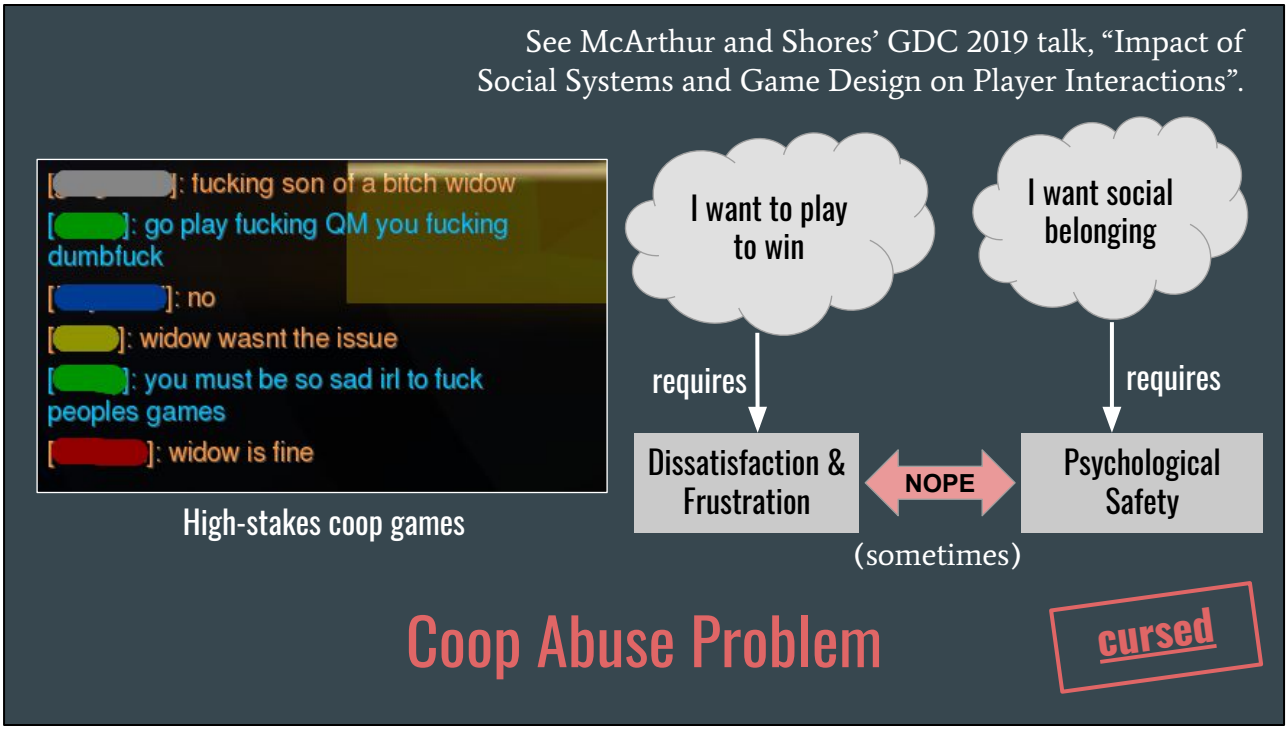

The problem is that the desire to win is in some sense fundamentally unsatisfiable. Players are said to believe a game is fair if they win 70% of the time. In a PvP game this is obviously not achievable. Dissatisfaction and frustration are common. If a player is also lacking anger management skills, and tends to blame others over themselves, they can tend toward verbal abuse. This abuse is in natural conflict with the psychological safety necessary to achieve a sense of social belonging.

I am by no means saying that players are not responsible for their actions. We need to be willing to tell the most abusive players, particularly those who harm marginalized people, that they don't have a place in our games. But the work shouldn't exist solely on the community management side. We can do a lot with the games themselves to reduce the tendency for abuse. Naomi McArthur and Kenny Shores had (independently of this work) a talk on a similar topic this year.

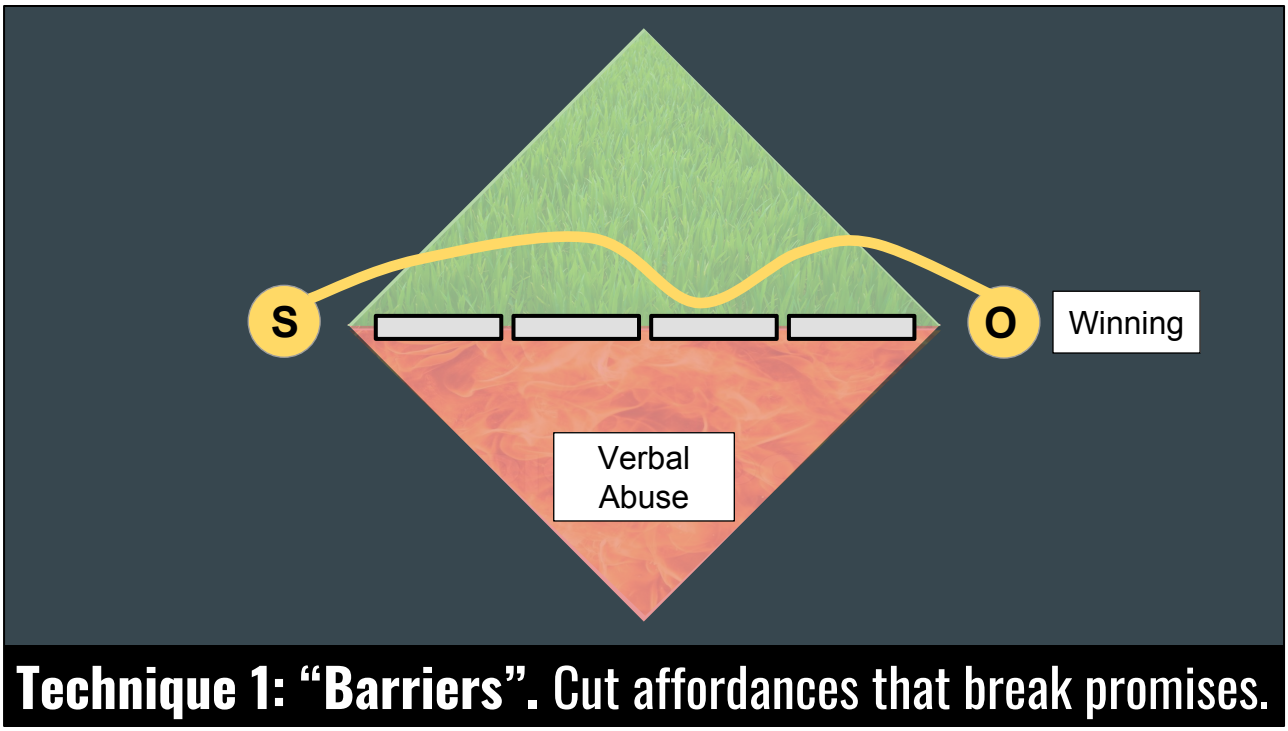

Technique 1: “Barriers”. Cut affordances that break promises.

So, how can we remove up affordances for verbal abuse?

Think about how these two players move through the gamespace together.

We literally remove the player actions that let them move into undesirable territory, so the golden path is forced to move.

Keep in mind not *any* change of affordances will work! You have to change options that allow players to engage with the problematic experience. And this means changing some minor-to-major assumptions about what the core of the game is. While hopefully keeping the absolute core.

Coop Abuse - Technique 1: Barriers

Approach: Limit player communication.

Sacrifice: The rich social channel for relationship-building.

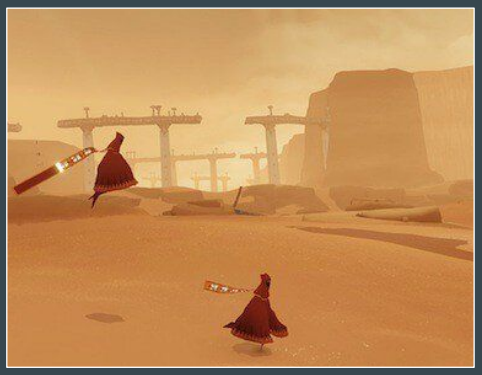

For instance, we can simply limit communication.

Journey is a game almost about this idea. You can walk, jump, and make pretty pings at each other. That's it. And still players have beautiful experiences together. Of course, journey is not a high-stakes game. But we can see promises of this kind of option in Apex Legends. Its ping system is so robust that you could build an entire game around it, disabling linguistic communication entirely (Although Apex Legends does not go this far).

The sacrifice here is the communication channel that allows relationships to build. Disclosure ends before it even begins,, if all you can do is communicate about the game itself.

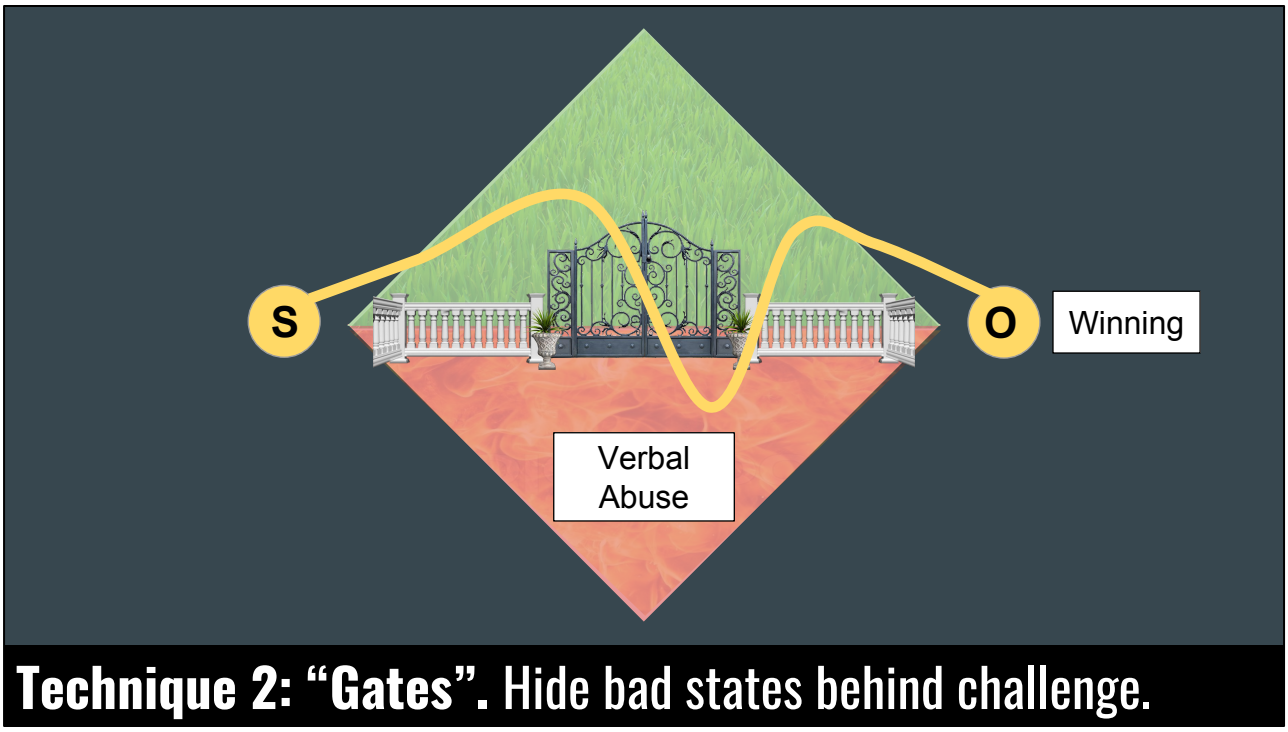

Technique 2: “Gates”. Hide bad states behind challenge.

How can we make verbal abuse challenging, rather than altogether impossible?

Coop Abuse - Technique 2: Gates

Approach: Limit individual responsibilities.

Sacrifice: Some of the fantasy of harmonious cooperation.

We can limit players' individual responsibilities.

Heroes of the Storm does this to the typical MOBA model. For example, gold is shared across players, diffusing responsibility, and even making it more difficult to discern how well each individual player is performing. If I don't know who's underperforming, it's harder to flame on individual players.

A more extreme example comes from games like Werewolf and Mafia - really any social deduction game in which the good players don't know they are on the same side. If I don't even know who is on my team, how am I to abuse them for failing to help me? Of course, once the game ends, that's an entirely different matter.

	Barriers	Gates	Carrots	S’mores
Coop Abuse				

So we’ve seen two more examples.

	Barriers	Gates	Carrots	S'mores
Quantified Creativity			●	●

Now let's see the other two techniques, on our final problem.

Creative games with goals

Quantified Creativity Problem

This might be a familiar story: when the Sims came out, I was enthralled. I played with objects, places, but of course more than anything, the characters. I manipulated them into stories then discovered surprising outcomes. That's the dream of the Sims: player storytelling.

At least, it started that way. But it wasn't long before the game gave me a few goals. Something comforting to latch on to, a sense of accomplishment. And for some players, like me, this consumes the experience. All my effort was quickly directed to the rat race, just trying to secure a good life for my family, a fancy house, and a sense of purpose for my poor protagonist. The game kept me intensely involved, but it became a game, and creativity/storytelling/investment fell by the wayside. What's worse, when I finally reached that coveted astronaut dream, I had nowhere else to go. I puttered around for a while, until I just stopped playing. I imagine it was similar to the post-entry depression experienced by real astronauts.

I'm not saying all players will fall for this trap. Clearly they won't. But to the extent that the game *succeeds* in giving players extrinsic motivation, it typically *fails* to preserve what makes the game great in the first place.

Creative games with goals

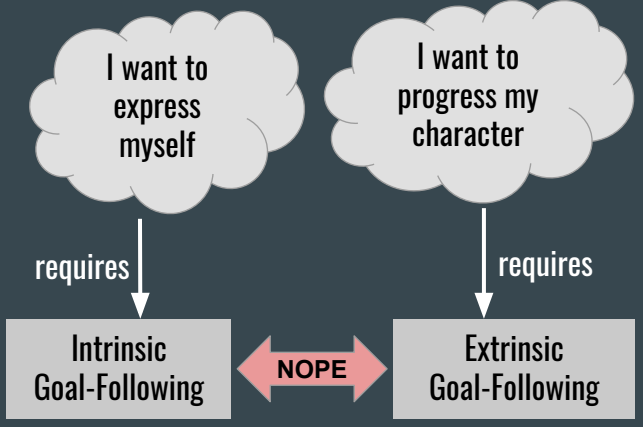

Quantified Creativity Problem

cursed

My desire to express myself requires me to focus on my intrinsic goals. But my desire to progress requires that I follow extrinsic goals - which have a psychological` tendency to subsume interest in intrinsic intrinsic goals. Note that we see a similar phenomenon in exploration games.

So, you can eschew the extrinsic goals, but then you risk giving up a huge number of players and circumstances of play. Can you have both? Not unless the designer gives something up. This is where our sacrifices will come in.

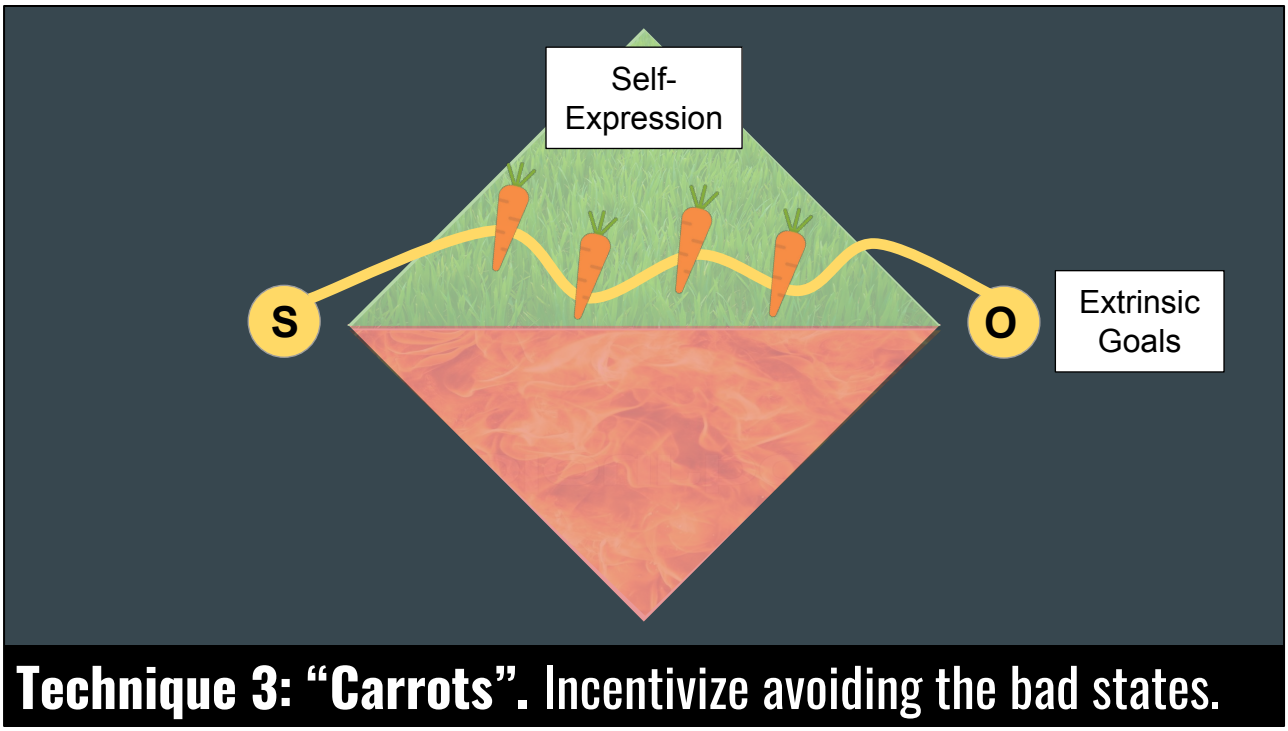

Path 3: Carrots. Incentivize avoiding the problematic states.. Games define or suggest user objectives, rather than being put to task for an existing user goal. We can exploit this by giving players the tools to opt out of the undesirable experience, *and* encouraging through incentives, shiny distractions, or literal encouragement. Note we can’t expect players to opt out solely because it’s unpleasant. It may be a step toward their goal, and players will rarely optimize for their own enjoyment.

Quantified Creativity - Technique 3: Carrots

Approach: Give players concrete reasons to be creative.

Sacrifice: The safety of unconstrained creativity.

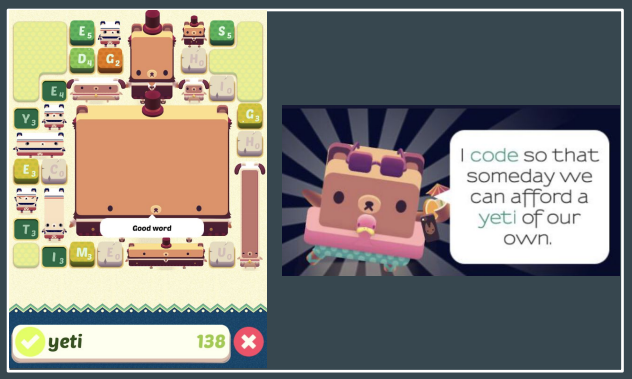

If a player's tendency is to optimize their behavior, can we incentivize them to opt out of doing so? Yes, by making predictability dangerous. This is common in PvP games: location in an FPS, build order in an RTS, attacks in a fighting game. Creativity and exploration gain explicit strength. Clash of Clans looks not entirely unlike the Sims, but you gain explicit advantages for being interesting. If you build a creative, off-meta base layout, you're more likely to resists the attacks of players. Yes it's a different genre; it's an extreme move. But depending on what you're looking for, it might capture the core.

We did this on Alphabear, by layering a creative sort-of-game on top of the goal-oriented game, and requiring you to beat the latter to enable the former. It's fun to spell interesting words, but it's often opposed to the primary goal of high scores. We created a kind of "extra-game" that rewards these behaviors. After each round, the game presents mad libs using spelled words. These could be shared with other players, not for formal value, but for social validation. Players reported this meaningfully changing their play experience. They weren't just spelling to win, they were choosing words playfully, imagining the mad libs that would result. We've taken an intrinsically creative game, dampened it with an optimization-focused layer, then built another intrinsically creative game around *that*.

The sacrifice here is the safety of unconstrained creativity. In neither case is the player simply expressing themselves. Instead they are expressing themselves pointed toward a problem. That's a very different mindset, and certainly less

expressive.

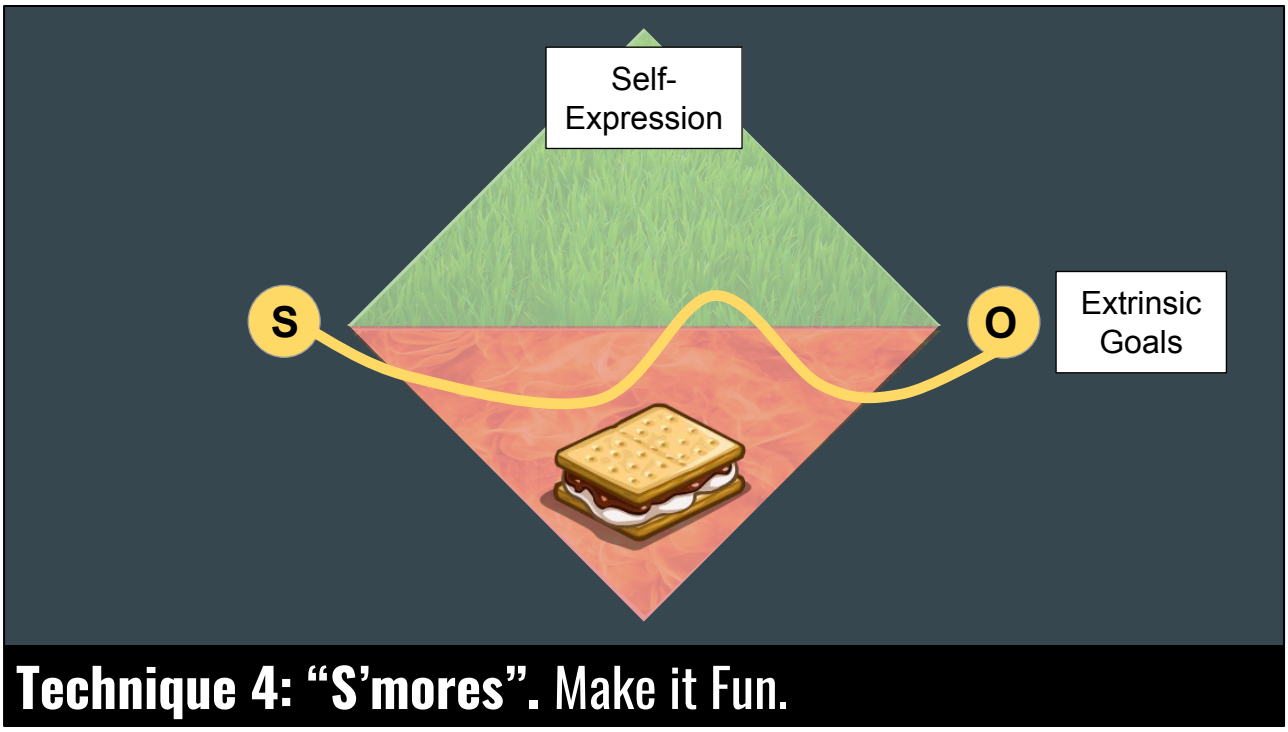

Technique 4: “S’mores”. Make it Fun.

For our final example, can we take a creativity game and simply lean into a player’s tendency to not be expressive?

Quantified Creativity - Technique 4: S'Mores

Approach: Create rich, deep goals around the simulation.

Sacrifice: Most of the weirdo self-expression.

Sure, we can take the rules and constraints and systems of the simulation and lean into them, making them so rich that playing them to win or for progress *is* satisfying enough.

And guess what? This is what a 4x game *is*. You take the deep simulation of e.g. SimCity, and you create so many incentives and goals and constraints that it transmutes into something else. It becomes a strategy game - a game about winning. Creative expression takes a backseat, but you end up with something beautiful, borne out of the same rich nouns and verbs that make creative games so fun.

The sacrifice here is even more of the weirdo self-expression that we lost in the previous example. You can play Civ as a purely creative game, but it strains against it. You've ended up making a totally different kind of game than the one you intended. But sometimes that's the better game. Particularly if your original game was impossible.

	Barriers	Gates	Carrots	S’mores
Quantified Creativity			✓	✓

So there we go.

	Barriers	Gates	Carrots	S'mores
FFA Politics	✓	✓	✓	✓
Coop Abuse	✓	✓		
Quantified Creativity			✓	✓

These are all the examples we've seen. Hopefully that gives you some intuition for how to use these techniques on your own.

Wrapping Up

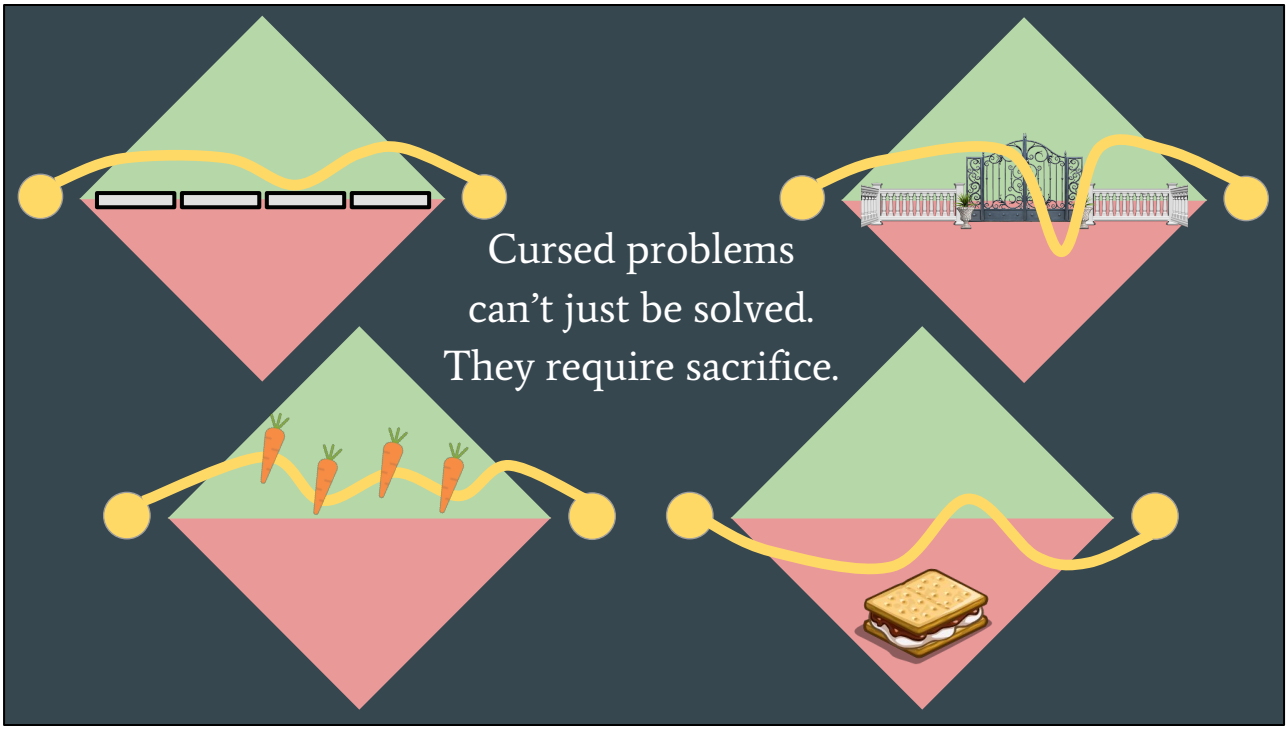

The four techniques I've described are surprisingly prevalent - or perhaps not surprising, given how general they are. Wall off the problem, add challenge, change incentives, or make it fun. Each involves some kind of sacrifice.

Unanswered cursed problems are the highest-hanging fruit.

They’re also the ripest.

One more note: we’ve talked about what to do when you run into a cursed problem. But sometimes it’s worth it to chase them. Because cursed problems are so difficult, they are often underexplored.

Challenges

Balanced player-generated content in multiplayer games.

PvP games in which all players feel there was a just outcome.

Mystery and discovery in the age of the internet.

So here are a few challenges. Problems people have told me are impossible. And they're right! But with a twist and a compromise, anything is possible. And behind that door lies a potentially huge creative space.

- Multiplayer games with player-generated content that isn't totally unbalanced.
- PvP games that get around the fact that most players expect to win 70% of the time.
- Games with a sense of mystery, surprise, discovery, and individual exploration, that maintain that feeling despite the Internet's tendency to lay all secrets bare.

Give these problems a try, see where you end up.

So hopefully you’re a little more prepared now. Whether you tackle one of those challenges head on, or find yourself with a cursed problem against your wishes, you’re in a position to not get lost. That desert really is empty.

The trick is to give up. To release the curse by abandoning the desert. By look over your shoulder, and seeing the lush forest of possibility just around the corner.

Thanks!

Questions?

Alex Jaffe
@blinkity

Extra Slides

Cursed Problems I Discussed

- Free-for-All Politics
- Quarterbacking
- Skill Inflation
- Commodified Reward
- Life Disruption
- Coop Abuse
- Quantified Creativity
- Player Content Problem
- PvP Win Rate Problem
- Explorer Problem

Cursed Problems I Didn't Discuss

- **Decision impact:** I want to make an important decision now / I still want to make an important decision later
- **Performance variance in high-score games**: I want to get a high score now / I still want to get a high score later (but there's a natural cap).
- **Loot hedonic treadmill**: I want the most awesome loot / I still want the most awesome loot, later.
- **Social Simulation that works w/ limited concurrency**: I want to play when is convenient to me / I want to play with a large community.
- **Evergreen games that don't alienate new players**: I want a simple game / I want a complex game.
- **Player Fragmentation Problem:** I want to play in my mode of choice / I want to play with a large pool of players.
- **Pay to Win Problem:** I want to pay to play the game my way / you want the game to feel fair and meaningful.

I'll also leave you with a challenge. We're lucky to work in such a rich medium, where unpassable obstacles an opportunity to rethink assumptions, rather than a true blocker. Stop assuming that something isn't possible. With a twist and a compromise, anything is possible. Here are some of my favorite hard problems that I haven't given up on. Try the above framings. See where they take you!

- Interactive storytelling that obeys traditional narrative arcs.
- Balancing player-generated content in a multiplayer progression system.
- Real-world exploration games that don't incentivize upending your life.
- Digital social deduction games that don't rely on the high-bandwidth interactions of IRL.

Free-For-All Politics - Technique 2: Gates

Approach: Limit players' awareness of each others' success.

Sacrifice: Some of the potential tension of buzzer beaters.

Another example of limiting awareness of success is Ticket to Ride. Softer touch limitation than Playstation All-Stars'. Players can generally see the score, but there are secret objective cards that modify the score.

Coop Abuse - Technique 4: S'Mores

Can we make toxicity fun for all parties? Yes, by bringing it into the magic circle.

Endorsed Griefing

Endorsed Shit-Talking

Can we make toxicity fun for all parties? Often we can, by bringing it into the magic circle. Griefing, in particular, spans a wide range of enjoyability, from seemingly safe games in which glitches allow players to abuse each other, to opt-in PvP in which the threat of in-game (combat) harassment is part of the appeal. The idea of a dog-eat-dog world in which players are potentially at the mercy of others is, for better or worse, a genuine player fantasy, and many games lean into this. PvP zones in an MMO and invading in Dark Souls both are good examples of this pattern. They allow the toxic behavior but bring it to the surface, in the process giving up some degree of true co-op nature.

In the case of verbal toxicity, a novel pattern is to limit players' communication, but to give them explicit communication tools that might otherwise be toxic. E.g. shit-talking emotes, as in Arena of Valor. The game's endorsement of these emotes defangs them, allowing players to trade in the fantasy of shit-talking, without as much risk of trauma. It's like mom giving you a list of curse words you get to use with your friends. Maybe fun at first, but it kind of loses the spicyness.

Quantified Creativity - Technique 2: Gates

Can we provide *challenging* goals that don't discourage creativity?
Yes, by using intentionally opaque systems.

Animal Crossing: New Leaf

Devil May Cry

If incentives distract players from their intrinsic interests, can we obfuscate those incentives to make explicitly chasing them impossible? Animal Crossing does just this with its "Perfect Town" score. The explicit metric is never stated, though it is hinted at. For some players, this score acts as a motivator, but mostly to do what they already wanted to do: make a more beautiful town. If explicit rules were given, many players would simply optimize for those goals and cease to be expressive.

The Devil May Cry series rewards stylish combos similarly. You *could* learn to optimize for it, but most players simply perform stylish, creative combos, as they see them, trusting that the game will rewards them for it.

Gif from https://comicvine.gamespot.com/images/1300-4836440

Decision Impact Problem

"A game is a series of interesting choices" - Sid Meier?
Special case: can a game have many "high-impact" moments, which can substantially change win chance?

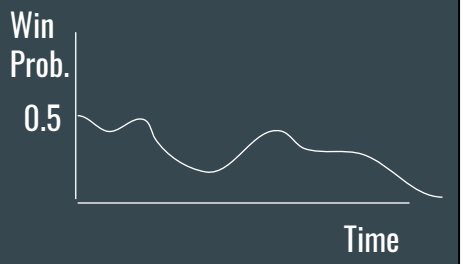

Flat Scoring

Rubberbanding

Snowballing

This is the broadest of today's cursed problems. They say Sid Meier said "a game is a series of interesting choices". Turns out no one has much clue who said this first. http://flashofsteel.com/index.php/2008/07/07/quote-misquote-cite/ But the concept resonates with many designers. I'm going to focus on one special case: can a game have many "high-impact" moments, which can substantially change win chance? This is often critical in a game one plays to win.

And that's the rub. Is "a large number of impactful decisions" paradoxical? The quantity of decisions necessarily limits the impact of each. Elias et al. (Characteristics of Games) frame this question in terms of win probability over time. In a goal-oriented game, an impactful decision is one that substantially adjusts the win probability. But the more biased the current probability gets, the less impact a *typical* decision can have on it. If most of your options at turn T drastically change the win probability, then your current state by definition was very uncertain.

We see this problem arise in practice, across so many games. Games with flat scoring, like basketball, sometimes *feel* flat - each decision is equally relevant, and there are so many that it can feel like a battle of attrition.

Other games attempt to address this by propping up potentially low-impact moments with artifice. Mario Kart'S *beloved* blue shell does this with rubber banding. As Elias et al. point out, doing this simply de-values the impact of early decisions that superficially seem to matter. It puts the weight primarily on late decisions, and savvy players eventually come to realize this.

Conversely, games with substantial snowball mechanics, like MOBAs and RTS's, put more weight on early decisions. It's certainly possible to come back from an early game deficit, but it's more the exception than the rule.

So this is the Decision Impact problem. How do we make a game that avoids having players make decisions that feel low-impact? As before, some games accomplish this; but they usually do so through a fundamental compromise.

Decision Impact - Technique 1: Barriers

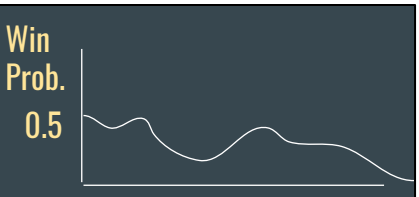

How can we excise low-impact decisions?

By using degrees of success, dampening the satisfaction of a binary success.

High-Score

Wide Ranking

Can we directly excise low-impact decisions? The constraints of the problem itself show us the way here. When trying to completely remove the bad thing, you really have to loosen those constraintsThink back to the graph we saw before. In my framing, high-impact decisions move the win probability. And on average, this probability drifts towards 1 or 0, once it starts changing. So to remove low-impact decisions, we have a couple options.

First, we can reduce the number of decisions. This isn't necessarily so bad! Short games can be exciting and high-stakes, like the indie fighter Samurai Gunn. What's more, the game needn't even be short. It just needs few decisions. Take One Night Ultimate Werewolf - a variant of the classic social deduction game, Werewolf/Mafia. Players make at most one decision individually, spend a great deal of time discussing, then make a single decision together, that decides the entire outcome of the game! Each of these games make a sort direct claim against the expected many-decision baseline of their genre!

Second, there's another way to deal with a graph that trends toward 0 or 1 too fast. Expand the magnitude of the graph! In other words, create more range of possible performances than win-loss. A game in which players attempt to maximize a score can help make every decision feel meaningful; each contributes to a potentially higher score. Arcade games have known this for a long time, as do modern speedrunners.

Unfortunately, serious players often end up setting their own binary goal, of beating the previous score. Note that in a PvP setting, it's often difficult to get players to care about degrees of winning. 1v1 games naturally imply a binary question to be answered. Uneven matchups sometimes motivate the stronger player to play for

heavy wins. More commonly, multiplayer games compel players toward the highest possible rank.

Decision Impact - Technique 2: Gates

How can we use *obfuscation* to reduce the pain of low-impact decisions?
By weakening heuristics, making all states *ambiguously* impactful.

Hidden Information

Intractability

How can we use difficulty to keep players out of the low-impact decision blues? There are two approaches here. One is to make it difficult to reach low-impact states; the other is to make it difficult to detect them.

In contrast, a game of high *uncertainty* makes evaluating win probabilities difficult. This uncertainty can come from randomness, hidden information, simultaneous play, or intractable complexity. (As Greg Costikyan discusses in “Uncertainty in Games”.) In each form, enough uncertainty weakens can weaken heuristics and make every decision feel *potentially* meaningful. Even while a god’s-eye-view of the game would already know the likely outcome. What you lose in these cases is a sense of control and manageability that many games of mastery thrive on.

Decision Impact - Technique 3: Carrots

How can players opt out of low-impact decision points?

Opt out of playing them.

Opting into early win

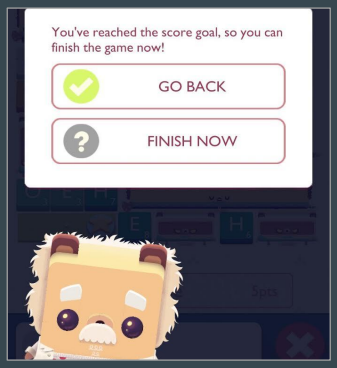

Risking early loss

We think of the impact of a decision point as something that happens to a player; they end up in a state and have to deal with it. So what does it mean for them to opt out? There's actually a common case here; the player recognizes the low impact of a state, and is given the opportunity to escape it!

Honestly, it's hard to conceive of how you'd give a player the option to opt out of *every* low-impact state...

We struggled with a low-impact decision problem in Alphabear, a word game by Spry Fox that I worked as a designer on. Player skill varied widely, as did the extent to which players chose to grind for power. As a consequence, it was hard to tune missions to the precise score targets that would produce tension around win/loss. We hand-wrung quite a bit on how to make an elegant system that would avoid this. But for Alphabear 2, a decision was made (after I left) to take a very simple approach. (By no means the first game to take this approach.) Simply give players an "instant win" button once they've exceeded the necessary score target. Yes, you lose tension. But for most players, it works! They have some side goals they can only achieve by completing the full mission. But if they want to barrel through a too-easy mission, they're able to do so with maximum quality of life. They just opt out.

In League of Legends, the game is interspersed with pivotal moments at which a team on the back foot can take a big risk to potentially turn the game around, but more likely accelerate their loss. For example, consider baron steals. (Even a single low-level player has the potential to steal baron, *if* they are able to get in. (But they

risk feeding the enemy even more.) In other words, from a low-win-chance state, they take an action that has a high probability of decreasing their win chance and reducing the expected time to game-end, but also has some probability of substantially increasing their win chance. It's a hail mary. These moments inject bursts of tension into the game, and let players choose not to draw the game out, even if it slightly decreases their overall win chance.

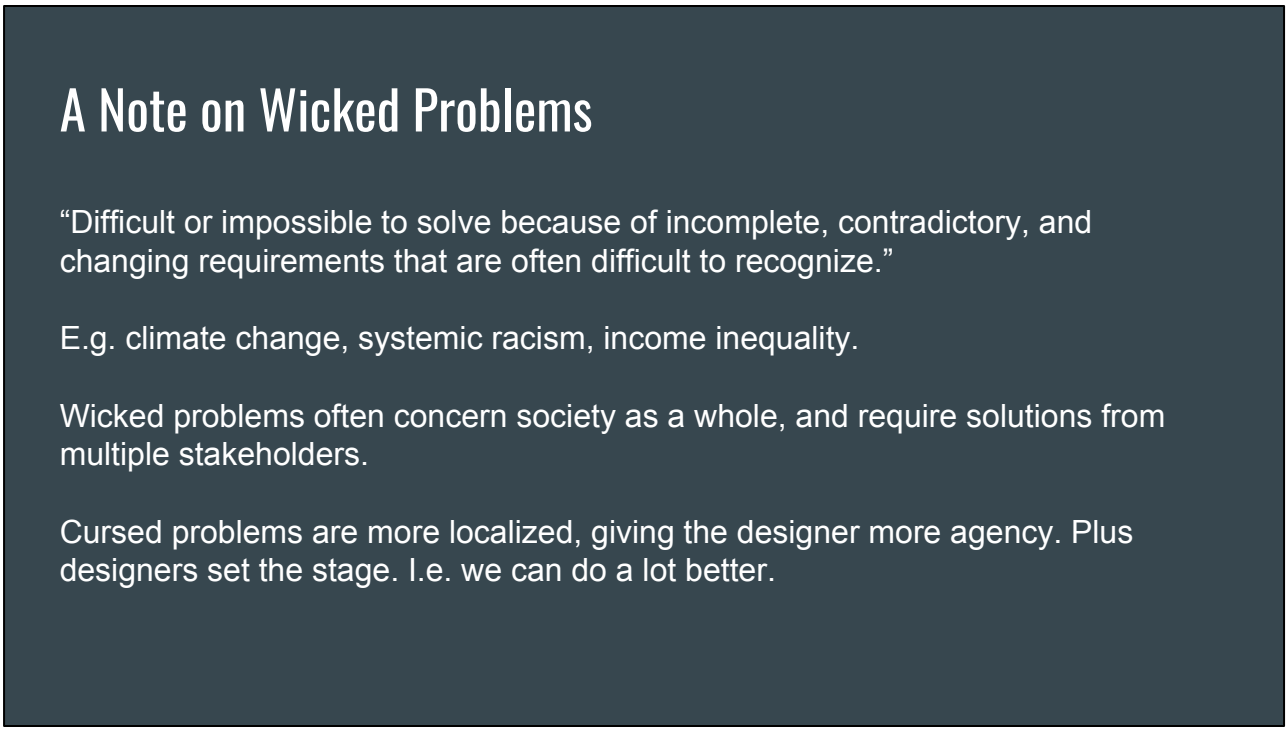

I chose the name Cursed Problems in homage to wicked problems, in the social sciences. A wicked problem is one that is “Difficult or impossible to solve because of incomplete, contradictory, and changing requirements that are often difficult to recognize.” That’s pretty on point for our conversation, right? But it’s important to distinguish them. Wicked problems are usually concern large-scale interpersonal interaction: politics, economics, environment, etc. Those who aim to intervene on wicked problems usually have less agency than game designers have on cursed problems.