

GM DON'T LIST #1: MORPHING REALITY

[by Justin Alexander – August 7th, 2017](#)

I got back from Gen Con yesterday. All in all, I had a really great time, picked up a bunch of cool stuff, and then came home to find a stack of Kickstarter games on my front porch. (Which I'm referring to as Gen Con Twelfth Night now that it's happened three times in a row.)

But I did run into something of a bad streak of luck this year when it came to actually *playing* games. Although there were a couple of really great games, for the most part I was mired in complete duds. In a few cases, this was the result of a bad scenario, but for the most part I was suffering from some truly atrocious GMing.

This caused me to reflect on the fact that GMing is one of those crafts where you very rarely get to see another practitioner's work. There are, of course, some groups where players will cycle through the GM's chair, but this seems to largely be the exception rather than the rule. As a result, I suspect that many (possibly most) GMs exist in a bubble, and this isolation limits their ability to recognize (and correct) their shortcomings.

I've written a great deal of advice for GMs here on the Alexandrian, and most of it is positive in nature. I like to tell you about the nifty stuff that you could be doing or adding to your games. But now I'm going to spend a little time taking the opposite approach: I'm going to talk about the stuff that you should NOT be doing.

The examples I'll be using here are primarily drawn from some of the bad experiences I had at Gencon. But since I'm not looking to specifically call anybody out, I'll be anonymizing the details.

GM DON'T #1: MORPHING REALITY

Player: Shit! Okay, we run out the other door!

GM: There is no other door. The only way out is the way you came in.

Player: What the hell happened to the other door?!

There's nothing more frustrating for a player (nor more guaranteed to shatter their engagement with the game) than watching the reality of the game world shift like a mirage in front of their metaphorical eyes.

In some cases, of course, this morphing of reality is the result of poor communication: The GM forgot to mention the ogre that is now pounding in their skull when they unwittingly tried to run past it. Or when the GM said that the doors were at the "end of the hallway" he meant they were facing each other across the width of the hall while the players assumed he meant they were standing side-by-side looking down the length of the hall.

Far more problematic, however, is when it becomes clear that the GM's own mental picture of the game world is inherently unstable. For example, there was a science fiction scenario where a ship's

compartment was filled with vacuum... or not. It oscillated randomly over several rounds of combat. (And if you're thinking that this would be distressing for anyone in that compartment, you would be correct. Maybe you're the character who carefully took the time to make the preparations to enter a vacuum only to be upstaged by someone else rushing into the compartment because the vacuum has vanished. Or maybe you're the character suddenly exposed to the vacuum because you entered the compartment after being assured that everything was fine now until the GM suddenly remembered that the vacuum was supposed to be there. Either way, everything else about the game will quickly become blotted out by your palpable frustration.)

I think it can be argued that maintaining the integrity of the game world is, in fact, the GM's most fundamental task. Everything else flows from that singular, unifying vision. And without that integrity meaningful choice ([the defining characteristic of the roleplaying game](#)) becomes impossible.

Perhaps the most common form of morphing reality is geographic: Distances that double or quadruple in size. Ogres who can somehow simultaneously be standing right next to two characters who are nowhere near each other. Hallways that appear and disappear from the floorplan. But it's a problem which can be found in any aspect of the game world: NPCs who change their appearance. Organizations that flip-flop between panopticon omniscience and bumbling cluelessness. Spells that vary in efficacy depending on the GM's mood.

On that last note, it's also important to note that, in addition to descriptive consistency, the GM should also strive to achieve *mechanical* consistency. As Ben Robbins' elegantly states in "[Same Description, Same Rule](#)":

The game world is imaginary. It does not exist except in the minds of the participants. Each person has their own mind and their own imagination, which makes it all the more important to make sure there is a consensus, that you are all operating in the same fictitious world and in agreement about how things work. Consistency makes that easier, inconsistency makes it harder.

To use an example from M&M, the players encounter one machine gun that uses a normal attack roll, and then later they encounter another machine gun that uses an Area attack instead (automatic hit, Reflex save to reduce damage). Conceptually the two machine guns are identical – one is bigger but otherwise the same.

A player sees the second machine gun before it fires and says "a ha, I will dodge to increase my Defense, which will make me harder to hit!" Logical but completely incorrect, because that player doesn't know that the second machine gun uses a rule mechanic that has nothing to do with Defense.

(...)

There's a simple fix for this:

The same description should never be modeled with two different rules. If you want to use a different rule, there should be a different description.



IMPROVING YOUR VISION

RPGs are the theater of the mind. They're improvisational radio drama. Achieving consistency means holding complex pictures of the game world in your mind while simultaneously juggling all of the other things the GM needs to be doing at the gaming table.

That's not easy.

The degree of complexity in that mental image that can be successfully managed will vary a lot between GMs. (This is also a skill, of course, and you'll generally find that you'll improve at it over time.) So the first thing a GM needs to do is *know their limits*. And once they know those limits, they can find ways to push beyond them.

For example, I know that I, personally, can't handle fights with simultaneous action resolution if there are more than 8-12 characters involved. Beyond that limit, I can't keep all of the disparate actions in my head at the same time and figure out how things would play out. So when that situation began cropping up frequently in my OD&D open table campaign, I responded by creating a mechanical structure that [split the round into multiple resolution phases](#) so that I was handling smaller groups of characters at any given moment.

If you struggle with keeping the geography of a locale consistent, sketch out a quick map. (This doesn't have to be a hyper-accurate blueprint: It just needs to cement the idea that the kitchen is here and the bedroom is over there.) Jot down notes on your NPCs to keep their appearance and characterization consistent. Keep a campaign journal so that you can track continuity between sessions.

IMPROVING YOUR CONVERSATION

Once you've got things straight in your own mind, you can work on improving communication with your players. Start by making a mental (or physical) note of moments when your players become confused by your descriptions. Review those moments: How could you have phrased things differently – or what details could you have added – in order to make things clearer? Over time you'll figure out which phrases (like “at the end of the hallway”) are too vague and how you can make them more precise.

(Note that precision does not necessarily equate to greater length.)

Creating clarity might also mean using visual references: Drawing diagrams, handing out photographic references for major NPCs, hanging a map of the city on the wall, and the like.

TRIAGE AT THE GAME TABLE

While you're working on all of that, keep in mind that no matter how skilled you become mistakes and misunderstandings will still happen. So give some thought about how you'll handle that confusion when it happens.

First, be alert for signs that a misunderstanding has occurred. Usually a dead giveaway is if the players are proposing actions which don't make sense to you. I've [talked about this](#) at length before and proposed a general principle:

If you don't understand what the players are trying to achieve with a given action, find out before adjudicating the action.

Second, give the benefit of the doubt to your players. Your vision of the game world is not precise and the situation of the game world is dynamic (even if your mechanics are breaking it down into sequential turns).

For example, there's an ogre fighting in the middle of a room and one of the players declares that he's going to run past the ogre. Clearly in their vision of the situation there's enough room for them to safely do that. Maybe in your vision of the room things aren't so clear-cut. But if you just have the ogre take a step to the left as he swings his club at Athena, then you could easily imagine Horatio rushing past him. So let it happen.



Third, find ways to compromise between your vision and what they want to accomplish.

Maybe it really is impossible for Horatio to just run past the ogre: It's a small room or maybe the ogre is specifically trying to prevent people from getting past it. Instead of just declaring it impossible, however, look at what the player is trying to accomplish (get past the ogre) and then offer them a way to do that which is consistent with your vision of the room: Maybe they can make a Tumble check to get past it. Or maybe Athena could deliberately bait it out of the way.

This process of compromise isn't just a specific application of [the "Yes, but..." principle of GM rulings](#) (although it is), it also smooths the players' correction of their mental picture. These mental pictures, after all, are built up from an aggregate of detail. By offering options for accomplishing their goals, you're encouraging them to focus on the additional details you're adding (the room is smaller than you thought) instead of on the rejection of their previous vision. (It's a subtle distinction, but in my experience it's significant.)

Finally, when the wheels come completely off the bus don't be afraid of allowing minor retcons to resolve the discordance.

If I'm following the advice above, I find this most often occurs when the consequence for the misunderstanding doesn't immediately manifest itself. For example, "If I'd known the ogres were close enough to get here before my next turn, I never would have stopped to pick up the idol!"

Even with the passing of some short span of time, it's usually still not too difficult to just back things up a step, correct the action taken under a misapprehension, and then move forward. But as chains of cause-and-effect become more complicated you do have to balance the potential discordance of the retcon against the discordance of the player's disconnection from the game world. (Also bear in mind that there is a difference between "the character didn't know" and "the player didn't understand": If the player thought the ogres were several hundred feet away when they were actually only a few dozen feet away, that's one thing. But if Horatio didn't realize that the ogres could traverse several hundred feet in a single round because of the cheetah totems they're wearing, that's a completely different thing.)

Once again, it can be useful to consider the compromise of a negotiated retcon: When Horatio grabbed the idol he was granted a brief vision, so you're not going to allow a retcon that wipes that moment out completely. But maybe you'll allow Horatio to avoid getting cut off by the ogres if he drops the idol and makes a run for it with a successful Athletics check.

GM DON'T LIST #2: ROLLING TO FAILURE

by Justin Alexander – September 30th, 2016



I was playing in a sci-fi game and my hacker wanted to set up a dead man's switch on the environmental systems in a space station: When the melee characters launched their assault, I'd pull the dead man's switch. The GM had me roll to set up the dead man switch... and then roll every single round to maintain it... and then roll again to throw it even after I'd set it up successfully.

Eventually, given the endless series of checks, I inevitably failed.

Laying aside the fundamental misunderstanding of what a dead man switch is, the negative effect of this sort of thing can be quite severe in an RPG session. In this case, the group quickly realized that we should never, ever try to make a plan: If we just improvised something, it would be resolved in a single roll and we'd have a chance to succeed. If we actually put together a plan, on the other hand, it would just invite lots and lots of dice rolling until the plan failed.

The solution is fairly simple: [Let It Ride](#). Have the character make a single skill check that determines the ultimate success or failure of their endeavor.

Another solution is a complex skill check (making multiple checks until X successes are achieved). These tend to be very elegant in dice pool systems, and when you want multiple checks to be made they're an effective framework for allowing that without the all-or-nothing of a single check ruining the entire attempt.

GM DON'T #2.1: FAILURE IS POINTLESS

The flip-side of rolling to failure is the "roll pointlessly until you succeed" thing. For example, you'll often run into games where the PCs need to unlock a door: There's no time pressure and no consequences for failure, and yet the GM will sit there and have the PCs roll over and over and over again until they finally succeed.

One way to deal with that is something like the Take 20 mechanic: If you *can* eventually succeed at this, then we can assume that you *will* eventually succeed at it and we can move on. Letting it ride can also

solve this problem by providing the opposite outcome: Your failure on this Open Locks check tells us you are simply not good enough to pick this lock at this time and in this way. (The single check determines your relationship with the lock and until you can substantially change the situation, your character is going to be stymied by that lock).

GM DON'T #2.2: TOO MANY SEARCH CHECKS

A somewhat related problem is when a [multi-step action resolution](#) gets broken down into too many discrete parts. This can take many forms, but the most cancerous form I've seen in the wild came from GMs who took the Search guidelines for 3rd Edition D&D way too seriously. Those guidelines specified that it took a full action to search a 5-foot square. That's a useful guideline for combat (when you might want to know how much area you can search during a single round), but some GMs took this to mean that you needed to make a separate Search check for every 5 foot square. So if you searched a 10-foot-wide hallway that was 40 feet long, you'd have to make sixteen (!) separate Search checks.

This isn't rolling to failure because each chunk is a legitimately separate task. (Failing to search in Square #1 doesn't mean you won't find anything in Square #2.) But it murders pace – which is either directly undesirable or undesirable because it discourages players from using the specialties affected by the problem.

The solution here is to collect the tests into meaningful chunks: Searching an entire room (or even suite of rooms) is obvious. Alternatively, if they want to search the dungeon hallways as they move along, let the result of the check ride until they either make a meaningful choice to do something other than search down the hall OR until that check result produces a result (either success or failure) that they can recognize as such (i.e., until that check either finds a trap or secret door or until it *fails* to do so and the trap happens to them or the ambush pours out of the secret door they missed).

Sometimes you'll end up with a player who demands multiple checks. In some cases this is because they, too, are following bad mechanical advice (like the “make a check for every 5 feet” misinterpretation of the Search rules). In other cases, it's a manipulation of metagame information (“I know I rolled poorly, so let's have that only apply to this one specific area and then I'll make another check”). Often it's because they're irrationally trying to manage risk (“I'll only search this little chunk so that if I roll poorly the effects will be minimized” – which doesn't make sense because your odds of discovering any given hazard remain unaltered, but that doesn't mean people don't do it).

Most of the time your response to this is fairly simple: You tell them no.

The exception would be the rare instance where it's actually *effective* pacing to stretch out the mechanical resolution. Like a slow motion shot in a film, these are the times when specifically highlighting each small, discreet, tension-filled moment serves to escalate the crisis and leave the table on the edge of their seats.

Identifying these moments is a gut-check, not a science. For example, I was just about to say that it would *never* be Search checks down a dungeon hall... but then I realized that there actually *was* a time that I followed a player's lead in the [Tomb of Horrors](#) to separately search every single inch of corridor because that mechanical resolution was so completely *right* in capturing the paranoia and terror the group was experiencing in that moment.

LET THE GOOD TIMES ROLL

On a closing note, let's be clear that not *every* series of sequential rolls with a non-discrete outcome is rolling to failure. We've already discussed the situation where each individual check is a separate, meaningful accomplishment. But it's also true that, for example, combat isn't a roll to failure even though it involves multiple checks culminating in a single outcome of life-or-death.

It's also useful to note that rolling to failure can be an *effective* choice if you're actually looking at a situation where failure is assured and the interesting question is how long a character can stave off that failure. For example, how long can you stay conscious in a vacuum? How long can you hold the door against the werewolves pounding on it from the other side?



GM DON'T LIST #3: RESOLUTION DITHERING

[by Justin Alexander - November 14th, 2016](#)



GM: Okay, you come up over the horizon of the station and you can see the trench up ahead. Three rebel fighters go roaring past.

Annie: I signal my wingmen to follow my lead and drop in behind them.

GM: Sure. You fall into their 6 o'clock and hit the thrusters, zooming up behind them.

Annie: I target the lead rebel pilot and take my shot!

GM: He's dancing around in the ray-trace of your targeting computer.

Annie: The Force is strong with this one. I pull the trigger!

GM: The walls of the canyon are really racing past you. All this amazing superstructure just whirring by in a blur.

Annie: Great. I take the shot.

GM: Suddenly that old YT freighter you'd planted the tracking device on earlier comes roaring out of deep space! It shoots! [rolls some dice] One of your wingmen explodes!

Annie: What?!

GM: What do you do?

Annie: I... take my shot?

GM: Your other wingman, distracted by the appearance of the new enemy, loses control! They smash into your wing, careen wildly, smash into the wall of the canyon, and explode! Your own stabilizers have been damaged and you go hurtling out into deep space!

This kind of resolution dithering – where the players have declared their actions, but the GM isn't allowing them to actually take and resolve those actions – is incredibly frustrating.

Sometimes the dither is caused by the GM prematurely asking the players what they want to do – after hearing the proposed action they realize that there's additional information that they want or need to convey. (Or, if they're improvising, details or cool ideas which popped into their head during the time that it took for the player to respond.)

Other times the dither occurs because the GM is waiting for someone to say the thing he wants them to do: Something cool is going to happen when someone tries to open the door, so any other action people propose will be put on pause until somebody in the group opens the door. (This also naturally leads to a narrower case in which only actions that would disrupt what the GM has planned are ignored – you can do anything unless it gets between them and that door.)

Another common form of dithering occurs when a GM responds to a declaration of action by discussing other options that are available. For example, I was playing in a cyberpunk game where I said I wanted to hack an electronic lock. The GM responded by pointing out that I could also kick the door down or just send my slither-bot under the door or physically pick the lock or...

Ultimately, when a player declares an action the GM needs to resolve that action and then describe the new situation: They need to move forward so that the next set of actions be cleanly declared. (The only exceptions are [if the GM doesn't feel they have enough information to resolve the action](#) or [if the declared action appears to be based on a misunderstanding of the current situation](#). In either case, the GM should still be seeking the necessary clarification as quickly as possible and then driving forward into the resolution.)

Resolution dithering often becomes obfuscated when the GM can jump between multiple PCs, leading to a muddle where the GM can get an action declaration from one PC, not fully resolve it, move onto the next PC, get another action they don't fully resolve, and then repeat cyclically – kind of bouncing around the group without ever moving the action forward. This seems particularly prevalent with neophyte GMs (possibly because their lack of confidence manifests as an unwillingness to make the sort of definitive declarations required of action resolution), and the resulting quagmire can be difficult to diagnose.

GM DON'T #3.1: THE REVERSE RESOLUTION RING

What I refer to as the *reverse resolution ring* is a kissing cousin with resolution dithering and, for me, is even more frustrating to experience as a player.

For example, I was playing in a game of *The One Ring*. The GM would describe a situation – like a guard dog growling as the party drew near – and I would say something like, “Okay, I'm going to grab some of the fresh venison from the deer we killed this morning and I'll toss it to the dog to distract it.” The GM takes note of that, but then proceeds around the table collecting action declarations from the other players.

So far, this is probably fine: Getting a collective understanding of what everyone is doing before figuring out how it would all play out together can actually be a really good technique for a GM to learn.

But where the reverse resolution ring kicks in is when a form of recency bias causes the GM to resolve the proposed actions in the opposite order from which they were declared (starting with the last person they

talked to and then working their way backwards to the person who actually kicked things off). This is a problem because, at some point during those declarations, the other players will often say something like:

“Oh! That sounds good! I’ll dig some meat out of my pack, too!”

Or:

“I shoot the dog with my crossbow.”

The latter negates the original declaration by solving the problem in an alternative way. The former ends up basically stealing the original idea (even when the player saying it was just trying to support what they saw as a good solution to the problem) – the copycat gets to be the one to actually do the cool idea.

In either case, the GM is essentially stealing spotlight time. They’re punishing the player who took initiative, which is directly problematic because that’s demoralizing and unfair to the player affected, and indirectly problematic because it will eventually have a corrosive effect on the willingness of the entire table to step up. Even if it’s just a subconscious reaction, eventually you’ll end up with something that could easily be misidentified as analysis paralysis, but is actually just a hesitation to pull the trigger when it’s just as likely to end up shooting you in the head.

(It actually reminds me of something that crops up in live theatre: One actor will come up with a funny bit of business or line reading. Other actors will see it and think, “That’s hilarious!” And then they’ll end up duplicating the bit in their own scene, which can often happen earlier in the play than the original actor’s bit. These derivative bits are often not as funny and only serve to sap the riotous humor of the original – which is often built on the straight takes which are supposed to precede it. It’s the director’s responsibility to make sure that this sort of undercutting *does not happen*. But I digress.)

The reverse resolution ring can get truly cancerous when it turns into an endless ring: The GM goes through the ring once asking declarations, goes backwards through the ring resolving actions, and then – since they’re back at the beginning of the ring – they ask that last player, “So, what do you want to do next?” ... only to then go forwards through the ring again getting everyone else’s declarations. The GM can even convince themselves that they’re “balancing” things – this guy went last, so let’s find out what he wants to do first. But that player is now systemically screwed, doomed to forever get upstaged by the rest of the group until something disrupts the current pattern.

In closing, however, I will mention the exception which proves the rule: A reverse resolution ring can be an effective technique when it’s used to model initiative. In other words, when the GM asks those with the lowest initiatives to declare their actions first and then resolves from highest initiative down. The “punishment” is now modeling the poor initiative result, and grants a strong benefit to those with a high initiative result.

GM DON'T LIST #4: THOU SHALT NOT HACK

[by Justin Alexander – October 3rd, 2017](#)

A cyberpunk character concept I would dearly love to play some day is that of the uber-competent hacker: Case from [Neuromancer](#). Batman's [Oracle](#). Edward from [Cowboy Bebop](#). Boris Grishenko from [GoldenEye](#). Luther Stickell from the [Mission Impossible](#) movies. Half the main characters from [Ghost in the Shell](#).

In my ideal version of the character, I'm the guy who stays in the van three out of five times, providing overwatch and support for my teammates while they mount their raid.

So why haven't I played this character?

Because what I've discovered is that a surprisingly vast number of GMs seem to consider the entire concept of using hacking to solve a problem to be some sort of anathema. So even when I've *tried* to play the character concept, I've ended up not actually being able to play the concept.

- Hack the security cameras to scope out the interior of the building you're raiding? Can't do it.
- Hack the security guard's cellphone to track her movement? Impossible!
- Play R2-D2 in a *Star Wars* game and hack an electronic lock? No way. Pull out your lockpicks!


Hack the rigging ports on the pursuit car to seize control of it? Obviously no one would *want* you to remotely seize control of a vehicle, so they would build a perfect security system that was completely unhackable, and therefore you can't hack it.

(That last rationalization seems to crop up a lot. It's like saying that obviously no one would *want* you to poke them with a sword; ergo it's impossible to hit someone who's wearing full plate.)

Quibble here and there with the plausibility of some of these scenarios in particular settings, but I've seen this behavior even in settings and games which include mechanics for handling these specific types of hacks! What I'm talking about is a systemic pattern of behavior in which the hacker basically can't do their thing. It's the equivalent of finding an *antimagic field* everywhere you go in a D&D game, except that I've found it to be a peculiarly ubiquitous attitude.

Of course, flat out denial isn't the only way this manifests: Setting disproportionately high difficulty numbers or using [roll to failure](#) techniques are probably the most common versions, actually.

I've found this particularly pernicious in many convention scenarios: The designers of the game want to show off its breadth, so they include a hacker archetype pregen. But the volunteer GM running the scenario subscribes to the doctrine of Thou Shalt Not Hack, so the pregen is a trap and the person picking it finds themselves sidelined for four hours.



The worst case scenario is, sadly, one of my favorite games: [Eclipse Phase](#). Wanting to show off everything that's possible in this cool, kitchen sink transhuman setting, the designers regularly include an infomorph pregen: A character without a body who exists only as a digital construct and can only take actions through the Mesh network.

Combine that with a GM who doesn't allow any meaningful action to take place through the Mesh network (which I've seen happen either first- or second-hand in no less than four convention scenarios) and you have a character who literally can't do *anything*.

Many of these GMs don't seem to be consciously aware of what they're doing, so you'll even find them saying things during character creation like, "Who wants to be hacker?" I used to hear that and think, "Okay! This guy is going to actually let me hack!" But, oddly, no. They recognize on a conscious level that a team of cyberpunk characters is supposed to have a hacker, but when it comes to actual play the hacker nevertheless finds themselves stymied at every turn.

EMBRACING THE HACK



I suspect part of the problem here is that a lot of GMs reflexively cling to the modes of play they learned running D&D dungeon crawls. Their expectation for how a facility raid is supposed to play out features people physically sneaking around and getting ambushed by security guards, and the hacker's attempt to grab the security cameras disrupts that expectation. Their vision of the game world (inaccurately) doesn't include hacking, so the hacker's solution to any given problem comes out of left field, and the GM reflexively shuts it down.

This is, obviously, a form of railroading: A preconceived idea of not just how a specific problem is meant to be solved, but a broad preconception of how entire *classes* of problems are supposed to be solved.

So the solution to this problem is relatively simple: Don't do that.

Conversely, however, hacking shouldn't be a magic button that can trivially solve all problems. When that happens, it creates [a spotlight problem](#) where the hacker upstages every other character and flattens the challenges presented by the scenario.

To counteract this problem, there are a couple things the GM should do. First, check the potential consequences a hacker faces: They should be comparable to those faced by other types of action. (Just as the hacker should not find it impossible to hack an automated car; the hacker themselves should not

benefit from a foolproof firewall.) Second, check your [vectors](#): Make sure that “solving” the scenario requires a multi-step resolution and, importantly, make sure that hacking can’t be used to trivially solve all the vectors.

The most obvious example of this is, “I can’t hack that system until you plug in my remote router!” But it can become an easy trap to always design scenarios in which the team does a bunch of stuff and “unlocks” the hacker so that the hacker can then win the day. Look at ways in which hacks are invaluable at the beginning and in the middle of scenarios.

Also remember that you don’t always need to lock these things in: Players hot-swapping in vectors you’d never thought of to solve their problems is what makes the game fun. Generally speaking, the rest of the group will find ways to advocate for plans which feature the strengths of their own characters if you give them the chance. You can encourage that by creating scenarios which require multiple problems to be resolved simultaneously. Also experiment with using [hard scene framing techniques](#) to move the action “onsite”, which will discourage the players from lingering in remote “planning” sequences where the hacker (and only the hacker) is capable of taking direction action.



GM DON'T LIST #5: NOT KNOWING THE RULES

[by Justin Alexander – January 3rd, 2018](#)



Infallibility is not, in fact, a requirement for a game master. Indeed, the idea that the act of GMing requires some sort of savant is a pernicious one which was, sadly, robbed the world of many fabulous GMs and many tables filled with happy gamers.

With that being said, one of the GM's responsibilities is, in fact, to provide a certain level of rules mastery. How you achieve that level of mastery is largely dependent on your own personal study habits. For me, the typical procedure is:

First, read the rulebook cover-to-cover. If you haven't read a rule at least once, then you've really got no chance of getting it right.

Second, [prepare a comprehensive cheat sheet for the system](#). The process of organizing and compiling the cheat sheet is, by itself, a great way to get a grasp on how they work and relate to each other (and also sussing out those minor mechanics you would otherwise gloss over). Once you're done, of course, the cheat sheet becomes invaluable at the actual gaming table, artificially supplementing your own knowledge. As we'll see, being able to quickly and accurately reference information is almost as good as knowing it off the top of your head.

(I've also talked before about how the [hierarchy of reference](#) can be used to progressively gain system mastery.)

Third, when I think it might be warranted (or fun!), I'll also run "playtest" one-shots using the system. These are a great way for both the GM *and* the players to gain familiarity with the system and work out its kinks before diving into a long-term campaign. (For players, I've tangentially found that this familiarity often makes for a richer and more engaging character creation process. Knowing how a game works provides really valuable context for the mechanical decisions you make when building your character.)

BUT WHY?

Some may wonder why this rules mastery is important. I've even met GMs who, for nearly incomprehensible reasons, take great pride in being largely ignorant of the rules. (This seems related to the school of thought which maintains that the rules of an RPG are just kind of a pleasant fiction that the players improv vaguely around / the GM uses only when necessary to reign with an iron fist.)

First, presentation and **pacing**. Nothing deflates excitement or undercuts tension at the gaming table faster than, "Hang on, let me just figure out how the rules for this work." [The Art of Pacing](#) mostly discusses macro-scale pacing, but pacing at the micro-scale is just as important: Keeping things flowing smoothly; maintaining (and escalating) the mood; sustaining player focus and attention. All of these things require the rules to flow out smoothly, cleanly, and accurately not only to minimize friction, but also because high quality rules that are effectively applied will *enhance* these things.

Speaking of which, **quality rulings** require both knowledge and comfort with the rules. Any master craftsman or artist knows the importance of being intimately familiar with your tools, and the art of the GM is no different. A good GM will make the rules sing, finding ways to combine and recombine them to achieve (and help their players achieve) delightful and unexpected things. But you have to fully understand your tools before you can start truly playing with them.

Third, **consistency**. In many ways, this is actually just a special case of [GM Don't List #1: Morphing Reality](#). If the GM doesn't know the rules, then their application of the rules will become inconsistent and unpredictable. This inconsistency results in the game world acting in weird and unpredictable ways, which inevitably frustrates the players: They see a lock and expect that they'll be able to use their Criminal skill to pick it because that's what they did last time; but this time the GM decides (or realizes) that it should *actually* require an Infiltration check to pick a lock and the players discover that they've sent the wrong person to deal with the problem.

Finally, when the GM doesn't know the rules — and isn't using them correctly — it preemptively **shuts down certain styles of play**. For some players, these elements of play are very important; for others less so. But either way, their loss will generally result in a flattened and less interesting gaming experience.

Not infrequently when I'm discussing these issues, these styles of play will be dismissed by the narrow minded as just "goofing around with mechanical widgets". But it's not that simple. Yes, there are those who play roleplaying games, in part, to have the satisfaction of overcoming (or outsmarting) specifically mechanical challenges. But mechanics permeate every aspect of an RPG, and their effect can be felt in many different styles of play. For example, there is satisfaction and enjoyment to be had in building a character who is very good at something and then doing that job well (just like the satisfaction of any job done well). When the rules suddenly shift and the mastery that you *should* have had suddenly ceases to exist, that can be an incredibly frustrating experience for players.

(And, in this sense, you may realize that [GM Don't List #4: Thou Shalt Not Hack](#) is, in fact, a special case of this general rule.)

GM DON'T #5.1: IGNORING THE RULES

As a corollary, it's also important that GMs don't habitually ignore the rules.

As I can already sense hackles rising across the internet, let me make it clear what I'm NOT talking about:

- **House rules.** You're not ignoring the rules when you decide to explicitly change them in order to better your game.
- **Variants stat blocks.** If you decide to give an orc a +1 *sword* or bump up a troll's Strength score, that's not ignoring the rules either. (For some reason there are people who think so, or who categorize this as "cheating". These people are, frankly, insane.)

Now that I've hopefully soothed some hackles and raised a different set of them, let's delve into this a little bit.

The main thing to notice is that when you ignore the rules you are actually stumbling directly into almost all of the exact same problems that occur when you're simply ignorant of them: Consistency necessarily deteriorates, which subsequently tanks the quality of your rulings and also creates the same frustrations from players depending on consistency in order to understand both the game world and their characters.

If you consistently find yourself ignoring (or wanting to ignore) a particular set of rules, that's an indication that those rules are fundamentally broken (at least for the experience you want to create) and you should be looking to *fix* them (or replace them entirely), not simply ignore them.

A common example of this are grappling rules. (Across most systems, really, but infamously so when it comes to virtually all editions of D&D.) And the solution is, in fact, to apply house rules [which make grappling appealing instead of a chore](#).

One particularly pernicious example of this which certain GMs endemically suffer from is, "I'm bored with combat let's skip it." (Or, really, any other aspect of game play. It's just that combat seems most common here.) This usually takes the form of resolving 1-3 rounds of combat normally and then saying, "Eh. Fuck it. Let's just sum up what happens and move on."

The GM's intention here is good: They sense that the game is getting boring and they want to fix it. But in doing so they systemically create a number of other problems:

- Characters built to enjoy their spotlight time during combat are being punished.
- Strategically clever and creative players often spend the first few rounds of combat setting up an advantageous situation that will give them a big, satisfying pay-off as the combat continues. By cutting combat off just as they finish their set up, the GM is perpetually blue-balling them.
- Because they're never certain exactly when (or if) a particular combat is going to be summarily dismissed, players become uncertain in their use of limited supplies. Burning a one-use potion or once-per-day ability only to have its use become irrelevant when the GM decides combat has become too "boring" to continue is incredibly frustrating.

All of these problems only get *worse* when the GM defines "boring" as "the PCs are winning", while remaining fully engaged and excited as long as his bad guys have the upper hand.

BUT RULE ZERO!

“But it’s the GM’s god given right to change or ignore the rules at their whim!”

Sure. But insofar as we agree that this is a power which a GM has, I would argue that its use should be considered, deliberate, and, above all, *limited*. More generally on this topic, I would tend to make three final observations:

Calvinball is a really funny joke, but it is, in fact, a joke. There’s a reason why games have rules, and RPGs are no exception. [System matters](#).

In my experience, the motivations GMs have for unilaterally ignoring the rules tend to be shitty ones. Virtually all of them, in fact, rhyme with “tailroad”.

But let’s assume that the GM has accurately identified a truly singular instance in which the rules should be ignored (instead of permanently changed) without letting their players know (instead of explaining the ruling they’re making and why it varies from the norm) in order to truly increase the table’s enjoyment of the game. Here’s my question:

What gifts the GM with the unique capacity to recognize when the application of a rule would be a bad idea for the game?

If you’d be equally happy with the other players at the table unilaterally deciding to fudge a dice roll or pretend that their skill rating is higher than it is or act as if their character has an ability that isn’t on their character sheet, then more power to you. But what I see at the table (and usually observe in these hypothetical discussions online) are hypocrites who simply feel that their opinion is infallible, but the judgment of everyone else at the table can’t be trusted.

TRIAGE AT THE TABLE

In reality, of course, nobody is perfect. Nobody is a walking encyclopedia. (Or, if they are, it’s the result of years or possibly decades of experience with a system.) Mistakes will be made. Rules will be forgotten or overlooked. That’s okay. The GM has to become comfortable with their fallibility so that they can deal with the consequences when they arise.

So what happens when you forget a rule at the table?

I’ve already mentioned cheat sheets. Permanently bookmarking frequently referenced sections of the book also helps. ([Post-It Memo Flags](#) are great for this.)


Also: Use the expertise of your players. Don’t be afraid to ask, “Does anybody remember how much damage a *fireball* does?” There are far too many GMs who are so terrified of the rules lawyer boogeyman that they won’t take advantage of the communal brainpower of the gaming group as a whole. (I’ve also found that some rules lawyers behave better when they can apply their rules expertise in this way. Not all, but some.)

Another very effective technique, particularly in combat, is to delegate someone else to look up a rule while you move onto and begin resolving the next action. You can then jump back to the original action when the rules reference is ready. (The multitasking keeps the game moving forward through the rules reference instead of creating a dead space.)

Finally, if a particularly obscure rule is escaping all efforts to clarify it, don’t be afraid to make an *ad hoc* ruling while making a note to come back and check what the actual rule is during the next break or after the session. It’s okay to trade strict accuracy to keep the pace up. (It’s also, in my experience, a good idea

to openly tell your players what you're doing. It doesn't hurt if you give the PCs the benefit of the doubt when making these sorts of rulings, either. [Default to yes](#), after all.)

Mistakes will be made and sometimes your current mastery will prove insufficient for the challenges of the moment. But as long as you handle these moments with openness, clarity, and goodwill, you'll come out on top. And, of course, the cliché is true: Every mistake is a learning opportunity. Every mistake can make you a better GM... if you let it.



What I'm interested in talking about today is the second variety of Choose Your Own Adventure technique, which I suppose we could call:

RUN-TIME CHOOSE YOUR OWN ADVENTURE

GM: You see that the wolf's fur is matted and mangy, clinging to ribs which jut out through scrawny skin. There's a nasty cut along its flank. It snarls menacingly at you. Do you want to attack it? You could also try offering it some food.


With run-time choose your own adventure, in addition to describing a particular situation, the GM will also offer up a menu of options for how the players can respond to it. In milder versions, the GM will wait a bit (allowing players to talk through a few options on their own) before throwing in his two cents. In the cancerous version, the GM will wait until a player has actually declared a course of action and *then* offer them a list of other alternatives (as if to say, "It's cute that you thought you had autonomy here, but that's a terrible idea. Here are some other options you would have come up with if you didn't suck.").

It can be an easy trap for a GM to fall into because, when you set a challenge for the PCs, you *should* be giving some thought to whether or not it's soluble, and that inherently means thinking through possible solutions. It's often very easy to just burble those thoughts out as they occur to you.

It's also an easy trap to fall into during planning sessions. Everyone at the table is collaborating and brainstorming, and you instinctively want to jump into that maelstrom of ideas. "Oh! You know what you could do that would be really cool?"

But you have to recognize your privileged (and empowered) position as the GM. You are not an equal participant in that brainstorming:

- As an arbiter of whether or not the chosen action will succeed, you speak with an inherent (and, in many cases, overwhelming) bias.
- You've usually had a lot more time to think about the situation that's being presented (or at least the elements that make up that situation), which gives you an unfair advantage.
- You often have access to information about the scenario that the players do not, warping your perception of their decision-making process.



The players, through their characters, are actually present in the moment and the ideas they present are being presented *in that moment*. The ideas that you present are interjections from the metagame and disrupt the narrative flow of the game.

Because of all of this, when preemptively suggesting courses of action, you are shutting down the natural brainstorming process rather than enabling it (and, in the process, killing potentially brilliant ideas before they're ever given birth). And if you attempt to supplement the options generated by the players, you are inherently suggesting that the options they've come up with aren't good enough and that they need to do something else.

So, at the end of the day, you have to muzzle yourself: Your role as the GM is to present the situation/challenge. You have to let the players be free to fulfill *their* role, which is to come up with the responses and solutions to what you've created.

As the Czege Principle states, "When one person is the author of both the character's adversity and its resolution, play isn't fun."

But more than that, when you liberate the players to freely respond to the situations you create, you'll discover that they'll create new situations for *you* to respond to (either directly or through the personas of your NPCs). And *that's* when you'll have the opportunity to engage in the same exhilarating process of problem-solving and roleplaying, discovering that the synergy between your liberated creativity and their liberated creativity is greater than anything you could have created separately.

WITH NEW PLAYERS

This technique appears to be particularly appealing to GMs who are interacting with players new to roleplaying games. The thought process seems to be that, because they're new to RPGs, they need a "helping hand" to figure out what they should be doing.

In my experience, this is generally the wrong approach. It's like trying to introduce new players to a cooperative board game by [alpha-quarterbacking](#) them. The problem is that you're introducing them to a version of a "roleplaying game" which features the same preprogrammed constraints of a board game or a computer game, rather than exposing them to the element which makes a roleplaying game utterly unique — the ability to do *anything*.

What you actually need to do, in my general experience, is to sit back even farther and give the new players plenty of time to think things through on their own; and explicitly empower them to come up with their own ideas instead of presenting them with a menu of options.

This does not, of course, mean that you should leave them stymied in confusion or frustration. There is a very fine line that needs to be navigated, however, between instruction and prescription. You can stay on the right side of that line, generally speaking, by framing conversations through Socratic questioning rather than declarative statements: Ask them what they want to do and *then* discuss ways that they can do that, rather than leading with a list of things you think they might be interested in doing.

WITH EXPERIENCED PLAYERS

You can, of course, run into similar situations with experienced players, where the group has stymied itself and can't figure out what to do next. When you're confronted with this, however, the same general type of solution applies:

A few things you can do instead of pushing your own agenda:

- Ask the *players* to summarize what they feel their options are.
- In mystery scenarios, encourage the players to review the evidence that they have. (Although you have to be careful here; you can fall into a similar trap by preferentially focusing their attention on certain pieces of information. It's *really* important, in my experience, for players in mystery scenarios to [draw their own conclusions](#) instead of feeling as if solutions are being handed to them.)
- If they've completely run out of ideas, bring in a proactive scenario element to give them new leads or new scenario hooks to follow up on.

Also: This sort of thing should be a rare occurrence. If it's happening frequently, you should check your scenario design. [Insufficient clues in mystery scenarios](#) and [insufficient scenario hooks](#) in sandbox set-ups seem to be the most common failure points here.

This problem can also be easily mistaken for the closely related situation where the group has too *many* options and they've gotten themselves locked into analysis paralysis. When this happens, it should be

fairly obvious that tossing even *more* options into the mix isn't going to solve the problem. A couple things you can do here (in addition to the techniques above, which also frequently work):

- Simply set a metagame time limit for making a decision. (Err on the side of caution with this, however, as it can be very heavy-handed.)
- Offer the suggestion that they could split up and deal with multiple problems / accomplish multiple things at the same time.


The latter would seem to cross over into the territory of the GM suggesting a particular course of action. And that's fair. But I find this is often necessary because a great many players have been trained to consider "Don't Split the Party" as an unspoken rule, due to either abusive experiences with previous GMs or more explicitly from previous GMs who don't want to deal with a split party. That unspoken rule is biasing their decision making process in a manner very similar to the GM suggesting courses of action, and the limitations it imposes often result in these "analysis paralysis" situations where they *want* to deal with multiple problems at the same time, but feel that they can't. Explicitly *removing* this bias, therefore, solves the problem.

You can actually encounter a similar form of analysis paralysis where the players feel that the GM is saying "you should do X", but they really don't want to. Or they'd much rather be doing Y. And so they lock up on the decision point instead of moving past.

Which, of course, circles us back to the central point here: Don't put your players in that situation to begin with.

GM DON'T LIST #7: PREEMPTING INVESTIGATION

[by Justin Alexander – March 21st, 2018](#)



GM: It's a pretty cheap lock, so it only takes you about fifteen seconds to pick it. You hear the satisfying click.

Rachael: Great. I'll slide my picks back into the hidden lining on my belt before opening the door and slipping through.

GM: You find yourself in the office of Sir Sebastian. An imposing, mahogany desk with a flared plinth dominates the center of the room. Heavy, velvet curtains with gold appliqué seem to swallow the light from the windows. Vivid, arsenic-green wallpaper render kaleidoscopic patterns on the walls. Give me a Search check.

Rachael: 25.

GM: Okay, you find a hidden compartment on the wall, which you open by tracing the patterns in the wallpaper. Inside you find a small, metal ball with black, acid-etched symbols covering its surface. Give me a Spot check.

Rachael: 18.

GM: You notice that there's a thin seam running around the center of the ball. Give me an Idea roll.

Rachael: 16.

GM: Okay, that just good enough. You realize that the ball can be rotated to form different patterns with the symbols. You experiment for a minute,

and find a sequence that causes the ball to pop open. Inside you find Marie Artaud's ring.

Rachael: Great. I'll take the ring, close the ball, and get back to the party before I'm missed, making sure to lock the door behind me.

Hopefully the problem here is immediately apparent to you: The GM is cutting off the player's investigation of the scene by preemptively calling for skill checks. The PC effectively ends up in a kind of "autopilot mode" during which the game ceases to be truly interactive and the player is rendered into a passive audience that can only watch the character's actions playing out.

It's rare (although, unfortunately, not unheard of) for this error to be carried out in quite so egregious a fashion, but I've found that its less pronounced variants are *shockingly* common.

THREE-TIERED PERCEPTION

Probably the most common version of this problem that I've seen is when the GM preemptively calls for a Search check or similar mechanic. At a minimum, however, a good GM needs to be able to distinguish between three different levels of character perception:

1. Automatic Perception
2. Spot-type Perception
3. Search-type Perception

I'm using skill names from 3rd Edition D&D, but this remains true even in games which don't mechanically distinguish between these categories. ([Pathfinder](#), for example, is 3rd Edition's kissing cousin, but lumps both Spot-type and Search-type perception into a single skill.)

If you're familiar with the [Art of Rulings](#), you may notice how these fall into its three core principles:


1. Passive Observation is automatically triggered
2. Player Expertise activates Character Expertise
3. Player Expertise can trump Character Expertise

Automatic Perception and Spot-type Perception both fall into the category of Passive Observation: Automatic Perception is the stuff that literally anyone standing there will observe. (If you want to think about it [in purely mechanical terms](#), it's the stuff that requires a DC 0 Spot check to notice.) Spot-type Perception is the stuff that people *can* notice while just standing there, but might not. (Spot checks are an example of this, but so are Knowledge checks: Anyone can see the large flag hanging on the wall, but only some people will recognize what nation the flag belongs to.)

Search-type Perception falls into the second category, being an example of Player Expertise activating Character Expertise: This is the stuff you can't see by just standing there. You need to go *do* something in order to see it / learn it.

Beyond this basic core, there are a few advanced techniques to consider.

Matryoshka Search Technique: This is something I've discussed in a dedicated post as a [Random GM Tip](#), but beyond the threshold of the basic Search-type Perception, you can begin to see the game space as nested layers of interaction.



You can actually see this in the example above: Rachael needs to search the room to find the hidden panel. She needs to figure out how to open it. Then she needs to examine the ball inside and figure out how to open that. There's not one threshold of interactivity; there are many, each nested inside the other.

Superman's X-Ray Vision: Special abilities (particularly always-on special abilities) can cause some items to swap between the different perception type categories for specific characters. This can result in Rachael and Teresa having different perceptual relationships with a given game space.



I'm Just That Good: What if you're really, really, *really* good at spotting stuff? So good, for example, that you might be able to notice the hidden panel in the wall from across the room whereas other characters would need to physically interact with the wall to notice it.

In some ways, you can actually think of this as a variant of the character possessing a particular special ability (it's just that their "special ability" in this case is being really, really good at noticing hidden things).

I actually mechanically instantiated this into 3rd Edition: In my house rules, if you beat the Search DC by +20 while making a Spot check, you'll notice the hidden feature as if you had actively searched for it (either directly, if possible, or through some form of tertiary indication if not; you may note that the latter is effectively introducing a Matryoshka technique). You can do similar stuff with, for example, exceptional successes in [Eclipse Phase](#) or point spends in [Trail of Cthulhu](#).

You might be wondering why this is "okay". Why is this any better than the example of the GM preempting them? Aren't you still skipping interactive steps?

You are, in fact, still "skipping" steps. But you're doing so as a reward for character ability. It's similar to a wizard "skipping" sections of the dungeon by using a *passwall* spell: Yes, you're bypassing the "intended" or "natural" path of progress, and there are things you're losing or missing out on as a result. But you're gaining a different (and important!) benefit.

That's why this is an advanced technique: You need to understand the rule in order to know when you can (and should!) break it.

ASSUMING ACTION

The preemptive Search check, however, is just one specific example of the GM making an *anticipatory ruling*; a ruling in which they assume that the player will make a particular choice and, therefore, skip past the step where the player actually *makes* that choice.

In this context, you can actually interpret the problem as a scene-framing issue. As described in the [Art of Pacing](#), the GM needs to identify empty time – i.e., time in which the player is neither making interesting choices nor experiencing the consequences of those choices – and frame past that empty time to the next meaningful choice. What’s happening here is that the GM is incorrectly skipping past meaningful choices.

The problems with this are manifold:

- It hurts immersion as the player loses control of their character.
- It prevents the player from actually *playing the game* as the loss of control results in a loss of interactivity. In this it’s similar to [alpha-quarterbacking in co-op board games](#).
- It prevents the player from making a different and unanticipated choice. The GM is not omniscient, so even when they assume that there’s only one “good” choice to be made, it doesn’t follow that this is the choice which *will* be made.
- On the other hand, the GM is a little *too* omniscient. They are biased by their design of the encounter and the wider knowledge of the scenario, which may blind them to the actual thought process the player/character is experiencing.

In this, you can see a pattern of problems similar to run-time choose your own adventure (as seen in [GM Don’t List #6](#)).

It is, perhaps, unsurprising that so many errors in GMing technique share common roots. And, conversely, that the solution to those errors are all rooted in a similar ideology.

GM DON'T LIST #8: MYSTERIES WITH NO CLUES

[by Justin Alexander - September 15th, 2018](#)

As I mentioned way back in [GM Don't List #1](#), I am generally anonymizing the examples I give in these essays. I'm not looking to shame specific Game Masters, and I'm not interested in punching down. My goal is not to initiate some sort of witch hunt; it's to educate and discuss.

This essay, however, is an exception. The problems I had this time ultimately originated with a published product, and a published product is fair game.

The book in question is [Strange Revelations](#), an adventure supplement for *The Strange*. When [I reviewed Strange Revelations](#), I mentioned that I would be using at least 8-9 of the ten scenarios in the book at my gaming table.

The tenth scenario was "Venom Rising".

SPOILERS AHEAD FOR VENOM RISING

"Venom Rising" is a hot mess of a scenario.

The premise is that an NPC has been framed for blowing up several industrial facilities and the PCs need to clear her name. So why is everyone convinced she did it? What's the evidence that she's guilty?

There isn't any.

So right there you have a broken premise, right? It's the equivalent of a D&D adventure where you're supposed to kill the ogre that's been harassing the village and, when you ask where he is, the mayor says, "That's his corpse in the corner there."

Which brings us to the scenario itself, which consists of the PCs visiting the suspect's living quarters and each of the industrial sites that were blown up. These scenes play out in one of two ways:

- The PCs show up, fight 4-5 bad guys, search the location, and **find no clues**.
- The PCs show up, search the location, **find no clues**, and *then* fight 4-5 bad guys.

And here's how the scenario premise is supposed to resolve itself:

- The PCs see a picture of the suspect and realize that she has a third arm.
- They do the equivalent of an internet search, find the publicly broadcast video feed from one of the attacks, and note that the masked terrorist who blew up the facility doesn't have three arms.

So when I said that there wasn't any evidence incriminating the suspect, that was inaccurate: It's not just that there's no evidence incriminating her; it's that the evidence which is *publicly available before the PCs begin their investigation* indicates that someone else did the crime.

Except it's even worse than that because the suspect's robotic third arm is pictured as *detachable*. Which means that the evidence which is supposed to be "exculpatory" isn't actually exculpatory.



Ultimately, the module tells the GM to resolve the situation by having the person actually responsible for the bombings simply show up and do something egregiously guilty-looking directly in front of the PCs for no particular reason.

Simply reading through the scenario it was obvious how it would play out at the table: The PCs would go from one location to the next. They would never be able to accomplish anything (because there was nothing to accomplish), and would thus grow slowly more and more aggravated with their frustrated investigations. Then the bad guy would arbitrarily show up and the scenario would end with a final, relieved whimper.

In short, the scenario was completely unsalvageable. I metaphorically tossed it to one side, focused my attention elsewhere, and basically never thought about it again.

... which proved to be a mistake. Because when Gen Con rolled around last month, I wanted to play *The Strange*. Because I had literally forgotten all about "Venom Rising", I didn't recognize the title and, therefore, didn't realize that Monte Cook Games was running a scenario from *Strange Revelations*. In fact, I was over an hour into the session before I realized what was happening.

And the scenario did, in fact, play out *exactly* the way I thought it would: In frustrating irrelevance with an unsatisfying whimper.

(Although to the credit of both creative play and the actual GM who was running the session, it wrapped up in a memorable fashion with the PCs ultimately sympathizing with the "bad" guy's political manifesto and making arrangements to hook him up with freedom fighters who could help him achieve his political goals. As one of the players put it: "It was the first time in a decade of roleplaying that convincing the villain to use better PR has solved the problem!")



TRIAGE AT THE TABLE

Such a result (usually minus the memorable moment) is, in fact, inevitable when the GM creates a mystery with no clues.

You wouldn't expect this to be a particularly common phenomena, but I've encountered it with a surprising regularity over the thirty years I've been gaming. It usually doesn't manifest itself with the utter purity demonstrated by "Venom Rising", but it often crops up in less severe forms or within certain sections of a large scenario.

I suspect the problem is born from a fundamentally inverted understanding of how mystery scenarios are structured: We think of mystery stories as being defined by the *lack* of information (because they are stories of finding something which is unknown), which erroneously leads us to design scenarios which *lack information*.

The primary solution, of course, is understanding that mystery scenarios are actually structured on *acquiring information* and liberally applying the Three Clue Rule during scenario design. That's a topic, of course, which is thoroughly covered in the [Three Clue Rule](#) essay, and I don't think I need to further belabor the subject here.

What I will note, however, is that sometimes you can unintentionally find yourself in this scenario during actual play: You prepped clues, but the PCs missed them. Or you simply made a mistake and the PCs have ended up in an investigatory cul-de-sac where they're unable to make any progress.

Triaging this situation at the table involves accurately maintaining and monitoring your [Revelation List](#):

- Make a list of all the conclusions the PCs need to make in the course of their investigation.
- List all of the clues that exist for each of those conclusions.
- Use the list during actual play to track which clues the PCs have received and – importantly! – which clues they appear to have missed.

(In some cases they'll double back and find clues they initially seemed to miss, but you can't count on that happening.)

What you need to be keeping an eye out for are the revelations for which the PCs have missed all of the necessary clues. (In practice, you really want to start paying attention at the point where they've missed all but one of the possible clues. Don't let yourself get completely backed into a corner before working to get yourself out.)

Some revelations may be nonessential. (Redundancy, particularly once you start working with full-fledged [node-based design](#), is not unusual.) Regardless, a missing revelation should be like an alarm bell going off behind your screen. Or maybe a Check Engine light would be the better analogy: The scenario may not be immediately bursting into flames and crashing, but you do need to get it into the shop ASAP for some tender love and care.

One of the best tools to have in your arsenal is a proactive element in the scenario which you can use to deliver new clues. (More discussion of this in the *Three Clue Rule*, but you can never go too far wrong having some thugs kick down the PCs' door.)

If everything has gone completely to hell, this proactive element can – as in “Venom Rising” – simply be the bad guy showing up for his final confrontation. This technique, however, is *incredibly* dissatisfying for the players. You're better off having the proactive element be at least one step removed so that the players can actually make the final conclusion for themselves (this is similar to [Matryoshka search techniques](#)).

On the other hand, you should also consider just letting the investigation *fail* sometimes. As GMs we're often terribly hesitant to do so, but if player choices are going to be meaningful (and they should be), then sometimes the consequences of those choices will be failure. Clinging to that moment – forcing the PCs to wallow in the failure – can often be more frustrating and harmful than simply allowing the failure to happen: The cultists complete their ritual and Innsmouth slides in the sea. The bank robbers escape with the cash. Doctor Nefaria successfully travels back in time and takes over Yugoslavia in 1982.

The good news is that the consequences of failure usually beget even more interest than success: The disappearance of Innsmouth bring Mythos cults into the national headlines, leading to a burst in youth Cthulhu gangs and Nyarlathotep-inspired rebellion. The bank robbers use the cash to fund an arms deal with Mexican drug lords. The once two-bit villain Doctor Nefaria has suddenly become your campaigns' Doctor Doom and the PCs have no one to blame but themselves.

Of course, they can only blame themselves (instead of blaming you) if you've designed the original mystery *with the necessary clues*.

GM DON'T LIST #9: FUDGING

[by Justin Alexander - September 28th, 2019](#)



No.

Bad GM.

No cookie.

Okay, we've been talking about things GM's shouldn't do for awhile now. So let's talk about the elephant in the room: Fudging.

The most common form of fudging, and that from which the technique takes its name, is changing the outcome of a die roll: You *fudge* the result. If the die roll is done in secret, then you can just ignore it. If it's done in the open, then you can invert the result by tweaking the modifiers involved. More advanced fudging methods can include stuff like adding extra hit points to a monster's total in order to keep them alive.

But, regardless of the specifics, fudging is when a mechanical resolution tells you one thing and the GM chooses to ignore the rules and declare a different outcome.

JUSTIFICATIONS FOR FUDGING

Okay, let's talk about the reasons GMs do this. All of these, of course, ultimately boil down to the GM not liking something that the resolution mechanics are telling them. The question is *why* the GM is unhappy with it.

#1 - Railroading. This one is pretty straightforward: Railroading happens when the GM negates a player's choice in order to enforce a preconceived outcome. Enforcing failure (so that the PC can't do what the player wants) is a really common way of railroading the game, and fudging is a really easy way to enforce failure.

See [The Railroading Manifesto](#) for a lengthy discussion of this topic and all the reasons why railroading is terrible and you should never do it.

#2 - To prevent a player character's death. Or, in some cases, GMs will only fudge if it's to prevent a total party kill – the death of ALL player characters. TPKs tend to kill campaigns (at least those not built around [open tables](#)), and lots of people would prefer to fudge the outcome of a fight (particularly if they feel that it's just due to "bad luck" or whatever).

See [The TPK Gamble](#) for a specific discussion of this.

#3 - To make the story "better." The most infamous version of this is, "But they can't kill the Big Bad Guy now! He's supposed to survive to Act III!"

I say *infamous* for good reason here: Players hate this shit with the fiery passion of a thousand burning suns. And you basically can't throw a stone in RPG circles without hitting someone who has a story about the time their GM pissed them off by doing it. Check out [The Principles of RPG Villainy](#) for a better alternative.

#4 - To correct a mistake you've made. Maybe you've been screwing up a mechanic for the whole fight and it's made things much harder for the PCs than it should have been. Or you accidentally doubled the number of guards when the fight started. Or, going even further back, maybe you just screwed up the encounter design and something that should have been easy for the PCs is actually incredibly difficult. So you fudge something to bring it back in line with what it was supposed to be or should have been.

This is actually pretty understandable, and I discuss the difference between openly retconning a mistake and silently retconning a mistake in [Whoops, Forgot the Wolf](#). But you can easily find yourself slipping from "fixing a screw-up" to "enforcing a preconceived outcome" here and end up back in railroading. So use caution.

DON'T FUDGE

In the end, all fudging is the GM overriding a mechanical outcome and creating a different outcome which they believe to be preferable (for whatever reason).

Over the thirty years I've been doing this, however, I've learned that many of the most memorable experiences at the table are the result of the dice taking you places that you never could have anticipated going. **Fudging kills those experiences.**

But what if the mechanical outcome really is terrible and would make both you and your players miserable?

If you and/or your players truly can't live with the outcome of a dice roll, **then you made a mistake by rolling the dice in the first place.** You need to [focus on fixing that problem](#).

This applies beyond individual dice rolls, too. If you don't want the PCs to die, for example, why are you framing scenes in which death is what's at stake? (This is a rhetorical question: GMs do this because D&D teaches them to (a) frame lots of combat scenes and (b) make the default stakes of any combat scene death.)

[The Art of Pacing](#) talks about the scene's agenda being the question which the scene is designed to answer. (For example, "Can Donna convince Danny to go into rehab?") If the question is, "Will the PCs die?" and the answer is always, "Absolutely not." then the scene is drained of meaning and becomes a boring exercise.

This is why, when the players figure out that the GM is fudging (and they will), it deflates tension and robs them of a legitimate sense of accomplishment. What was once meaningful is suddenly revealed to be meaningless. And this is the biggest problem with fudging: **It may fix an immediate problem, but it will inflict permanent damage on everything.**

In a very real sense, fudging is a betrayal of trust. And once you, as the GM, lose the players' trust, it becomes virtually impossible to regain it. Fudging ends up tainting everything you do: It removes the real magic of an RPG campaign and turns it into a cheap magic trick. Once the players spot the trick (and, again, *they will*), the magic vanishes entirely and you're left with a hollow experience.

Regaining their trust and making them believe in the magic again is really difficult.

TRIAGE AT THE TABLE



Here's my controversial rule of thumb:

*The more you fudge, the shittier you are as a GM – either because you **are** fudging or because you **need** to.*

If you're not just fudging to be an asshole and screw over your players, then you're ultimately fudging in order to fix something that has gone wrong:

- You adjudicated the resolution poorly.
- You designed the scenario badly.
- You screwed something up and need to correct it.
- You're using a set of rules which creates results you and/or your players aren't happy with.

And so forth.

This is not to say that you should never fudge. Mistakes happen and we don't need to live with those mistakes in the pursuit of some unrealistic ideal. But every time you *do* fudge, you should view that as a failure and try to figure out how you can fix the underlying problem instead of just continuing to suck in perpetuity:

- Don't roll the dice if you can't live with the outcome. (And, ideally, learn how to still create *meaningful* stakes instead of just skipping the resolution entirely.)
- Figure out how to design robust scenarios that don't break while you're running them.
- Create house rules to permanently fix mechanics that are creating undesired results. Or, if the system is completely out of line with what you and your players want, swap to a different system.

And so forth.

Next time you find yourself in a position during the game where you feel it's necessary to fudge, I want you to do a couple of things.

First, ask yourself: **Is it truly necessary to fudge in this moment?** Is it necessary to reject the improvisation prompt of the mechanical resolution's outcome, or can you find a way to work *with* that outcome to create something interesting and enjoyable? At the stage in the resolution process where you're [narrating outcome](#), you usually still have a lot of power as the GM. An easy example of this is failing forward: Instead of the PC failing in what they wanted to do, they succeed with a negative twist or consequence.

But also, to a certain extent, just take a moment to second guess yourself: The outcome which you initially think *cannot possibly happen*, often *can* happen. It's just not what you expected or would have done of your own volition. Try to push back that initial moment of rejection and really, truly think about what the outcome would be and whether there's interesting and cool stuff that lies beyond that outcome.

Second, ask yourself: **Can I just be open and honest with my players in this moment?** Instead of secretly fudging the outcome, could you just explain to the players that, for example, you screwed up the encounter and things need to be retconned a bit?

And maybe you can't! There are circumstances where you're better off plastering over the cracks of your mistake with a cheap magic trick instead of damaging the players' immediate immersion and engagement with the game world. It's not ideal, but sometimes that's the best you can do for right now. You'll just have to learn from your mistakes and do better next time.

CODA

If you're still a proponent of fudging, let me ask you a final question: Would you be okay with your players fudging their die rolls and stats and hit point totals?

If not, why not?

If you truly believe that fudging is necessary in order for you to preserve the enjoyment of the entire table, why do you feel you know better than the other people at the table what they would enjoy?

Think about it.

THE FUDGING COROLLARY: NOT ALL DICE ROLLS ARE MECHANICS

[by Justin Alexander - September 30th, 2019](#)



In [GM Don't List #9: Fudging](#) I discuss why GMs should avoid fudging, and if they do end up needing to fudge, why they should view that as a failure point in their game and a learning opportunity to figure out how they can become a better GM.

The most common form of fudging is changing the outcome of a die roll (the roll was a failure, but you say it was a success, or vice versa), but it's widely understood that there are also other ways to fudge mechanical results (increasing a creature's hit point total, for example).

What I think is less commonly understood, however, is that not all the dice rolls you make in a roleplaying game are mechanics, and it's not actually fudging to change or ignore those dice rolls. Specifically, procedural content generators. Such generators can actually use any number of randomization techniques (for example, [here's a method for using CCG cards to generate adventures](#)), but since we've already got dice laying around the typical RPG table most procedural content generators just use those.

WHY ISN'T THAT FUDGING?

If you're struggling to understand why changing the outcomes of a procedural content generator isn't the same thing as fudging a mechanical resolution, let's take an extreme example. I'm prepping a scenario for my next session and I need a name for an NPC. So I pop open the [Random Name Generator](#) at Behind the Name, select for random surnames, click the button, and get:

Ivonne Eógan Masson

For whatever reason (maybe personal aesthetic, maybe because the Masons are already established as major power brokers in the city and I think it's interesting this random generator has unexpectedly

connected this NPC to the clan), I decide to drop the second “s” from “Masson” and name the character Ivonne Eógan Mason.

Did I just fudge?

Frankly speaking, no. Not by any reasonable/functional definition of the term.

What if instead of tweaking the outcome I actually just ignored it and rolled again by hitting the “Generate a Name!” button a second time? Still no.

What if I move this interaction from prep to actual play (I need to come up with a new NPC’s name on the fly, so I randomly generate one and then tweak it)? Still no.

What if the random name generator is published in the game’s rulebook? Still no.

This isn’t fudging not only because it would make the concept of “fudging” so broad as to be meaningless, but also because treating the outcome of a procedural content generator as a straitjacket or legally binding contract is to fundamentally misuse the procedural content generator. Using a procedural content generator is more like coating the bottom of an agar plate with a growth medium: As its exposed to your creative subconscious, the growth plate begins to accumulate a bunch of random creativity and odd synchronicities that begin to grow and thrive. (Ivonne Mason, for example, is a very different character than Lea Colton or Caroline Bone specifically *because* each of those random names provides a different creative stimulus.) Treating the outcome of the procedural content generator as if it were inviolable scripture, on the other hand, is like sterilizing the agar plate; it completely short-circuits the process.

ALL MECHANICS ARE PROCEDURAL CONTENT GENERATORS!

“Ah ha!” you say. “But aren’t *all* resolution mechanics actually procedural content generators, the results of which are meant to be creatively interpreted by the GM and other players? Is not the [narration of outcome](#) the same thing as taking a randomly generated group of bandits and creating the [Blood Shield Bandits](#)?”

Basically, no. There’s a similarity of process (roll dice, interpret results), but the function of resolution mechanics and procedural content generators are in many ways actually inverted: A resolution mechanic takes generally non-specific creative input and creates specificity (often literally a binary pass/fail state). A procedural content generator, on the other hand, *produces* non-specific creative input and expects the GM to create the specificity.

Because the processes involved are similar and because “specificity” can be a sliding scale, you *can* use procedural content generators as resolution mechanics (sterilizing your agar plates) and vice versa. But because the tools are designed for one thing and you’re using them for something else, the result is usually like using a screwdriver as a hammer.

GRAY AREAS

This is not, however, to say that there are no gray areas which lie along the boundary between resolution mechanics and procedural content generators.

You can see this perhaps most clearly when a game takes something which is traditionally a procedural content generator in other systems and makes it a hard-coded mechanic (or closer to being a mechanic) instead.

For example, in *Apocalypse World* characters are created from playbooks. For example, if you want to play a Gunlugger, you take the Gunlugger playbook and it instructs you, "To create your gunlugger, choose name, look, stats, moves, gear, and Hx." Each of those categories then has a specific list of things. This becomes a surprising gray area: Many people, conditioned by other RPG character creation systems, looked at the provided list of names as a resource that could be used or ignored. (Many editions of D&D have similar lists of elven names, for example. *Over the Edge* provides a list of AI Amarjan names. And so forth.) *Apocalypse World*, however, specifically seeks to enforce setting through non-traditional mechanics, and so I've played at tables that instead interpreted this as a mechanical requirement: You must choose your name from the provided list of Gunlugger names (which are distinct from the list of, for example, Hardholder names, thus asserting setting). I'm actually still unsure what D. Vincent Baker's intention was.

You can find another gray area in the dungeon stocking procedures of the original edition of *Dungeons & Dragons*. While the rules allow for some discretion on the part of the GM in the distribution of treasure and monsters, the line between advice, procedural content generator, and actual mechanic is very fuzzy and open to a lot of interpretation. This is even more true when it comes to wandering monsters in OD&D: A random encounter check is often interpreted as a procedural content generator by modern GMs, but in OD&D it's stated as a straight-up mechanic.

The [*Mythic Game Master Emulator*](#) is also an interesting example: In order to emulate the role of the GM, the *Mythic* system basically adds a lot of binding structure onto a suite of procedural content generators. But despite the gray area this creates, the system still draws a fairly strong distinction between the output of the resolution mechanics and the output of the GM emulator (i.e., the procedural content generators). If you want to get a really clear feeling for how using resolution mechanics and using procedural content generators differ from each other, spending an afternoon playing around with *Mythic* as a cap-system for your favorite RPG can be very illuminating.

GM DON'T LIST #10: IDEA ROLLS

by Justin Alexander – May 28th, 2020



Here's my random tip for using Idea rolls as a GM:

Don't.

Let me start by explaining what I'm talking about: In *Call of Cthulhu*, an Idea roll "represents hunches and the ability to interpret the obvious." In some of the older scenarios published for the game, this roll would actually be used to prevent players from having their characters take certain courses of action because the character wouldn't know to do them – sort of aggressively preventing [player expertise from trumping character expertise](#).

There are some obvious problems with that, too, but what I'm interested in right now is the far more common technique of using the Idea roll to tell players what they "should" be doing. For example, if the players are talking about how they can get an audience with a casino owner, the GM might call for an Idea roll and say, "You could disguise yourselves as high rollers." Or when the PCs stumble onto a bloodstained altar in the center of a stone circle, the GM might call for an Idea roll and then say, "You could try putting that idol you found earlier on the altar!"

Even in games that lack a specific mechanic like this, you may see similar techniques improvised (usually with some form of Intelligence check).

GM-INITIATED IDEA ROLLS

The basic function of the Idea roll is essentially like using a walkthrough in a video game: You don't know what to do, so you have to consult a guide that can get you past the point where you're stuck. A GM-initiated Idea roll, though, is often more like having an obnoxious friend witting with you who's played the game before and simply WILL NOT shut up and let you play the game for yourself.

If you're a GM prepping a scenario and you come to a place where you think an Idea roll will be necessary, that's a really clear sign that you need to DO BETTER. Saying, "I need an Idea roll here," is basically saying, "I have designed a scenario where the players are going to get stuck here." Instead of prepping an Idea roll, figure out some way to redesign the scenario so that the players *won't* get stuck there. (The [Three Clue Rule](#) will often help.)

What about run-time Idea rolls? In other words, you're currently running the session, you can see that the players are irreparably stuck, and you need to fix the problem. Well, there are two possibilities:

First, they're not actually stuck, in which case you don't need to use an Idea roll.

Second, they ARE stuck and definitely need help to get unstuck. In which case, you shouldn't be rolling the dice because failure is not actually an option: You *need* to give them information. Therefore you should not be rolling to see whether or not they get it.

PLAYER-INITIATED IDEA ROLLS

On the other side of the screen, a player-initiated Idea roll is generally more viable: This is basically the players sending up an emergency flare and saying, "We're lost! Please send help!" To return to our analogy of the video game walkthrough, this is the player who has been stymied to the point where they're no longer having fun and just want to be able to move on in the game.

In my experience, it should be noted, what such players are looking for is often not the *solution*; what they are looking for is an *action*. They feel stuck because they don't know what they should be *doing*. A [Matryoshka search technique](#), therefore, is often a great way to respond to this.

Something else to look for is the clue that they've overlooked. Not necessarily a clue they haven't *found*, but one which they don't realize is actually a clue, which they've radically misinterpreted, or which they've completely forgotten they have. For example:

- "You realize that *patent leather* can also be used for furniture, not just shoes."
- "While S.O.S. could be a cry for help, couldn't it also be someone's initials?"
- "You suddenly remember that you still have Suzy's diary in the pocket your trench coat. Didn't she mention something about the color purple, too?"

[Trail of Cthulhu](#) innovated a cool mechanic along these lines for its Cthulhu Mythos skill: You can use this skill to "put together the pieces and draw upon the terrible knowledge that you have been subconsciously suppressing, achieving a horrific epiphany. The Keeper provides you with the result of your intuition, sketching out the Mythos implications of the events you have uncovered."

There are two important features to this mechanic: First, it doesn't require a roll. (Again, if the players *need* help, then denying it to them on the basis of a dice roll doesn't make sense.)

Second, it has a cost: The sudden insight into the terrible realities of the universe will cost you Stability and, quite possibly, Sanity. Importantly, this cost is NOT exacted "if the player deduces the horrible truth without actually using [the] Cthulhu Mythos ability." The cost, in my experience, not only dissuades players from relying on the mechanic instead of their own ingenuity, it also enhances the sense of accomplishment they feel when they solve the mystery or gain the insight without using the mechanic.

The 7th Edition of *Call of Cthulhu* has similarly modernized the Idea roll, using a [fail forward](#) technique where a failure still get the PCs the necessary clue/course of action, but also results in some sort of negative consequence: Getting the clue might bring you to the attention of the bad guys; or you might waste weeks of time digging through a library before finally stumbling across the right reference; or, like *Trail of Cthulhu*, the insight might force a Sanity check.

Another cool technique it suggests, particularly in the case of failing forward, is to aggressively reframe the scene: Jump directly to the point where the PCs have followed the lead and gotten themselves into trouble as a result.

A final interesting variant here is to make the Idea roll concept diegetic instead of non-diegetic; i.e., to make it a decision the *character* makes instead of the player. In a fantasy setting, for example, the character might literally make a sacrifice to the Goddess of Knowledge in order to receive a divine vision.

GM DON'T #10.1: TELLING THE PLAYERS THE PLAN

Like an aggressive Idea roll on steroids, some GMs will go so far as to just literally tell the players what their characters will be doing for the entire scenario.

For example, I was playing in a convention one-shot where we were street samurai who got hired to be ringers on a Blood Bowl team in order to rig a high-stakes game. This was a really cool premise, turning the usual expectations of the game on its head and giving us an opportunity to explore how the PCs' heist-oriented abilities could be used in a completely novel environment.

Unfortunately, the session quickly went completely off the rails. Rather than letting the players make any meaningful decisions, the GM had pre-scripted every play of the game: We were reduced to simply rolling whatever skill had been scripted for us. (It didn't help that the rolls themselves were essentially pointless since the outcome of every drive and most of the plays had ALSO been planned ahead of time.)

This was an extreme example of something closely related to [GM Don't List #7: Preempting Investigation](#), but I bring it up here mostly because I've seen several GMs who use Idea rolls to similar (albeit usually less absurd) ends. These game are characterized by the players making an endless stream of Idea rolls, with the GM constantly saying things like, "Pierre [your character] thinks he should come back and check out the Le Petit Pont after dark." Or, "You could probably get a pretty good view from the top of Notre Dame. You'll need to figure out some way to get up to the top of the towers." Or even, as literally happened in one game, "Rebecca thinks she should stab the Archbishop in the chest." ("No, really, she thinks this is *really important*.")

Basically: Don't do this. Present your players with problems, not solutions. Give them the space to mull over a situation and figure out what they want to do (or what they think they *need* to do) in response to that situation.