Random GM Tip: Making Clues

by Justin Alexander – May 29th, 2021

In the *Three Clue Rule*, I assert that for conclusion you want the PCs to make in a mystery scenario, you should include at least three clues. Many people have had a lot of success designing and running scenarios with this advice, but as GMs begin working with this technique for the first time, it can often prompt the question:

How do you come up with all those clues?

So let's talk a little bit about what goes into making clues and a few of things I've learned over the years that you may find useful.

First, check out *Random GM Tip: Using Revelation Lists*. That article provides an in-depth look at how I organize my clue lists while designing and running scenarios, and that structure is going to be useful as you start laying the clues into your scenario.

Second, don't overthink it. A lot of clues will flow naturally from the scenario as you're designing it. I recommend doing an initial "structural pass" in which you're defining the necessary revelations and assign clues to them. If you can get three clues for every revelation at this point, that's fantastic. But if you're struggling to fill in all the blanks on your revelation list, that's just fine: Start working out the details of the scenario with your revelation list nearby and you'll often find clues easily along the way. ("Oh! Silvester would definitely know that Uncle Bob used to be in the marine corps!" or "Hey! I could put a matchbook from the Silver Rodeo in Susan's purse!") And if that still doesn't fill in all the clues you need, all this detailed knowledge of the scenario will make it a lot easier to come up with the clues you still need during a second brainstorming session.

Six Types of Clues

Conceptually, there are six types of clues.

Static clues are a specific piece of information which can be found in a specific way. For example, a bloodstain that can be found by searching a room.

Flexible clues are pieces of information that you know are there to be found, but which can be accessed in multiple ways. For example, imagine that there's valuable information in a computer database. The PCs could hack the database remotely, physically raid the building where the server is located; bribe, seduce, con, or otherwise subvert an employee with access to the database; and so forth.

Broadly speaking, PCs will learn about people, places, organizations, events, or other scenes where they can investigate whatever it is that they're investigating. In the terminology of node-based scenario design, these are nodes, and static and dynamic clues are placed within specific nodes.

Proactive clues, by contrast, come looking for the PCs. I talk about them at some length in the *Three Clue Rule* and also *The Secret Life of Nodes*, but Raymond Chandler provides the classic example: "A guy with a gun walks through the door." In other words, instead of the PCs identifying a node and going to investigate it, some nodes (usually people) will come to the PCs and bring clues with them.

Reactive clues, on the other hand, kind of exist "in the cloud," so to speak. They aren't assigned to specific nodes and they don't come to the PCs. Instead, these are clues the PCs will find through holistic investigation techniques that they can frequently use without specific prompting. Typical forms include canvassing, research, and divination spells. *Rulings in Practice: Gather Information* dives into these types of clues in greater depth.

(This terminology can be a little weird. You can think of it like this: If the clues are proactive, then the PCs are reactive – i.e., the clue walks through the door with a gun and the PCs have to react to that. If the clues are reactive, then the PCs have to be proactive – nothing is going to specifically prompt them to go looking for these clues.)

Because reactive clues are more-or-less totally dependent on the players spontaneously thinking to go looking for them, they can be very unreliable when designing an adventure. On the other hand, if looking for a particular type of reactive clue becomes a standard operating procedure for a particular group, then the exact opposite can be true!

For example, over the past couple of years I've spent a lot of time running the same *Call of Cthulhu* and *Trail of Cthulhu* scenarios for different groups. Groups with players who have read the rulebooks or have a lot of experience with the game will almost automatically go to the library or research the local newspaper morgues for references to whatever they're investigating because the rulebooks establish this as a standard operating procedure for investigators. Groups with that experience just… don't.

The *Technoir* roleplaying game is actually mechanically designed around the default action of "hit up one of your contacts." If the PCs don't know what to do next, they should go ask one of their contacts (and the game is designed so that the contact always has a clue or a job or some form of lead that gives them something to pursue). The first few times I ran the game, this worked flawlessly. Then I ran the game for a group that hadn't read the rulebook and didn't know to do this and the game turned into a grind for a couple of sessions until I realized that I just needed to literally tell new the players: If you're stuck, do the noir thing and hit up a contact!

In other cases, you'll find groups spontaneously developing their own standard operating procedures.

Dynamic clues are what the PCs find when they take an investigative action which should logically provide information despite the fact you didn't specifically prepare a clue for it. This is what the *Three*

Clue Rule refers to as permissive clue-finding. For example, the PCs decide to look around outside the house. You didn't specifically anticipate them doing that, but you know that the killer ran into the tree line at the far end of the property, so you conclude that the PCs could potentially find the killer's footprints.

However, not all dynamic clues are the result of you being blindsided during a session. In some cases, they can actually be designed into the structure of a complicated scenario where you can be fairly certain the PCs will take a particular class of action, but you can't really sure exactly what form that action will take. For an example of this, check out the *Dragon Heist Remix*, where, for example, the PCs might choose to research a particular faction:

> *If the PCs want to find a faction by doing general research, point them in the direction of one of the faction's outposts. (Each outpost will contain clues that point to Lairs, which are generally their ultimate goal.)*

And similar guidance is given for what happens when they track or interrogate bad guys. (You could also think of dynamic clues as being flexible clues that are flexible to the point of formlessness.)

Unassigned clues are basically what happens when you go whole-hog on permissive clue-finding: Instead of prepping specific clues, the GM only preps a revelation list. In some cases these revelations will be assigned to specific scenes (i.e., the fact Tony is hiding at a lake house can be found at the Silver Rodeo). In either case, the GM waits for the PCs to propose any investigation action, then looks at their revelation list and improvises an appropriate clue for that revelation.

I tend to be fairly skeptical of this approach:

- It puts A LOT of pressure on the GM's ability to improvise during sessions.
- It often produces wishy-washy scenarios, that tend to lack texture. (A good mystery scenario is often as much about what you don't find as what you do, and this technique tends to miss those beats in the story.)
- As mentioned above, many clues will tend to emerge naturally from the details of the scenario as you design it. So this technique usually goes hand-in-hand with underdeveloped scenarios.

I've run mystery scenarios that were entirely improvised, so the technique certainly can work. But I think you will almost always get better results by having a robust, reliable foundation and then improvising on top of it.

MAKING CLUES: STRUCTURE

When making clues, the first thing I do is think about what the PCs need to know structurally for the adventure to work. This is why I start with a revelation list: There may be a lot of other information that the PCs discover in the course of their investigation (that the villain beats her husband; the particular effects of the poison the killer is using; the cult's beliefs on the wisdom of cats), but I really want to keep my focus on that essential structure.

Then I look at what I know about the crime – or whatever it is that they're investigating – and think very specifically about what in the specific node I'm looking at could indicate the thing they need to know.

This may seem obvious, but I can get lost a surprising amount of the time floating around in a, "What clues are there?" haze. You really want to flip that around: Instead of thinking about what clues might be in this scene, focus on what you *need* the current scene to tell the PCs and then treat that as a puzzle or a problem to solve.

For example, the PCs are investigating a cabin on the lake. If you start by saying, "What clues would be in the cabin?" that's too broad. It's too vague.

But if you instead say, "I need a clue that points the PCs to Cai Lijuan," that's far more actionable:

- There's a crumpled up envelope in the wastebasket addressed to Cai.
- The property owner can identify Cai as the person who was renting the cabin.
- There's a box of cold pizza in the fridge. Cai's name is printed on the delivery label.
- The family vacationing in the next cabin down the lake met Cai and knows his name.
- Cai's car is still parked at the cabin; you can run the license plate or VIN numbers.

And so forth.

Making Clues: The Skill List

For inspiration, look at the skill list in the game you're using. Pick any skill. How could the PCs use that skill to get the information they need?

For example:

- Cryptography? Cai's left his diary, which he writes in a personal code, in the bedside table.
- Locksmith? There must be something Cai locked up here. Let's say there's a safe hidden behind a picture frame with documents identifying Cai inside.
- Photography? There's a USB stick with photos. Nothing identifying in the photos themselves, but you could check the EXIF data.

Obviously not every skill will be relevant to every piece of information. But if you have a particular piece of information in mind, running down the skill list will probably make specific skills jump out and provide ideas.

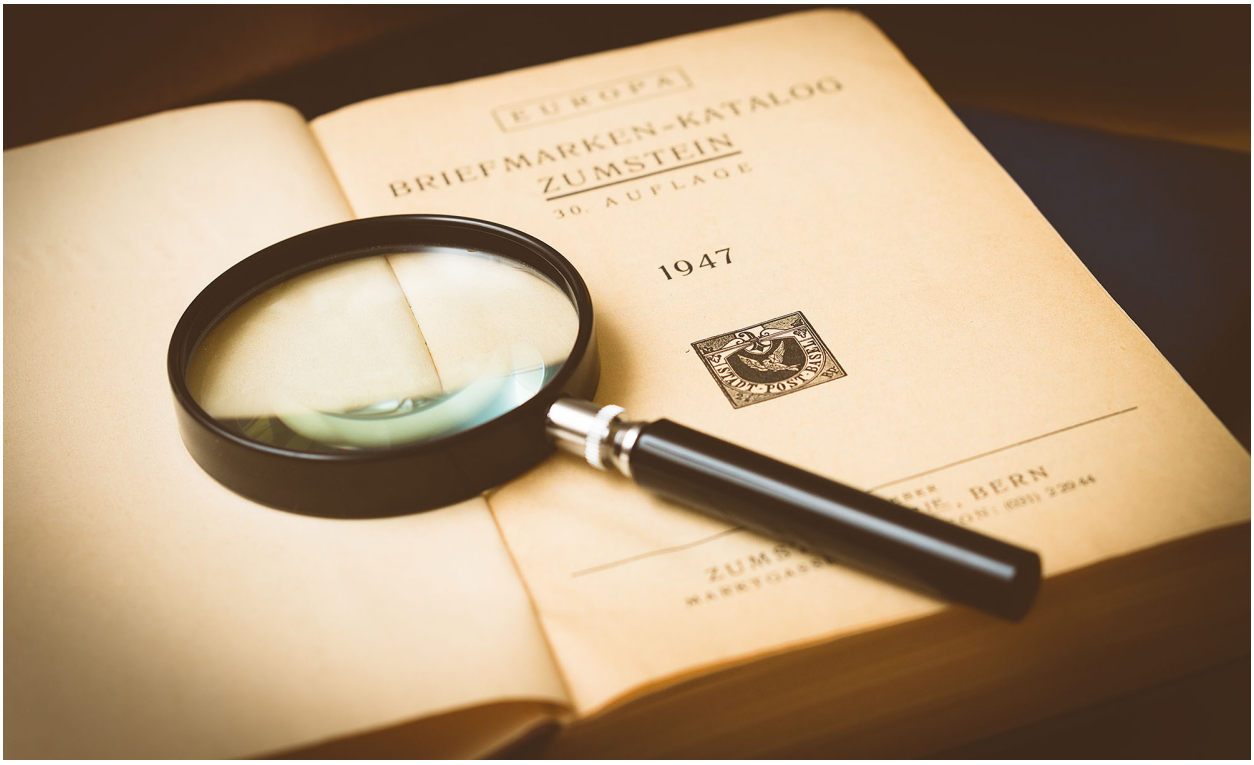

The Fallacy of the Perfect Crime

Something that can really trip you up when trying to create clues is the fallacy of the perfect crime.

For example, let's say the PCs find an incriminating note that gives them a vital clue.

… but wait. Why would the conspiracy let the PCs find the note? Why wouldn't they destroy it?!

Well, in some cases it's carelessness. In others, they're holding onto incriminating evidence that can give them leverage over their fellow conspirators. Sometimes they're planning to destroy it, but they just haven't yet. Sometimes the note has information they need for a job, so they hold onto them until the job is done. Maybe they have a sentimental attachment to the note because it was written by their dead wife. And so forth.

Alternatively, maybe they DID destroy the note and *that's* the evidence that the PCs find. Sometimes this is structurally cosmetic (the only information to be gleaned is that the information has been destroyed). Sometimes this can be a consequence of the players' actions, like in my *Dragon Heist* campaign where they knocked politely on someone's door and then waited for them to open it while, inside, the bad guy was shuffling all their papers into the fire. (One of the nice things about the Three Clue Rule is that you can destroy clues with a clear conscience when it's appropriate.) But often the PCs will be able to pull a clue out of the destruction (they burned the note, but this one enigmatic scrap remains!).

Of course, this logic also extends beyond incriminating notes.

Wouldn't it make sense for the murderer to wear gloves or wipe away their fingerprints? (Maybe they did and the clue is something else. Or they tried, but missed a partial. Or they're dumb. Or they panicked. Or they wanted to, but got interrupted.)

Wouldn't they have disabled the cameras? (Maybe they didn't spot the camera. Maybe they *did* hack the cameras and deleted the footage, but left traces from their hack that can be traced instead. And so forth.)

The perfect crime (or conspiracy), of course, would have flawless operational security and leave no evidence. But that's often not how it plays out in reality, and usually doesn't make a great scenario in any case.

As GMs designing scenarios, I think we're particularly susceptible to the fallacy of the perfect crime: Obviously if the bad guy knew they were leaving a clue, they would destroy it. You just thought of how the criminal could have left a clue, so obviously the criminal would have realized that, too!

What you actually want to do is the opposite: If you're thinking about the crime and can't find a clue that gives the PCs the information you need to give them… well, what extra mistake did the criminal make that will let you give them that information?

Varied Clues

As you're designing clues for your scenario, you'll want to make sure to include a wide variety of them. This is partly about creating a more engaging investigation. (If the PCs are just doing the same thing over and over and over again, that's just less interesting than an adventure where they're doing a lot of different things. And a puzzle isn't really puzzling if the solution is always the same thing.) But it's also structurally important: If all the clues are fundamentally similar, then it's not just that the players are doing the same thing over and over again; it's that the players HAVE to do that thing. And if they don't think to do it, then they'll miss all the clues.

The Three Clue Rule is built on redundancy, but clues which are overly similar to each other only provide a superficial redundancy. It's kind of like monoclonal agriculture: When all the bananas are clones of each other, they're all susceptible to the same pests and can be universally wiped out by a single disease. Just so with monoclonal clues, which can all fail simultaneously.

For example, if the PCs need to find out that Tony is hiding out at the Silver Rodeo, you might say:

- They can ask Tony's girlfriend, Susan.
- They can ask Tony's partner, Silvester.
- They can ask Tony's mother, Sara.

Those are three different clues. But if the players, for whatever reason, don't ask the people in Tony's life where Tony is hiding out — because they don't want to risk tipping him off, because they erroneously conclude that Tony wouldn't have told anyone where he was going, or just because they don't think of doing it — then that one failure will wipe out all of those clues.

Now, the principle of permissive clue-finding means that shouldn't necessarily get rid of these "redundant" clues. But for structural purposes, they can be grouped together and only "count" as a single clue for the purposes of the Three Clue Rule.

When you're looking to make varied clues, using clues with different forms is good. But the most important thing is that the clues should be found in different ways – different skills, different insights into how a crime scene should be investigated, and so forth.

As a final note, remember that the problem of monoclonal clues is limited to the clues pointing to a single revelation. It's fine to design a scenario with lots and lots and lots of clues coming from talking to people, as long as those clues are spread out across a bunch of different revelations (each of which has varied clues of different types also pointing to them).

MULTI-PART CLUES

A common error that I see GMs make when playing around with the Three Clue Rule for the first time is to mistake the Three Clue Rule for a kind of logic puzzle:

- This clue indicates that the killer was wearing a green sweater.
- This clue indicates that the killer was taller than six feet.
- This clue indicates that the killer had gray hair.

If you combine those clues, you know that the only person with gray hair who was taller than six feet and also wearing a green sweater was Peter! So Peter killer Tony's wife!

These clues can work if each uniquely points at Peter: He was the only one in the green sweater, the only one taller than six feet, and the only one with gray hair.

But if you need all three pieces of information – to eliminate the other people with green sweaters or gray hair or whatever – then this falls apart. Because you NEED all three pieces of information, that means that each piece of information (green sweater, six feet, gray hair) is actually a separate conclusion.

This doesn't mean that you can't design mysteries like this, though! Once you recognize that these are three separate conclusions, you can simply follow the Three Clue Rule: Have three clues pointing to the green sweater, three clues pointing to the killer's height, and three clues indicating the killer had gray hair.

This technique can add a satisfying dimensionality to your mystery scenarios, giving the players a clear sense that their investigations are building towards some central revelation. It can be particularly

effective for the Big Truth(s) in *X-Files*-type campaigns or identifying the location where the big conclusion of the campaign is going to take place. You can see an example of this in my *Eternal Lies Remix*.

FORMS OF CLUES

Let's wrap things up by looking at the specific forms that clues can take. This won't be a complete or encyclopedic coverage of the topic, largely because the cool thing about clues is that they can be virtually infinite in their form and variety, but hopefully I can provide a few ideas.

First, there are some broad categories of form that clues can fall into:

- Physical Artifacts
- Glyphs/data
- Bio-Signature
- Interrogation
- Surveillance

These, too, are not comprehensive. And the boundaries between them aren't exactly razor-sharp.

Here's a list of specific examples drawn from scenarios I've designed:

- Correspondence (letters, e-mail, etc.)
- Diaries
- Effluvium from a location (matchbooks, theatrical posters, tickets, etc.)
- Official reports
- Tracks
- Surveillance (of a person or location)
- Tailing someone
- Business cards
- Fingerprints
- DNA
- Blood type (including fantastical types like "Vulcan")
- Graffiti
- Financial records
- Tattoos
- Canvassing
- Video/audio recordings
- Mystic visions/strange dreams
- Shipping information (tracking, postmarks, return addresses, etc.)
- Books (including inscription and marginalia)
- Bureaucratic records/background checks