

Notes about the Gradwrite Writing Guide

What is the purpose of this guide? Of course, if you're planning to write a trivia pack (or a set of packs for a tournament), this writing guide contains a number of guidelines to help ensure your pack is fun to play and fair for players. Beyond just writing questions, understanding the principles to which we adhere can improve your performance as a player by providing you with a deeper understanding of how we craft our questions and a deeper understanding of the game as a whole. In publishing this guide—our official internal writing guide—we hope to level the playing field somewhat and provide anyone interested a solid grounding in the theory which informs our packs. Outside of the context of improving as a player, this writing guide can also serve as a decent introduction to the basics of trivia theory, for those who are interested in deepening their understanding of the game.

This writing guide is rather long. If you'd like, you can find an abbreviated version of this guide [here](#). If you're just looking to write a good pack, the abbreviated guide will get you pretty far. The abbreviation is exactly that, though: a summary of the things which are absolutely essential. Regardless of why you might be looking at this writing guide, reading the full version will provide you with a much more complete picture of when and how to apply principles most effectively. We strongly recommend that you read the entirety of the full writing guide if you have the time. Also note that this is a working document, so sections may be added, removed, or edited over time.

Part I of the writing guide concerns basic principles to which all of your questions should adhere. Part II deals with principles pertaining to particular question types. Later parts of the guide presuppose knowledge of the earlier parts.

All of the Guiding Principles in this guide are worded as commands. Outside of Gradwrite, we have no place telling you what to do and what not to do with your packs, so feel free to experiment! Almost all of these rules have exceptions. For what it's worth, though, we at Gradwrite and people at other trivia organizations have developed these principles over the course of years, and they shape the packs which you are accustomed to playing. Also, strict as the Guiding Principles may seem on paper, in practice there's always plenty of room for writers to stamp their own distinctive style onto their packs.

Finally, we'd like to make it clear that this writing guide is nowhere close to comprehensive. If you're interested, don't stop here! Efforts to elucidate exactly what makes a good pack have spawned a whole world of theory beyond what's contained in this writing guide.

We hope you find this writing guide helpful!

Gradwrite Trivia

PART I:

GENERAL GUIDELINES

Basic Conventions

There are a few basic rules that your questions should almost always follow. We'll just list these out for you:

Guiding Principle: It is better not to include a question than to include a bad question. If you can't figure out a way to get a question to conform to the principles in this guide, consider scrapping it. Don't delete your questions, though! Often, returning to old questions or getting others' opinions can save a question which might have initially seemed impossible to write well.

Why do we do this? It is better to have fewer packs (but packs that play well) than to have more packs which mostly play well, but have a few unfair questions throughout.

Guiding Principle: Questions should almost always be phrased as a question, not a statement, unless you've written a special category and explicitly explained its format to players. Even spelling questions should be phrased as a question (for example, "How would you spell the name of the best-selling brand of Japanese mayonnaise, 'kewpie'?"). The one major exception is for questions which ask for a subset of the possible answers (e.g., "name two of the three..."); you may word these as statements if necessary. Questions should be grammatically correct sentences, unless that has been explicitly explained.

Why do we do this? Phrasing questions as questions helps provide consistency to the game and clarity to players. Questions that refer to "this author" are usually harder to parse than questions that ask for "what author", for example. The effort required to reword a statement into a question is typically very small and always worth the added clarity for players.

Guiding Principle: Questions should always be one sentence, unless you've explicitly explained that format to players. This also means that questions should be short enough that players aren't struggling to parse the question, which often happens when writers try to stuff too many clues into a single question. In general, varying your question length is a good idea!

Why do we do this? The fast-paced style of trivia that Gradwrite writes requires the game to have something of a rhythm and proceed at a certain pace. Excessively long questions can interfere with the flow of the game. The game is most fun when there are a mix of long and short questions!

More importantly, though, players expect questions to be one sentence long. Generally speaking, we at Gradwrite try to avoid subverting players' expectations in ways that punish good habits. Being able to buzz in right at the end of a sentence is a skill worth cultivating, but including a second sentence after the first—a second sentence which likely reveals more useful clues—punishes players who have developed that skill through extensive practice.

Also, as with the previous principle, the effort required to rephrase a multi-sentence question into a single sentence is relatively small. If it really seems impossible to ask a particular question in one sentence, that may be a sign that the writer is trying to stuff too many clues into a single question.

Guiding Principle: Questions should *not* explicitly list out the set of potential answers like a multiple choice question on a quiz. This means no true or false questions. Questions which ask for a member of a small set (e.g., “Which Canadian territory...”) are fine.

Why do we do this? Multiple choice questions are aesthetically displeasing, and they require players to wait until a list of things is read out, so we avoid them when we can. They can also introduce an unwanted element of luck: a less-knowledgeable team can take a chance by buzzing before the entire list is read

Fortunately, there’s almost always a way to turn a multiple choice question into a regular question. Often, simply removing the list of options leaves you with a perfectly serviceable non-multiple-choice question. For a true or false question, you might pick one thing in the question as the answerline and reword the question so that the other parts of the question serve as clues: for example, from “True or false? The names of the three girls Gru recruits to trick Vector in *Despicable Me* are Margo, Edith, and Agnes.” to “What character recruits Margo, Edith, and Agnes to trick Vector in *Despicable Me*?”.

In the cases where it’s impossible (e.g., “Out of iron, helium, or chlorine, what element has the highest atomic number?”), it’s usually because the question clues a distinction which is arbitrary or invented—the fact that iron has the highest atomic number out of iron, helium, and chlorine is a fact of no importance, interest, or value to anyone. We wouldn’t want to include those types of questions for reasons quite independent of the fact that they’re in the multiple choice format.

Guiding Principle: Questions should be clear and precise—we don’t like to use terms like “related to” or “associated with”. Instead, precisely describe what that relation or association is.

Why do we do this? These kinds of descriptions leave a lot of room for interpretation rather than precisely describing what the relation is. It leaves room for other reasonable interpretations of the question (very possibly other correct answers).

As trivia writers, we want to encourage players to learn things. A question like “What two countries are associated with the Hala’ib Triangle and Bir Tawil?” is only interesting to someone who has already heard of the Hala’ib Triangle or Bir Tawil. By more precisely explaining that Egypt and Sudan both lay claim to the Hala’ib Triangle and both refuse to accept Bir Tawil as part of their territory, the question becomes meaningful to someone who might not have heard of either region before. Furthermore, a player who hears an imprecise question will form only a vague association between the names mentioned, if they remember the question at all. An interested player who

hears a precise question learns a real fact, and may even be inspired to learn more about the topic on their own time.

Difficulty

When you're just getting started, the hardest part of writing questions is almost always hitting a good difficulty level, because your perception of how hard a question is can often be quite different from how difficult it actually is for players.

The Golden Rule of Trivia Writing: Your questions are harder than you think they are.

Why is this the case? This mismatch can happen for a couple of reasons. The biggest one is that you tend to write questions about things you know about—naturally, these things feel easier to you than they are to the average player, and in general this leads to questions that are way too hard. Another reason could be exposure to past packs—just because a question hasn't shown up doesn't mean it's too hard, but more importantly, just because it's shown up before doesn't mean that it's easy!

How can writers accurately assess the difficulty of their questions, then? At Gradwrite, we like to think about difficulty in terms of which players are answering the question correctly.

Guiding Principle: When judging the difficulty of your questions, think about how hard the question will be, not how hard it should be. For example, a question about the fundamentals of a well-known, popular field will be easier than a question about the fundamentals of a less well-known field, because the second one will be converted (that is, answered correctly) less often.

Why do we do this? Thinking about how the question will actually play results in a more objective judgment. Even among people who have a good understanding of the general landscape of a particular field, there is frequent disagreement about how fundamental or “important” a particular fact is within that field—for example, it is extremely natural to feel that since you know something, other people who have similar interests “should know” the same thing. As a result, different writers and editors frequently have wildly different perceptions of how difficult a particular question is. Of course, thinking about conversion is far from a flawless way of gauging difficulty, and writers' biases still show through, but it's served us much better than other ways.

Judging difficulty by conversion also better aligns with what we're actually trying to achieve by thinking about difficulty in the first place. When compiling packs, we try to space out hard questions and put easy and medium-difficulty questions in between them so as not to make the experience a slog for players. What constitutes a “slog for players” has nothing to do with how important a particular fact is, though—it's about which questions are answered correctly and which are not. A team that just finished sitting through a brutally difficult Set of 4 doesn't care about whether or not they *should have* known any of the answers, they care that they *did not* know any of the answers. If the metric by which we judge difficulty says that that Set of 4 was balanced, it's not a very good metric. A metric which purportedly predicts which questions will be answered

correctly and which will not—that is, a conversion-based metric—is much more useful for what we’re trying to achieve and the way we use difficulty ratings when constructing packs.

This principle is especially important when considering questions that ask about course material from upper-year courses, particularly the sciences. A significant number of teams will not feature any players in grade 12: this automatically makes many questions about grade 12 sciences or things like philosophy skew hard, even if they are about core content from those courses.

In order to make sure your packs are difficulty-balanced, it’s helpful to have a scale on which to assess the difficulty of your questions. This allows you to check your difficulty balance on Sets, as well as to communicate about difficulty with other people working on the packs. At Gradwrite, we use a four-point difficulty scale to assess our questions, where our easiest questions are marked as 1 and our hardest questions are marked as 4. (We only ever use whole numbers to rate difficulty.) We typically say that a question marked as difficulty 1 is “1-gettable”, a difficulty 2 question is “2-gettable”, and so on. As explained above, this scale is a conversion-based scale. What that means is that difficulty doesn’t directly depend on the actual content of the question, but instead, on what fraction of the teams at a tournament will get your question correct. A 1-gettable question is a question that gets very high conversion, not a question that’s “fundamental to the field”.

The actual conversion numbers (that is, what percent of teams would get a question) that are attached to each difficulty rank on the four-point scale are not hard-and-fast. You might hear someone describe it as 90-100% of rooms (that is, games between two teams) should get a 1, 70-80% a 2, 30-50% a 3, and 10-20% a 4, but it is very hard to estimate these numbers for a given question.

More valuable, maybe, are heuristics. 1s should be really, truly, very easy—if you were to ask a 1-gettable question to 10 people who have never played trivia before, at least one of them should be able to get it right. In a game between two fresh-to-trivia teams, almost all of the 1s should be answered correctly. 2s are meant to be slightly tougher, but still solidly on the easier end: in a game between two of the stronger teams at a tournament, you should expect almost all of the 2s to be converted by one of the two teams. 3s are in the tougher range, so they should require a bit more specialist knowledge: not even all the top teams should be converting difficulty 3 questions, although some should be converted regularly in a game between two strong teams. 4s are the hardest questions we ask at all—everyone should struggle with 4s, and the people who do get them should have real specialist knowledge. It is perfectly normal for 4s not to be converted, even in games between strong teams.

If absolutely no one knows the answer to your question (0% of teams are “converting” it), it is probably too hard. Of course, though, there is an exception:

Guiding Principle (The “Good Trivia” Exception): Every once in a while, if a question asks about content that everyone, not just specialists in particular categories, will find interesting, it is alright to include regardless of its difficulty. The question should be interesting to pretty much anyone, regardless of other knowledge—the information contained in the question and answer alone should be enough to paint a vivid and interesting picture in the mind of almost every player. Questions that presuppose knowledge of a particular person, work, war, region, or concept rarely qualify, unless the presupposed knowledge is practically universal.

Why do we do this? The biggest thing that’s bad about a question going dead is that players get less out of the question—for the most part, trivia players like to learn, but when a typical hard question (i.e., one that requires players to recognize names or have deep knowledge of a topic to even understand what’s going on) goes dead, usually only a few players (if any) have the requisite knowledge to learn facts from the question. If the question has universal appeal, though, it’s fine if it goes dead, because everyone or almost everyone can get something out of it regardless.

Of course, any questions you include under the Good Trivia Exception will be among the most difficult questions in your packs.

When writing your packs, you’ll want a good mix of easy questions, medium-difficulty questions, and harder questions, but there should be more easy questions than medium and more medium than hard—with the exception of the few easiest questions you write, even your easy-medium questions will not be converted in some rooms. To keep the game flowing in every room, it’s best for most of your pack to be 1- and 2-gettable questions with a few 3- or 4-gettable questions peppered throughout. More detail and justification can be found in the Compiling section.

Don't swerve!

"The title character says she will 'buy the flowers herself' at the beginning of *Mrs Dalloway*, which was written by what author? A: Virginia Woolf". What's wrong with this question?

Definition: A swerve is a question which encourages knowledgeable players to buzz in early with a particular answer before "swerving" and asking for something else entirely.

Guiding Principle: When writing a question, you should always make sure that your question encourages players with more knowledge to buzz in early. Always double-check your questions to make sure they aren't swerves.

Why do we do this? In the above example, players who know a lot about *Mrs Dalloway* will buzz early in the question with "Mrs Dalloway" and get the question wrong. This is despite them having more knowledge and the skills to anticipate that the question will probably ask for what novel fits that description. Players want to be able to demonstrate their knowledge and their skill for anticipating where a question will go, but these sorts of questions punish players for having that knowledge and applying that skill.

How can you avoid accidentally writing swerves into your packs? One easy solution is to just put what you're asking for at the start of the question. In the *Mrs Dalloway* example, that would mean starting the question with "What author used the line...". Things get very boring, though, when every question starts that way, and it doesn't allow players to demonstrate their skill for predicting where a question is going. Good packs consist of a mix of questions that begin with an indication of what they're asking for (like "What author...") and questions that don't. How can you be sure, though, that those questions which don't make it immediately clear what they're asking for aren't swerves?

Rule of Thumb: If your question hasn't yet told players what it's looking for, any descriptions should be descriptions of the correct answer. We could fix the original *Mrs Dalloway* question from earlier by asking "The title character says she will 'buy the flowers herself' at the beginning of what novel by Virginia Woolf? A: *Mrs Dalloway*". This question is no longer a swerve because since the initial description is of *Mrs Dalloway*, the answer to the question is "*Mrs Dalloway*".

Of course, there are exceptions to this rule: for example, if what you're describing is clearly too complicated to be the answer to the question, you might describe something more complicated with the hopes that players realize you're looking for something more simple. This is up to your discretion, but just keep in mind what you'd feel like as a player hearing the question. Always reread and rewrite your questions to avoid swerves.

Micropyramidality

Consider the question “The bosons that carry the electromagnetic force are what massless particles, the quanta of light? A: photons”. Instead of phrasing it this way, we might instead ask it like “What quanta of light are the bosons that carry the electromagnetic force? A: Photons”. This second question is worse than the first, because it isn’t micropyramidal.

Definition: a pyramidal question is one which puts harder clues before easier clues.

Where does this terminology come from? “Pyramidality” is a term also used in quizbowl, a different trivia format, to refer to how the earliest clues in a long paragraph-like question are much, much harder than the ones at the end. You can imagine the small number of players that know hard clues as the top of the pyramid, and the bigger number of players that know easier clues like the bottom of the pyramid.

If that’s pyramidality, what’s micropyramidality? While we still follow the principle of pyramidality when writing Gradwrite questions, they’re only a single sentence long. This means that we need to pay attention to even smaller details (hence *micropyramidality*) when we think about pyramidality. Sometimes the ordering of a handful of words or using a longer or shorter connecting phrase between clauses can create small differences that give more knowledgeable players the edge they need to buzz in faster.

Guiding Principle: Questions should almost always be micropyramidal. Sometimes, your questions might have only one phrase/clue—this is fine! In that case, you just don’t need to worry too much about micropyramidality.

Why do we do this? It takes some pretty good knowledge to know that photons are bosons that carry the electromagnetic force. Comparatively, it takes less knowledge to know that photons are the quanta of light. By putting “quanta of light” before “bosons that carry...”, the question would prevent players who know more about photons from gaining an edge and buzzing in earlier, since they’ll be buzzing at the same point as those that only know about photons being “quanta of light.” Writing micropyramidal questions lets us reward players with more knowledge.

Unique Identification

Definition: a uniquely identifying (or “UI”) question is one which pins down exactly one answer.

This doesn’t mean that the question can’t have multiple things listed as acceptable in its answerline (e.g., an author’s pen name and their real name), as long as they’re all no more than different names for the same thing. It does mean, though, that it can’t have any alternative answers which refer to distinct things. As an extreme example, a question like “Who was the Prime Minister of Canada in 2025?” couldn’t work, because both Justin Trudeau and Mark Carney—two distinct people—are both correct answers. Questions that ask for multiple things or for part of a set of things (e.g., “What two men served as Prime Minister of Canada in 2025?” or “Name either of the two men who served as Prime Minister of Canada in 2025”) are fine, as long as that set of answers are the only answers that satisfy the question.

Guiding Principle: Every question should be uniquely identifying. Even if you think your question has a *best* answer, that isn’t good enough—the way the question is phrased needs to be such that there can’t be any way a player could argue for a different answer.

Why do we do this? If multiple distinct things all satisfy a question, writers have a choice: accept all of them, or accept some and reject others. Choosing some to accept and some to reject is unfair, though—by assumption, *all* of those answers satisfy the question. Even if some answers feel “better” than others, someone who answers with one of the “worse” answers has still provided a correct answer to the question, and they should be rewarded for that, because they demonstrated just as much knowledge as someone who gave the favoured answer. This would make it seem as though the better option is to accept all of the correct answers. Questions that accept multiple distinct things are misleading for players, though. If the question asks for one answer without explicitly asking for a subset of correct answers, it implies that it is the only correct answer. Accepting multiple answers implicitly feeds players incorrect information.

Since all questions need to be UI, we typically don’t like to write questions where the answer is a description. It also means that clues which assert that something is the “most famous” or the “most notable”, or anything else that’s not 100% objective, need to be supplemented with at least one other clue which is uniquely identifying on its own and leaves no room for opinion. These clues (even supplemented by others) can also feel bad for players because different players may have different ideas of what is the most famous.

Making sure your questions are UI is often the most time-consuming part of fact-checking. For example, if you’re citing an accomplishment someone has made, you need to double check that no one else has made the same accomplishment! The best way to be thorough with this is to search each piece of information in the question and see if there are any other answers

which that clue could refer to. If there are other possible answers, you need to ensure that the other clues you have in the question together rule out that potential alternate answer. You should also ensure that early non-UI clues don't bait players into buzzing with a certain obvious answer and only ruling it out in later UI clues. Whenever possible, each of the clues in a question should be UI on their own.

Useful Trick: If you're struggling to make a question UI, consider cluing something about the way the answer is spelled—the number of letters, the first letter, or something of that kind. Use this sparingly, since it can sometimes feel shoehorned in and shouldn't be an excuse to not write an otherwise not uniquely identifying question, but it can be a good tool to prevent those one or two other answers from slipping in.

Write interesting questions!

Not all of your questions need to be novel and revolutionary; sometimes, a well-written question about a well-known topic is just what players want in a pack, and these kinds of questions can be interesting in their own way (remember, just because most people know something doesn't mean everyone knows it!). Combining questions about well-known subjects in interesting ways can also give them a fresh new feel (perhaps in a mixed set—see the Writing Sets section for more detail). There are also some specific ways to go about adding some extra excitement to your questions:

Useful Trick: Ask about things that you haven't heard asked about before (and that you care about). This can lead to very hard questions, but if you can cover new topics in reasonable ways, this often leads to very well-received questions. Think about topics that you think are underrepresented in trivia—there might be a lot of askable content there that we just haven't used yet.

Useful Trick: Ask about things that have been asked about before, but in a new way or with new clues. Take a different perspective in approaching the same question topic.

Useful Trick: Look around you: items in your house, things you do throughout the day, signs you see—almost anything can be asked about in a trivia question. This inspiration can lead to some cool questions that players didn't even imagine might be asked!

Useful Trick: Look at old packs—but sparingly. Old packs can provide plenty of inspiration, but usually when you're doing this, you want to combine it with some of the strategies above too. Can you use an idea from a previous pack, but change the answerline? Can you keep the answerline but use different clues or think about the topic from a different angle? Maybe there's an old question that you think would be pretty clearly bad to include directly in a tournament because it has been asked about far too much. Is there a way you can add new information or change the way it's asked to recast those old questions and teach players something new?

You can also try entirely new ideas. Wordplay had a ton of innovation at Gradwrite Spring 2024, and players really liked this! Think about entirely new question styles and types—although you probably won't find something that will show up in every pack, a new format within a Set of 3 or 4 can add a lot of variety to the tournament.

Unfortunately, questions about truly novel topics are often very difficult—well-known things often seem less interesting simply because they're well known. How can you write questions at a reasonable difficulty level while still including interesting content?

Useful Trick: Put harder clues at the beginning of a question and easier clues at the end. For example, you might write a question about mitochondria which first talks about oxidative phosphorylation, before then mentioning that they're the powerhouse of the cell.

Why do we do this? Putting harder clues at the beginning before dropping some easier clues at the end still lets you include the interesting (but perhaps difficult) content that you want to while still ensuring that players without specialist knowledge can still get the question. More detail and justification can be found in the Micropyramidity section.

Sometimes, though, a question is hard because not many people have heard of the answer at all. There are also ways around this:

Useful Trick: If an answer is too obscure, try flipping the answerline and a clue in the question.

For example, instead of “What Viennese art museum is home to Gustav Klimt’s *The Kiss*? A: Belvedere”, you might write the easier “What city’s Belvedere art museum is home to Gustav Klimt’s *The Kiss*? A: Vienna”.

Why do we do this? In the Belvedere example, the second question is *much* easier than the first one, while still including all the same content. People who know a lot about the history of art in Vienna will get to show off their knowledge on both questions, but players who don’t know or can’t remember the name of the Belvedere can still figure out the answer to the second question by knowing a little bit about Gustav Klimt.

Finally, you should make the most of your difficult questions.

Guiding Principle: Use your hard questions as a way to get interesting content into the packs.

This means eliminating questions which are “artificially” hard—questions which, for example, require players to remember specific, non-famous years, or which ask for the names of minor characters in works of fiction.

Why do we do this? Often, in order to properly tune the difficulty of a Set, Split, Double Jump, or pack as a whole, writers find themselves in need of a particularly hard question. In these situations, it can be tempting to write a question like “In what year did Ivan Pavlov win the Nobel Prize in Physiology or Medicine? A: 1904”. This is certainly a very difficult question, and it would fill in the gap nicely! It is, however, extremely boring.

It’s important to think about *why* your questions are hard. For example, Ivan Pavlov is very well known. Someone who knows a little bit about him and the context of his work might be able to place his work around the beginning of the 20th century, even. The only reason this question is hard is because pinpointing the specific year is difficult. This may not be a bad thing in and of itself—perhaps it’s plausible to say that someone truly interested in Ivan Pavlov would know when he won the Nobel Prize (though this is also contentious). It does mean, though, that the question is pointless to anyone who knew that Pavlov worked at the start of the 20th century—sure, after hearing the question and its answer they now know that he won in 1904 specifically, but that’s not a particularly interesting fact. The same goes for questions that ask for the name of a minor character from a work of fiction—a name alone is meaningless. Often, players will know exactly

who the question is asking about (indeed, perhaps they've even read the book, or watched the movie, or otherwise consumed the work of fiction) but they will not be able to come up with the name of the character in time. When it's finally revealed, they have learned nothing.

Aside from the active sense in which artificially hard questions are boring, they're also a waste of space. As discussed in the Difficulty and Compiling sections, hard questions should be pretty rare compared to easy and medium questions. Slots for difficult questions are unbelievably precious! It's a real shame when one of those slots is taken up not by something interesting, but by a question which gives players no new knowledge and instead tests whether they can recall fiddly details.

When you find yourself in need of a hard question, instead of writing something artificially hard, do some research! Often, the most interesting questions get written when writers do their own research and stumble across something they didn't know before. Going further afield is also a great way to find new topics, rather than delving even deeper into the same exhausted subjects that have been asked about time and time again.

Formatting Details

For the most part, how you choose to format your questions and answerlines doesn't matter, as long as everyone working on your packs is in agreement (for example, if you're writing for Gradwrite, you should follow the format outlined in the Gradwrite Format Guide). Some decisions, though, can have a real effect on the way the game is played:

Guiding Principle: Answerlines should leave no room for interpretation. Things like “accept reasonable equivalents” should not be included.

Why do we do this? In a fair game, questions test players, and ideally only players. When the reader is asked to make a judgment call, whether a team gets points for a question suddenly depends on the reader in the room they were randomly assigned to play in. It is extremely common that one reader will consider a given answer to be a “reasonable equivalent” of the listed answer, but, in a different room with a different reader, a different team will give that same answer and have it rejected. In order to eliminate this dependence on the reader and ensure the fairness of the game, answerlines should be as precise and objective as possible.

Guiding Principle: Don't include something in an answerline unless you think it's plausible that not including it would cause problems in a game. This means not including alternate acceptable answers unless you think someone might really give that answer in-game, and not including answers to explicitly reject unless you think that the reader might accidentally mishear that answer as the correct answer.

Why do we do this? At game speed, scanning through a long answerline can take a long time. When a reader takes significantly more than a second to look through an answerline, it can give some teams an unfair advantage—often, the reason that a reader hesitates is because the given answer was close, but not quite acceptable, and the reader needs to double-check the full answerline to make sure. This hesitation can clue the other team in that they should give a similar answer. Regardless of whether that's the cause of hesitation, the extra time that the reader takes gives the other team more time to think. Fairness aside, this hesitation can also just slow down the game. In order to keep the strain on the reader minimal, we like to keep our answerlines concise.

While it is true that “Pablo Diego José Francisco de Paula Juan Nepomuceno María de los Remedios Cipriano de la Santísima Trinidad Ruiz y Picasso” is a correct answer to a question about Pablo Picasso, the chance of a player actually giving that answer in-game is vanishingly small, and so it's not worth increasing the amount of time it takes to read through the answerline. On the other hand, it does seem very possible—likely, even—that someone could answer “North Atlantic Treaty Organization” to a question about NATO, so it is worth including in the answerline. If, in-game, a player does give an answer which wasn't included in the answerline for concision, they can always

challenge, and if their answer is indeed correct then their team will get the points they are owed.

The same reasoning applies for why we like to exclude most explicit instructions to reject certain answers. At Gradwrite tournaments, staff are instructed to adhere as strictly as possible to the answerlines, even if they think the answerline is incorrect. This means, for example, that there is no need to explicitly tell readers to reject “Guyana” for a question whose answerline is “**French Guiana**”—since “French Guiana” is entirely underlined, the reader has no reason to do anything but reject a partial answer. The only case where explicit rejection instructions are necessary is where a reader might mistake one answer for another—for example, since Guyane is another name for French Guiana, the answerline might look like “**French Guiana** (accept: **Guyane**)”. In this case, since a reader could reasonably mishear “Guyana” as “Guyane”, it would be helpful to include “(reject: **Guyana**)” in the answerline.

Perhaps less important, but still important, is making sure the words and names in your questions are pronounced correctly.

Guiding Principle: Include pronunciation guides in your questions. Most people are not familiar with IPA, so we find the clearest pronunciation guides list syllables as they’re commonly written in English, where capitals indicate stress: for example, “Mixolydian [“MIX-oh-LID-ee-an”]”. We strongly recommend including pronunciation guides on all spelling questions, even if it seems as though the word would be impossible to mispronounce.

Why do we do this? If a reader mispronounces a key word or name, the teams in that room can be severely disadvantaged (or advantaged, for spelling questions) compared to other rooms. Especially for spelling questions, the consequences of a mispronunciation can make a real difference on the game. Given the minimal amount of effort required to write a pronunciation guide for a difficult word or name, it’s best to err on the side of caution.

PART II:
PARTICULAR
QUESTION TYPES

Writing Sets, Splits, and Unrevealed Sets

At Gradwrite, we typically use the term “set” as a catch-all term for Sets of 3, Sets of 4, Splits, and Unrevealed Sets. (If we’re talking about a Set of 3 or Set of 4, we’ll say “Set of 3” or “Set of 4”). We also distinguish between two types of sets:

Definition: A pure set (or straight set) is a set whose questions are all about the same topic. For example, a category called “Set of 4: Celtic Myth” whose questions all in fact do test knowledge of Celtic mythology (rather than, for example, a set where one of the questions is about Celtic myth in pop culture) is a pure set.

Definition: A mixed set is a set whose questions test different areas of knowledge. For example, a category called “Set of 4: Red Things”, where the thing linking the questions together is not their topic but something shared between their topics, is a mixed set.

Guiding Principle: Splits should always be pure sets, and the title of a Split should always be very clear. For example, you wouldn’t want to have a Split whose title is “Power” and then have it be a mixed set about the word “power” in different contexts. Even if it is a pure set about, say, power in physics, you should make sure to title the split “Power in Physics” to eliminate the possibility of a player mistaking it for a set about power in some other context.

Why do we do this? When players choose splits, they should know what they’re getting into. It’s no fun to divide up your team only to find that the Split titles were misleading, and so now some players are stuck on a Split that they don’t know anything about. It also introduces an element of randomness which takes away from the fairness of the game—if players aren’t quite sure what they’re getting into, it comes down to the other skills of each player (not the ones that the team thought they were selecting for when they chose splits) who gets the questions in a Split.

This doesn’t mean that you should avoid mixed sets altogether! Mixed sets are a great opportunity to show how clever and creative you are. At Gradwrite, we try to write most (but certainly not all) of our Sets of 3, Sets of 4, and Unrevealed Sets as mixed sets.

Guiding Principle: Mixed sets should include as many different areas of knowledge as possible.

Why do we do this? Linking 3 or 4 different areas of knowledge to the same topic is very impressive. If you’re writing a mixed set that links together mixed topics, it should earn its place—cluing the same topic multiple times makes it seem as though there really isn’t that much to link together after all.

As with all questions, difficulty control is one of the most important parts of writing a good set.

Guiding Principle: Almost every set should have at least one 1-gettable question and at least one 3- or 4-gettable question. In general, the other questions in the set should be

somewhere between 1- and 3-gettable. In cases where there are exceptions to this rule, they come from leaving out the 3- or 4-gettable question, never the 1-gettable question. It is not good enough for a set's *average* difficulty to be reasonable—every set must have a 1-gettable question.

Why do we require a 1-gettable question? Without any 1-gettable questions in a set, it becomes extremely likely that the entire set—three or four questions in a row—goes dead in a number of rooms. This feels terrible for the players in those rooms.

It may seem that a set with low average difficulty—for example, a set where all the questions are 2-gettable—should not be subject to this requirement. Even four 2-gettable questions in a row can reasonably all go dead in a number of rooms, however. This is especially true in light of the fact that a pair of teams which fails to convert three 2-gettable questions on a particular topic probably doesn't know very much about that topic, and so doesn't have a great chance at converting the fourth.

Why do we require a 3- or 4-gettable question? True, easy questions are fun, but in order to meaningfully distinguish between top teams, packs need questions which require a bit more specialist knowledge. If all the questions in a Set of 4 on Celtic myth are easy or medium-difficulty, a team which has middling knowledge of Celtic mythology might fare just as well as a team that has very deep knowledge of the subject, which is unfair to the more knowledgeable team. Also, it's fun to show off deep knowledge! Adding a hard question in each set allows for these benefits without making the entire set feel punishing for non-specialists.

When, then, is it fine to leave out the hard question?

Rule of Thumb: If a mixed set would be made worse by making one of the questions harder, it's alright for all of the questions to be easy or medium-difficulty.

Why is this OK? As discussed in the “Write interesting questions!” section, artificially difficult questions are no fun for anyone, so if it really doesn't seem like there's a naturally hard question for your mixed set, it would be best not to have to force something. For a pure set, this reason isn't good enough, for the reasons outlined in the previous Guiding Principle. You need to do more research or figure something out that's between artificially and naturally hard. However, the problems that present themselves for a too-easy pure set—namely, that specialists don't get credit for their deep knowledge—aren't a problem for mixed sets: the way they link questions together is rarely ever something that someone could intentionally specialize in. Nobody specializes in facts about things that are red, for example.

Useful Trick: Before writing the questions for a set, plan out the answerlines you want to use and the intended difficulty of each question. This lets you make sure that your set will be difficulty-controlled before you take the time to figure out the wording of the

questions. If you can't find an answerline that you could write a 1-gettable question for, it may be a sign that your set is too hard.

Beyond difficulty, you should also make sure that your sets don't feel uninspired. A big part of this is using interesting clues and ensuring that the hard questions in the set aren't artificially hard, but the way a set feels often depends strongly on how the questions work with one another and with the set's theme. Not every set is going to be revolutionary, but there are a couple of things you can check for to make sure they don't feel flat.

Guiding Principle: Questions should *not* be a single word or phrase. In particular, try to avoid sets of the form "given X, name Y".

Why do we do this? First and foremost, single-word and single-phrase questions present practical difficulties. Consider a Set of 4 with the following structure: given a country, players must name the capital of that country. Now suppose that there is someone on each team who has memorized all of the capitals of all of the countries in the world. It is in the best interest of each of those players to buzz in as soon as the reader says anything at all—perhaps even a syllable is enough for them to figure out what country is being asked about. Often, in anticipation, players will buzz in before the reader has even started speaking. Cases like these can make it extremely difficult for the reader to decide what to do. These types of issues are most pronounced when the list being asked about is memorizable (e.g., countries and their capitals, elements and their symbols, etc.), but they can ruin the experience for players just as badly even for lists that are in principle not memorizable (e.g., books and authors).

From a more theoretical and aesthetic perspective, single-word and single-phrase questions are boring. It feels as though the writers had an empty place in their pack and didn't know what to put there, so they decided to fill it in by writing 3 or 4 questions of exactly the same structure. Regardless of how much effort is put into choosing the theme or answers for single-word or single-phrase sets, they invariably come out sounding uninspired, lazy, and boring. These types of sets also tend not to test degrees of knowledge very well.

Guiding Principle: Whenever possible, try to have diverse answerlines within a set, and make sure that your questions clue their answers in different ways. This means that even in a set like "Split: African Rivers", not all of the answerlines should be rivers. Even just one or two questions that ask for a different type of thing (e.g., a country, a city, etc.) can make a big difference. It also means that even if your answers aren't all rivers, there should be at least a couple of questions in the set that don't boil down to just testing associations between cities and the rivers they lie on.

Why do we do this? A set feels monotonous if players have to sit through what feels like multiple copies of the same question. This problem is exacerbated if a player or team isn't inclined to be able to get that type of question, and can be even worse if players

who otherwise have relevant knowledge are locked out by the format. For example, a player might know cool facts about rivers, but can't associate them with nearby cities. They will sit through the entire set on rivers without any opportunity to buzz. Cluing sets in different ways keeps players on their toes.

Also, when all of the answerlines in a set are the same type of thing, it can seem as though the entire set is an exercise that needs to be gone through. It feels as though the writers needed to fill in a set and so they picked (for example) a few rivers and a fact about each of those rivers. Even if the rivers were chosen carefully and creatively, that doesn't always show through to players hearing the set. This is even more so the case for sets which clue all of their answers in the same way. It's the same problem as the one that arises for single-word or single-phrase questions, albeit with much less disastrous consequences.

Guiding Pineapple: Sets should be something over and above their constituent questions.

There are lots of ways to achieve this, but we'll highlight two of the most common:

Tight theming can make a set feel inspired. For example, "Set of 4: Religions" is pretty boring—at that point, it might be better to just spread the questions out as Blitzes throughout the rest of the packs—but the specificity of "Set of 4: African and Afro-Caribbean Religions" makes it compelling. (This is just an example, and it doesn't mean that your set titles need to be super long! Indeed, wherever possible, concision in set titles is a good thing.). A good rule of thumb is that the questions should adhere to a tight enough theme that you should *not* be able to spread the questions of the set throughout the packs as Blitzes without players noticing.

A more subtle way of making your sets interesting is by making it possible for players to buzz on some questions in the set earlier than they would be able to if the question was a Blitz. For example, as a Blitz, the question "What is the NATO phonetic code word for the letter Y? A: Yankee" is impossible to buzz on until the last word. In the category "Set of 4: The Letter Y", though, players who are paying attention can buzz in by "code word", or even as early as "NATO" with a bit of anticipation. This strategy typically works best in mixed sets, but if done well it can work in pure sets too. Be careful, though: there's a fine line between questions which reward attention and questions which are just antimicropyramidal.

In short, there should be some reason why the questions in your set deserve to be in a set together, rather than just being used as Blitzes.

Why do we do this? There's nothing actively bad about boring sets, but they are, after all, boring. A boring set here and there isn't the end of the world, as long as the questions are well-written and the content of each question is interesting, but having boring set after boring set can really make a pack feel flat. A majority of sets feeling inspired is a hallmark of a masterfully-crafted pack.

In general, you want to make it seem as though you put in real effort to make your packs interesting.

Guiding Principle: Only write an Unrevealed Set if there's no way to write it as a Set of 4.

Why do we do this? Unrevealed Sets are less aesthetically appealing than Sets of 4, and to a certain extent they interrupt the flow of the game. These are minor reasons, but the amount of effort required to transform an Unrevealed Set into a Set of 4 is typically very small, so it's almost always worth it to just write the Set of 4 instead. Typically, the only reason we write Unrevealed Sets is because there are certain topics on which it might not be possible to write a difficulty-controlled Set of 4. Having the answer to the first question be the category for the rest of the set gives the writer a lot of freedom, because often the best-known thing about a particular subject is its name.

Guiding Principle: Within a set, write and order questions to vary question length.

You might choose to write longer questions in order to include deeper clues in a micropyramidal structure, or explain a more complex topic. On the other hand, shorter questions are a great way to keep players sharp on the buzzer. Keep in mind that the length of the question on the page is not always a perfect indicator of how long a question will be when played. If you expect that most teams will be buzzing on the first few words of a long question, it will play as a short question.

Why do we do this? Varying length tends to lead to sets that are engaging for a couple of reasons. For one, variation in length means players need to stay alert to questions that might require more patience, or questions that require a quick, clever buzz. This effect results in players getting to demonstrate a breadth of gameplay skills. More abstractly, the rhythm and pace of gameplay is a big part of the character of the game, so careful selection of rhythm and pace can make a big difference to how fun the game feels.

Writing Double Jumps

Definition: The first question of a Double Jump is called the first jump. The second question in a Double Jump is called the second jump.

Guiding Principle: The first jump of a Double Jump should be easy—ideally, 1- or 2-gettable.

Why do we do this? The only thing that separates the way a Double Jump plays from the way a Blitz plays is that the second jump is only available to the team that correctly answers the first jump. In whatever rooms the first jump goes dead, the entire Double Jump plays like a pair of Blitzes, which defeats the purpose of writing it as a Double Jump in the first place. As such, we like to keep the expected number of rooms which miss the first jump low.

Guiding Principle: The second jump of a Double Jump should be clearly harder than the first jump. A good rule of thumb is that the difficulty rating of the second jump should be at least 2 points higher than the difficulty rating of the first jump. The hardest questions in a tournament will often be among the second jumps of Double Jumps.

Why do we do this? Suppose the second jump of a Double Jump is around the same difficulty as the first jump: that is, a team which has the knowledge to get one will likely have the knowledge to get the other. Since the team that gets the first jump gets exclusive access to the second jump, and since by assumption they likely have the knowledge to convert the second jump, it is as though the first jump is worth 20 points—the points from getting it correct and the points from their uncontested opportunity to get the easy second jump. Aside from the fact that in general bigger swings are less desirable, if we wanted to have one question be worth 20 points, we would simply award 20 points for the one question, rather than going through the trouble of the Double Jump structure.

Guiding Principle: Double Jumps should have a coherent, tight theme. At weakest, this means that the second jump should clue the same knowledge as the first: no pivots from ocean currents to electrical currents. A good (though certainly not infallible) rule of thumb is that someone who knows the answer to the second jump should also know the answer to the first jump. At strongest, tight theming means that the first and second jump should create an unspoken hyperspecific theme beyond just the first answer being related to the second.

Why do we do this? As with sets, tight theming is inherently desirable because it is aesthetically appealing. Coherent theming is even more important for Double Jumps, though, because of their formal structure. A team which has correctly answered a first jump about a particular topic has earned the opportunity to answer another question *about that same topic*. As argued earlier, second jumps should be really quite tough. If the second jump carries on the theming of the first, the team that got the first jump gets

the opportunity to show how deep their knowledge is on that topic, having "earned" the right to by answering the first jump correctly. However, if the second jump does not carry on the theme of the first, that team has just been given a hard question on a random topic for no "earned" reason. The first jump is like a promise—the team deserves a fair chance at getting the second jump. Also, changing the knowledge base increases the probability that the second jump will play like a really hard and less fun blitz.

Useful Trick: Write Double Jumps by extending Blitzes. If you ever find yourself short of Double Jumps, you can take a hard Blitz, write another, easier Blitz about the same topic, and use that new Blitz as a first jump and the original Blitz as the corresponding second jump. Alternatively, if you have an easy Blitz about a topic you want to explore further, you can use the original as the first jump and write something more difficult for the second jump.

Useful Trick: Write Double Jumps by salvaging from unusable pure sets. If you were planning to write a pure set but couldn't manage to calibrate the difficulty properly, consider taking the easiest question in the set and the hardest question in the set and pairing them to form a Double Jump. Be careful that the theme is still coherent without the set title, though.

Writing Jackpots

At Gradwrite, we award 30 points for a correct answer on Part 1 of a Jackpot, 25 for Part 2, 15 for Part 3, and 10 for Part 4. Each team has two chances to buzz throughout a Jackpot. As with anything else in this guide, you can play around with the specific score numbers and rules, but you will notice that we have chosen a different scoring scheme than the one used by some other organizations. In particular, we chose to reduce the value of the first clue, as we felt that awarding 40 points resulted in Jackpots having a disproportionately large impact on the game.

On the surface, Jackpots are the question type which least resemble Blitzes, but in a number of ways they play very much like a set or a collection of Blitzes (and often more like Blitzes than Chains or Mystery Links). It's easy to forget to apply certain basic question-writing principles when writing Jackpots—in particular, unique identification and difficulty calibration.

Guiding Principle: Every part of a Jackpot should be uniquely identifying. Moreover, each clue should be useful in a meaningful way to someone who knows a lot about the answer.

It's hard to describe exactly what this means, but it's pretty clear what doesn't work; putting the answerline's birthday as the first clue, for example, is both not uniquely identifying and not that meaningful or interesting.

Why does each clue need to be useful? If there are clues which are useless to someone who knows a lot about the answer, those clues are almost certainly useless to everyone else too, so that clue is impossible to buzz on. Just as with individual Blitzes, the only reason why you might want to include a question which you expect 0% of players to know the answer to is the Good Trivia Exception.

Why does each clue need to be uniquely identifying? Theoretically, as long as the first clue of a Jackpot is uniquely identifying, the whole thing should be fine: after all, the answer needs to satisfy all four parts of the Jackpot. In practice, though, this is far from ideal. If a later part of a Jackpot does not on its own pick out exactly one answer, it ceases to be a useful clue: presumably, if a player is hearing a later part of a Jackpot, it's because they didn't know the clues in the previous part, and so if later clues only uniquely identify an answer in conjunction with earlier clues, those later clues require some knowledge of earlier clues in order to help players.

Guiding Principle: Each clue in a Jackpot should be clearly easier than previous clues. This has two parts: you should order the clues from hardest to easiest, of course, but it's also extremely important that the clues are of meaningfully different difficulties.

Why do we do this? Ordering clues from hardest to easiest is simply an application of pyramidality, which is discussed in detail in the Micropyramidality section. Why do the clues need to be different difficulties, though? What's wrong with, for example, a Jackpot whose Part 1 and Part 2 are of the same difficulty? Well, if that was the case, then almost everyone who would know the answer on Part 2 would also know the

answer on Part 1, and so they'll never get to Part 2, because they'll buzz on Part 1. This means that Part 2 is effectively a useless clue to most people who will hear it, which defeats the purpose of including it at all. More importantly, if Part 1 is of a similar difficulty to Part 2, the difference in points is not reflective of a meaningful difference in knowledge.

Guiding Principle: The first part of your Jackpot should be converted by very few teams, but most (if not all) teams should know the right answer by the end of Part 4. This means that the answer to any Jackpot needs to be something that's very well-known.

Why do we do this? Well, the first clue should be difficult because it's worth much more than a regular question, and you don't want to be handing out a ton of points for knowing relatively easy clues. The last clue should be easy for the same reason that we don't like to have sets which don't have any easy questions—it feels terrible for two teams to sit through an entire Jackpot that ends up being about something neither of them knew in the first place. Jackpots need to have fairly well-known answers to keep the game accessible.

The last challenge with Jackpots is that you need to be careful with the clues and wording you use to make it hard for teams to guess correctly unless they really have relevant knowledge.

Definition: A fraudable question is one whose answer can be guessed (or "frauded") by players who don't actually know the clues mentioned.

In general, we try to avoid writing too many fraudable questions, but fraudability becomes a much, much bigger problem when the point values are that much higher.

Guiding Principle: Jackpot clues, especially early Jackpot clues, absolutely cannot be fraudable. The most common things to check for when trying to determine if clues are fraudable are names and groups. For example, if you're asking for a country and have lots of names in the question, players may be able to recognize the language the names come from and guess correctly despite not knowing the actual clues. Similarly, if you're asking for a member of a group and the answer is the only famous member (or clearly the most famous member) of that group, players may decide to throw out a guess which just happens to be correct. The same problem arises with small groups, because players can use the process of elimination—a Jackpot that starts off by asking "What Canadian territory...", for example, will cause lots of problems.

Why do we do this? It's already somewhat problematic in, for example, a fraudable Blitz, where teams can get 10 points despite not having the intended requisite knowledge, but to get awarded 25 or 30 points for a lucky guess isn't an opportunity we want to provide to players. It's plainly unfair for a team to get a 50 or 60 swing without earning it by having deep knowledge.

Writing Chains

A good Chain is very fun to play, and can be one of the most reliable hallmarks of a good pack. Naturally, that means that good Chains are incredibly challenging to write! The process of writing a Chain—starting from an initial idea, building out the six questions needed, and finally refining all of the questions to be good trivia questions in their own right—will probably take longer than any other single group of questions you’ll write. On the other hand, it can be pretty easy to write a Chain that you think looks alright, but doesn’t end up working out as well as you’d hope. Here are some general rules your Chains should follow to be as good as possible:

Guiding Principle: The link should matter for every question. In a Chain, knowing the answer to the previous question *must* allow a player to buzz earlier than if they didn’t have that information.

This has two components. The first component is that the linking clue (the one related to the previous question) should be quite easy—if only a few players have the requisite knowledge to use the link, most players are locked out of the only Chain in the pack. Ideally, the linking clue should be the easiest clue in the question.

The second component is to make sure that the question is worded in a way that allows the player to use the link. A good way of doing this in a Chain is to place the linking clue at the very end of the question, so that players who are paying attention can anticipate where the question is going. For example, if the answer to a question was “Andersen,” I should follow it up with “What style of animation was used in films such as *Isle of Dogs* and *The Fantastic Mr. Fox* by filmmaker Wes Anderson?” instead of “Wes Anderson used what style of animation in films such as *Isle of Dogs* and *The Fantastic Mr. Fox*?” In the first example, a player thinking fast can probably buzz immediately after “style of animation” by knowing that Wes Anderson is most notable for using stop-motion animation, but in the second example the question just plays like any other Blitz, because any advantage from the link is immediately lost. You need to be careful that the question is still micropyramidal, though, because if easier clues are placed before the linking clue then players will buzz on the easy clues without using the link at all. This is why it’s helpful for the linking clue to be the easiest clue.

Why do we do this? If the link doesn’t matter, there’s no difference between a Chain and a normal Blitz question! Sure, it may set the player up to be able to use the link on the next question, but the original question is a waste of space in the Chain.

Guiding Principle: Every question in a Chain should explicitly include the previous answer. It’s easy, while writing, to forget to explicitly name the thing in the previous answerline because the connection between the previous answer and the way you incorporated the link into the question seems obvious.

Why do we do this? Players expect the connection to be explicit, and this assumption helps players anticipate where the question is. Also, you shouldn't have to know about the topic to recognize the link.

Useful Trick: Use homophones to get a link to work. For example, if the previous answer was “psi”, the Greek letter, you could try to clue psi again, or you could clue the singer PSY, which might make for a better category spread and avoid double-cluing. Use this trick carefully, though, because it can make it harder for players to use the link to anticipate where the next question is going.

Guiding Principle: The questions that make up a Chain should span as many different topics as possible.

Why do we do this? We want as many players as possible to have a chance at playing the only Chain in the pack! Including a wide spread of topics makes the Chain accessible to as many players as possible. It's also very impressive to see so many fields of knowledge woven together.

Guiding Principle: Make sure that no question in a Chain tests the same knowledge as the previous question. An easy way to build out a Chain is to have one topic, and ask several back-to-back questions about that topic which are all linked to one another—for example, following up “Who painted *Les Femmes d'Alger*? A: Pablo Picasso” with “What painting depicts the bombing of its namesake Spanish town, and was painted by Pablo Picasso? A: *Guernica*”. This makes for a bad Chain. A good way to check that you haven't accidentally written this type of Chain is to make sure that the questions test knowledge of different topics (e.g., history and science, rather than history and history).

Why do we do this? This type of link removes the fun of the lateral thinking involved with using the link. The links are too predictable, and most of the players at the table (in the example, those who don't know much about Picasso) are locked out of the only Chain in the pack.

Useful Trick: Write Chains collaboratively. Often, the different perspectives and knowledge bases that other people have allow them to think of innovative links that you wouldn't have been able to see yourself.

Writing Mystery Links

If Chains are the hardest questions to write, Mystery Links are a close second. Just as with a good Chain, though, a well-written Mystery Link is often the highlight of a pack.

A lot of the guidelines for good Chains have Mystery Link counterparts. In each case, the justification is more or less identical to that found in the Writing Chains section.

Guiding Principle: The link should matter. In a Mystery Link, a player who knows the link *must* be able to buzz earlier on the last few questions of the Mystery Link than they would be able to if they didn't know the link.

This has two components. The first component is that the link should be quite easy—if only a few players have the requisite knowledge to use the link, most players are locked out of the only Mystery Link in the pack. This means that the link needs to be a set with very well-known members. Small sets work especially well for Mystery Link.

The second component is to make sure that the questions are worded in a way that allows players to use the link. This means that the start of the question picks out one—and, crucially, only one—of the members of the linking set. For example, if the link was NHL teams, you wouldn't want to start a question with “What animal...”, because there are lots of NHL teams named after animals. Typically, questions which ask for a person are tough to apply the link to, because anything could be a person's name.

Starting the question with the “ask” (e.g., “What movie...”) or starting by establishing context (e.g., “In Greek mythology,...”) are usually helpful ways of starting a question late in a Mystery Link.

Guiding Principle: The questions that make up a Mystery Link should span as many different topics as possible.

Guiding Principle: Make sure that none of the questions in a Mystery Link clue the same knowledge as the link. One way to write a question for a Mystery Link about (for example) playable characters in *Super Smash Bros.* might be “What character is voiced by Ryan Reynolds in a 2019 film which sees him in the role of a detective? A: Pikachu”. Here, you're testing knowledge about Pikachu both in the question and in the link, which isn't great.

Similarly, tricks that work for Chains tend to be helpful for Mystery Links too:

Useful Trick: Use homophones to get a link to work. For example, if you're writing a Mystery Link on Greek letters, you could write a question about the singer PSY rather than trying to figure out a way to have the answerline be “psi”. Be careful, though, because this can make it harder for players to recognize or apply the link, because they'll be thinking of the answerline as spelled in the context of the question, not the link.

Useful Trick: Write Mystery Links collaboratively. Often, the different perspectives and knowledge bases that other people have allow them to think of innovative ways to clue link answers that you wouldn't have been able to see yourself.

Of course, there are also a number of principles unique to Mystery Links:

Guiding Principle: The answer in the context of the link should always be acceptable for any question in a Mystery Link. For example, if the link is songs on *The Rise and Fall of a Midwest Princess*, you cannot include a question whose answer is “midnight”, because the name of the song is “After Midnight”. Similarly, you shouldn't require players to include extra information: you also couldn't include a question in that Mystery Link which required players to say “decaf coffee” (rather than just “coffee”), because the name of the song is “Coffee”.

Even in cases where the answer in the context of the link is acceptable, you should try to make the answerline precisely what it should be for the link—if the answer is a person, and only the person's first name fits the link, consider starting the question “What is the first name of...”. This is awkward, but it serves the purpose of a Mystery Link better because it makes it easier for players to figure out and apply the link.

Why do we do this? We like our answerlines to be precisely what they should be for the link because it makes recognizing the link from the answers significantly easier for players—completing phrases or picking out the relevant parts of a not-entirely-relevant answerline is extremely difficult to do at game speed. Also, if the answer in the context of the link isn't acceptable, it can punish players who try to use the link.

Useful Trick: You can explicitly tell readers to reveal a particular answer. For example, suppose the answer to a question in a Mystery Link is Emma Wodehouse, but only “Emma” is relevant to the link. Since Emma Wodehouse is a fictional character, either “Emma” alone or “Wodehouse” alone should be acceptable. You can write “**Emma Wodehouse** (accept either underlined part, but reveal: **Emma**)”. This ensures that all rooms get to hear the answer intended for the link, even if the question is answered “Wodehouse”. Part of writing a good Mystery Link is balancing the use of this trick with “forcing” certain answerlines (e.g., “What is the first name of...”).

Guiding Principle: When ordering the questions in a Mystery Link, place them from the hardest to use the link at the beginning to the easiest to use the link at the end. Typically, this means putting the best-known answers in the context of the link near the end of the Mystery Link. For example, if you're writing a Mystery Link on Major Arcana tarot cards, you should put questions about The Tower or Temperance earlier than questions about Wheel of Fortune.

Why do we do this? Ordering the questions in this way gives players the maximum amount of time possible to figure out the link. Players who know the topic of the link very well can recognize it from the obscure link answers and gain an advantage on the

last two or three questions, while players who don't immediately recognize the link still have a chance to figure it out and use it to their advantage later on. If the questions which were easiest to apply the link to were placed at the beginning, the Mystery Link would play like a set of Blitzes for anyone who couldn't figure out the link immediately. The link can't be used before players figure it out, so the first question doesn't need to be written to value knowledge of the link, but as it becomes easier to know the link, that knowledge should become more valuable.

Guiding Principle: The questions near the end of a Mystery Link should be extremely difficult for a player who doesn't know the link. For a player who does know the link, the late questions should be quite easy.

Why do we do this? At Gradwrite, we've run into problems where very experienced and knowledgeable players are able, despite not paying attention to the link at all, to buzz on the questions in a Mystery Link as though they were Blitzes (and even to outbuzz players who have been paying attention and who know the link). This is against the spirit of a Mystery Link—it means that in some rooms, the category essentially boils down to a collection of Blitzes. Making the last couple of questions in a Mystery Link arbitrarily difficult for a player who doesn't know the link drastically decreases the likelihood of this problem arising.

Writing Mystery Links is very difficult, and sometimes (more often than with Chains) you'll find that a particular link is just not viable, and needs to be abandoned.

Useful Trick: If a Mystery Link doesn't seem to be working out, consider turning it into a mixed set. For example, if you can't figure out a way to write a Mystery Link about tarot cards, you might write something like "Set of 4: Also a Tarot Card".