

# The Best Essay

paulgraham.com · Paul Graham · 2024-03 · [source](#)

---

| | |

|

March 2024

Despite its title this isn't meant to be the best essay. My goal here is to figure out what the best essay would be like.

It would be well-written, but you can write well about any topic. What made it special would be what it was about.

Obviously some topics would be better than others. It probably wouldn't be about this year's lipstick colors. But it wouldn't be vaporous talk about elevated themes either. A good essay has to be surprising. It has to tell people something they don't already know.

The best essay would be on the most important topic you could tell people something surprising about.

That may sound obvious, but it has some unexpected consequences. One is that science enters the picture like an elephant stepping into a rowboat. For example, Darwin first described the idea of natural selection in an essay written in 1844.

Talk about an important topic you could tell people something surprising about. If that's the test of a great essay, this was surely the best one written in 1844.

And indeed, the best possible essay at any given time would usually be one describing the most important scientific or technological discovery it was possible to make.

[1]

Another unexpected consequence: I imagined when I started writing this that the best essay would be fairly timeless — that the best

essay you could write in 1844 would be much the same as the best one you could write now. But in fact the opposite seems to be true. It might be true that the best painting would be timeless in this sense. But it wouldn't be impressive to write an essay introducing natural selection now. The best essay now would be one describing a great discovery we didn't yet know about.

If the question of how to write the best possible essay reduces to the question of how to make great discoveries, then I started with the wrong question. Perhaps what this exercise shows is that we shouldn't waste our time writing essays but instead focus on making discoveries in some specific domain. But I'm interested in essays and what can be done with them, so I want to see if there's some other question I could have asked.

There is, and on the face of it, it seems almost identical to the one I started with. Instead of asking what would the best essay be? I should have asked how do you write essays well? Though these seem only phrasing apart, their answers diverge. The answer to the first question, as we've seen, isn't really about essay writing. The second question forces it to be.

Writing essays, at its best, is a way of discovering ideas. How do you do that well? How do you discover by writing?

An essay should ordinarily start with what I'm going to call a question, though I mean this in a very general sense: it doesn't have to be a question grammatically, just something that acts like one in the sense that it spurs some response.

How do you get this initial question? It probably won't work to choose some important-sounding topic at random and go at it. Professional traders won't even trade unless they have what they call an edge — a convincing story about why in some class of trades they'll win more than they lose. Similarly, you shouldn't attack a topic unless you have a way in — some new insight about it or way of approaching it.

You don't need to have a complete thesis; you just need some kind of gap you can explore. In fact, merely having questions about

something other people take for granted can be edge enough.

If you come across a question that's sufficiently puzzling, it could be worth exploring even if it doesn't seem very momentous. Many an important discovery has been made by pulling on a thread that seemed insignificant at first. How can they all be finches?

[2]

Once you've got a question, then what? You start thinking out loud about it. Not literally out loud, but you commit to a specific string of words in response, as you would if you were talking. This initial response is usually mistaken or incomplete. Writing converts your ideas from vague to bad. But that's a step forward, because once you can see the brokenness, you can fix it.

Perhaps beginning writers are alarmed at the thought of starting with something mistaken or incomplete, but you shouldn't be, because this is why essay writing works. Forcing yourself to commit to some specific string of words gives you a starting point, and if it's wrong, you'll see that when you reread it. At least half of essay writing is rereading what you've written and asking is this correct and complete? You have to be very strict when rereading, not just because you want to keep yourself honest, but because a gap between your response and the truth is often a sign of new ideas to be discovered.

The prize for being strict with what you've written is not just refinement. When you take a roughly correct answer and try to make it exactly right, sometimes you find that you can't, and that the reason is that you were depending on a false assumption. And when you discard it, the answer turns out to be completely different.

[3]

Ideally the response to a question is two things: the first step in a process that converges on the truth, and a source of additional questions (in my very general sense of the word). So the process continues recursively, as response spurs response.

[4]

Usually there are several possible responses to a question, which

means you're traversing a tree. But essays are linear, not tree-shaped, which means you have to choose one branch to follow at each point. How do you choose? Usually you should follow whichever offers the greatest combination of generality and novelty. I don't consciously rank branches this way; I just follow whichever seems most exciting; but generality and novelty are what make a branch exciting.

[5]

If you're willing to do a lot of rewriting, you don't have to guess right. You can follow a branch and see how it turns out, and if it isn't good enough, cut it and backtrack. I do this all the time. In this essay I've already cut a 17-paragraph subtree, in addition to countless shorter ones. Maybe I'll reattach it at the end, or boil it down to a footnote, or spin it off as its own essay; we'll see.

[6]

In general you want to be quick to cut. One of the most dangerous temptations in writing (and in software and painting) is to keep something that isn't right, just because it contains a few good bits or cost you a lot of effort.

The most surprising new question being thrown off at this point is does it really matter what the initial question is? If the space of ideas is highly connected, it shouldn't, because you should be able to get from any question to the most valuable ones in a few hops. And we see evidence that it's highly connected in the way, for example, that people who are obsessed with some topic can turn any conversation toward it. But that only works if you know where you want to go, and you don't in an essay. That's the whole point. You don't want to be the obsessive conversationalist, or all your essays will be about the same thing.

[7]

The other reason the initial question matters is that you usually feel somewhat obliged to stick to it. I don't think about this when I decide which branch to follow. I just follow novelty and generality. Sticking to the question is enforced later, when I notice I've wandered too far and have to backtrack. But I think this is the optimal solution. You don't want the hunt for novelty and

generality to be constrained in the moment. Go with it and see what you get.

[8]

Since the initial question does constrain you, in the best case it sets an upper bound on the quality of essay you'll write. If you do as well as you possibly can on the chain of thoughts that follow from the initial question, the initial question itself is the only place where there's room for variation.

It would be a mistake to let this make you too conservative though, because you can't predict where a question will lead. Not if you're doing things right, because doing things right means making discoveries, and by definition you can't predict those. So the way to respond to this situation is not to be cautious about which initial question you choose, but to write a lot of essays. Essays are for taking risks.

Almost any question can get you a good essay. Indeed, it took some effort to think of a sufficiently unpromising topic in the third paragraph, because any essayist's first impulse on hearing that the best essay couldn't be about x would be to try to write it. But if most questions yield good essays, only some yield great ones.

Can we predict which questions will yield great essays? Considering how long I've been writing essays, it's alarming how novel that question feels.

One thing I like in an initial question is outrageousness. I love questions that seem naughty in some way — for example, by seeming counterintuitive or overambitious or heterodox. Ideally all three. This essay is an example. Writing about the best essay implies there is such a thing, which pseudo-intellectuals will dismiss as reductive, though it follows necessarily from the possibility of one essay being better than another. And thinking about how to do something so ambitious is close enough to doing it that it holds your attention.

I like to start an essay with a gleam in my eye. This could be just a taste of mine, but there's one aspect of it that probably isn't: to write a really good essay on some topic, you have to be interested

in it. A good writer can write well about anything, but to stretch for the novel insights that are the *raison d'etre* of the essay, you have to care.

If caring about it is one of the criteria for a good initial question, then the optimal question varies from person to person. It also means you're more likely to write great essays if you care about a lot of different things. The more curious you are, the greater the probable overlap between the set of things you're curious about and the set of topics that yield great essays.

What other qualities would a great initial question have? It's probably good if it has implications in a lot of different areas. And I find it's a good sign if it's one that people think has already been thoroughly explored. But the truth is that I've barely thought about how to choose initial questions, because I rarely do it. I rarely choose what to write about; I just start thinking about something, and sometimes it turns into an essay.

Am I going to stop writing essays about whatever I happen to be thinking about and instead start working my way through some systematically generated list of topics? That doesn't sound like much fun. And yet I want to write good essays, and if the initial question matters, I should care about it.

Perhaps the answer is to go one step earlier: to write about whatever pops into your head, but try to ensure that what pops into your head is good. Indeed, now that I think about it, this has to be the answer, because a mere list of topics wouldn't be any use if you didn't have edge with any of them. To start writing an essay, you need a topic plus some initial insight about it, and you can't generate those systematically. If only.

[9]

You can probably cause yourself to have more of them, though. The quality of the ideas that come out of your head depends on what goes in, and you can improve that in two dimensions, breadth and depth.

You can't learn everything, so getting breadth implies learning about topics that are very different from one another. When I tell

people about my book-buying trips to Hay and they ask what I buy books about, I usually feel a bit sheepish answering, because the topics seem like a laundry list of unrelated subjects. But perhaps that's actually optimal in this business.

You can also get ideas by talking to people, by doing and building things, and by going places and seeing things. I don't think it's important to talk to new people so much as the sort of people who make you have new ideas. I get more new ideas after talking for an afternoon with Robert Morris than from talking to 20 new smart people. I know because that's what a block of office hours at Y Combinator consists of.

While breadth comes from reading and talking and seeing, depth comes from doing. The way to really learn about some domain is to have to solve problems in it. Though this could take the form of writing, I suspect that to be a good essayist you also have to do, or have done, some other kind of work. That may not be true for most other fields, but essay writing is different. You could spend half your time working on something else and be net ahead, so long as it was hard.

I'm not proposing that as a recipe so much as an encouragement to those already doing it. If you've spent all your life so far working on other things, you're already halfway there. Though of course to be good at writing you have to like it, and if you like writing you'd probably have spent at least some time doing it.

Everything I've said about initial questions applies also to the questions you encounter in writing the essay. They're the same thing; every subtree of an essay is usually a shorter essay, just as every subtree of a Calder mobile is a smaller mobile. So any technique that gets you good initial questions also gets you good whole essays.

At some point the cycle of question and response reaches what feels like a natural end. Which is a little suspicious; shouldn't every answer suggest more questions? I think what happens is that you start to feel sated. Once you've covered enough interesting ground, you start to lose your appetite for new questions. Which is just

as well, because the reader is probably feeling sated too. And it's not lazy to stop asking questions, because you could instead be asking the initial question of a new essay.

That's the ultimate source of drag on the connectedness of ideas: the discoveries you make along the way. If you discover enough starting from question A, you'll never make it to question B. Though if you keep writing essays you'll gradually fix this problem by burning off such discoveries. So bizarrely enough, writing lots of essays makes it as if the space of ideas were more highly connected.

When a subtree comes to an end, you can do one of two things. You can either stop, or pull the Cubist trick of laying separate subtrees end to end by returning to a question you skipped earlier. Usually it requires some sleight of hand to make the essay flow continuously at this point, but not this time. This time I actually need an example of the phenomenon. For example, we discovered earlier that the best possible essay wouldn't usually be timeless in the way the best painting would. This seems surprising enough to be worth investigating further.

There are two senses in which an essay can be timeless: to be about a matter of permanent importance, and always to have the same effect on readers. With art these two senses blend together. Art that looked beautiful to the ancient Greeks still looks beautiful to us. But with essays the two senses diverge, because essays teach, and you can't teach people something they already know. Natural selection is certainly a matter of permanent importance, but an essay explaining it couldn't have the same effect on us that it would have had on Darwin's contemporaries, precisely because his ideas were so successful that everyone already knows about them.

[10]

I imagined when I started writing this that the best possible essay would be timeless in the stricter, evergreen sense: that it would contain some deep, timeless wisdom that would appeal equally to Aristotle and Feynman. That doesn't seem to be true. But if the best possible essay wouldn't usually be timeless in this stricter sense, what would it take to write essays that were?

The answer to that turns out to be very strange: to be the evergreen kind of timeless, an essay has to be ineffective, in the sense that its discoveries aren't assimilated into our shared culture. Otherwise there will be nothing new in it for the second generation of readers. If you want to surprise readers not just now but in the future as well, you have to write essays that won't stick — essays that, no matter how good they are, won't become part of what people in the future learn before they read them.

[11]

I can imagine several ways to do that. One would be to write about things people never learn. For example, it's a long-established pattern for ambitious people to chase after various types of prizes, and only later, perhaps too late, to realize that some of them weren't worth as much as they thought. If you write about that, you can be confident of a conveyor belt of future readers to be surprised by it.

Ditto if you write about the tendency of the inexperienced to overdo things — of young engineers to produce overcomplicated solutions, for example. There are some kinds of mistakes people never learn to avoid except by making them. Any of those should be a timeless topic.

Sometimes when we're slow to grasp things it's not just because we're obtuse or in denial but because we've been deliberately lied to. There are a lot of things adults lie to kids about, and when you reach adulthood, they don't take you aside and hand you a list of them. They don't remember which lies they told you, and most were implicit anyway. So contradicting such lies will be a source of surprises for as long as adults keep telling them.

Sometimes it's systems that lie to you. For example, the educational systems in most countries train you to win by hacking the test. But that's not how you win at the most important real-world tests, and after decades of training, this is hard for new arrivals in the real world to grasp. Helping them overcome such institutional lies will work as long as the institutions remain broken.

[12]

Another recipe for timelessness is to write about things readers already know, but in much more detail than can be transmitted culturally. "Everyone knows," for example, that it can be rewarding to have kids. But till you have them you don't know precisely what forms that takes, and even then much of what you know you may never have put into words.

I've written about all these kinds of topics. But I didn't do it in a deliberate attempt to write essays that were timeless in the stricter sense. And indeed, the fact that this depends on one's ideas not sticking suggests that it's not worth making a deliberate attempt to. You should write about topics of timeless importance, yes, but if you do such a good job that your conclusions stick and future generations find your essay obvious instead of novel, so much the better. You've crossed into Darwin territory.

Writing about topics of timeless importance is an instance of something even more general, though: breadth of applicability. And there are more kinds of breadth than chronological — applying to lots of different fields, for example. So breadth is the ultimate aim.

I already aim for it. Breadth and novelty are the two things I'm always chasing. But I'm glad I understand where timelessness fits.

I understand better where a lot of things fit now. This essay has been a kind of tour of essay writing. I started out hoping to get advice about topics; if you assume good writing, the only thing left to differentiate the best essay is its topic. And I did get advice about topics: discover natural selection. Yeah, that would be nice. But when you step back and ask what's the best you can do short of making some great discovery like that, the answer turns out to be about procedure. Ultimately the quality of an essay is a function of the ideas discovered in it, and the way you get them is by casting a wide net for questions and then being very exacting with the answers.

The most striking feature of this map of essay writing are the

alternating stripes of inspiration and effort required. The questions depend on inspiration, but the answers can be got by sheer persistence. You don't have to get an answer right the first time, but there's no excuse for not getting it right eventually, because you can keep rewriting till you do. And this is not just a theoretical possibility. It's a pretty accurate description of the way I work. I'm rewriting as we speak.

But although I wish I could say that writing great essays depends mostly on effort, in the limit case it's inspiration that makes the difference. In the limit case, the questions are the harder thing to get. That pool has no bottom.

How to get more questions? That is the most important question of all.

## Notes

[1]

There might be some resistance to this conclusion on the grounds that some of these discoveries could only be understood by a small number of readers. But you get into all sorts of difficulties if you want to disqualify essays on this account. How do you decide where the cutoff should be? If a virus kills off everyone except a handful of people sequestered at Los Alamos, could an essay that had been disqualified now be eligible? Etc.

Darwin's 1844 essay was derived from an earlier version written in 1839. Extracts from it were published in 1858.

[2]

When you find yourself very curious about an apparently minor question, that's an exciting sign. Evolution has designed you to pay attention to things that matter. So when you're very curious about something random, that could mean you've unconsciously noticed it's less random than it seems.

[3]

Corollary: If you're not intellectually honest, your writing won't just be biased, but also boring, because you'll miss all the ideas you'd have discovered if you pushed for the truth.

[4]

Sometimes this process begins before you start writing. Sometimes you've already figured out the first few things you want to say. Schoolchildren are often taught they should decide everything they want to say, and write this down as an outline before they start writing the essay itself. Maybe that's a good way to get them started — or not, I don't know — but it's antithetical to the spirit of essay writing. The more detailed your outline, the less your ideas can benefit from the sort of discovery that essays are for.

[5]

The problem with this type of "greedy" algorithm is that you can end up on a local maximum. If the most valuable question is preceded by a boring one, you'll overlook it. But I can't imagine a better strategy. There's no lookahead except by writing. So use a greedy algorithm and a lot of time.

[6]

I ended up reattaching the first 5 of the 17 paragraphs, and discarding the rest.

[7]

Stephen Fry confessed to making use of this phenomenon when taking exams at Oxford. He had in his head a standard essay about some general literary topic, and he would find a way to turn the exam question toward it and then just reproduce it again.

Strictly speaking it's the graph of ideas that would be highly

connected, not the space, but that usage would confuse people who don't know graph theory, whereas people who do know it will get what I mean if I say "space".

[8]

Too far doesn't depend just on the distance from the original topic. It's more like that distance divided by the value of whatever I've discovered in the subtree.

[9]

Or can you? I should try writing about this. Even if the chance of succeeding is small, the expected value is huge.

[10]

There was a vogue in the 20th century for saying that the purpose of art was also to teach. Some artists tried to justify their work by explaining that their goal was not to produce something good, but to challenge our preconceptions about art. And to be fair, art can teach somewhat. The ancient Greeks' naturalistic sculptures represented a new idea, and must have been extra exciting to contemporaries on that account. But they still look good to us.

[11]

Bertrand Russell caused huge controversy in the early 20th century with his ideas about "trial marriage." But they make boring reading now, because they prevailed. "Trial marriage" is what we call "dating."

[12]

If you'd asked me 10 years ago, I'd have predicted that schools would continue to teach hacking the test for centuries. But now it seems plausible that students will soon be taught individually by AIs, and that exams will be replaced by ongoing, invisible micro-assessments.

Thanks to Sam Altman, Trevor Blackwell,  
Jessica Livingston, Robert  
Morris, Courtenay Pipkin, and Harj Taggar for reading drafts of

this.

|  
|