

How to Write Usefully

paulgraham.com · Paul Graham · 2020-02 · [source](#)

|

February 2020

What should an essay be? Many people would say persuasive. That's what a lot of us were taught essays should be. But I think we can aim for something more ambitious: that an essay should be useful.

To start with, that means it should be correct. But it's not enough merely to be correct. It's easy to make a statement correct by making it vague. That's a common flaw in academic writing, for example. If you know nothing at all about an issue, you can't go wrong by saying that the issue is a complex one, that there are many factors to be considered, that it's a mistake to take too simplistic a view of it, and so on.

Though no doubt correct, such statements tell the reader nothing. Useful writing makes claims that are as strong as they can be made without becoming false.

For example, it's more useful to say that Pike's Peak is near the middle of Colorado than merely somewhere in Colorado. But if I say it's in the exact middle of Colorado, I've now gone too far, because it's a bit east of the middle.

Precision and correctness are like opposing forces. It's easy to satisfy one if you ignore the other. The converse of vaporous academic writing is the bold, but false, rhetoric of demagogues. Useful writing is bold, but true.

It's also two other things: it tells people something important, and that at least some of them didn't already know.

Telling people something they didn't know doesn't always mean surprising them. Sometimes it means telling them something they

knew unconsciously but had never put into words. In fact those may be the more valuable insights, because they tend to be more fundamental.

Let's put them all together. Useful writing tells people something true and important that they didn't already know, and tells them as unequivocally as possible.

Notice these are all a matter of degree. For example, you can't expect an idea to be novel to everyone. Any insight that you have will probably have already been had by at least one of the world's 7 billion people. But it's sufficient if an idea is novel to a lot of readers.

Ditto for correctness, importance, and strength. In effect the four components are like numbers you can multiply together to get a score for usefulness. Which I realize is almost awkwardly reductive, but nonetheless true.

How can you ensure that the things you say are true and novel and important? Believe it or not, there is a trick for doing this. I learned it from my friend Robert Morris, who has a horror of saying anything dumb. His trick is not to say anything unless he's sure it's worth hearing. This makes it hard to get opinions out of him, but when you do, they're usually right.

Translated into essay writing, what this means is that if you write a bad sentence, you don't publish it. You delete it and try again. Often you abandon whole branches of four or five paragraphs. Sometimes a whole essay.

You can't ensure that every idea you have is good, but you can ensure that every one you publish is, by simply not publishing the ones that aren't.

In the sciences, this is called publication bias, and is considered bad. When some hypothesis you're exploring gets inconclusive results, you're supposed to tell people about that too. But with essay

writing, publication bias is the way to go.

My strategy is loose, then tight. I write the first draft of an essay fast, trying out all kinds of ideas. Then I spend days rewriting it very carefully.

I've never tried to count how many times I proofread essays, but I'm sure there are sentences I've read 100 times before publishing them. When I proofread an essay, there are usually passages that stick out in an annoying way, sometimes because they're clumsily written, and sometimes because I'm not sure they're true. The annoyance starts out unconscious, but after the tenth reading or so I'm saying "Ugh, that part" each time I hit it. They become like briars that catch your sleeve as you walk past. Usually I won't publish an essay till they're all gone till I can read through the whole thing without the feeling of anything catching.

I'll sometimes let through a sentence that seems clumsy, if I can't think of a way to rephrase it, but I will never knowingly let through one that doesn't seem correct. You never have to. If a sentence doesn't seem right, all you have to do is ask why it doesn't, and you've usually got the replacement right there in your head.

This is where essayists have an advantage over journalists. You don't have a deadline. You can work for as long on an essay as you need to get it right. You don't have to publish the essay at all, if you can't get it right. Mistakes seem to lose courage in the face of an enemy with unlimited resources. Or that's what it feels like. What's really going on is that you have different expectations for yourself. You're like a parent saying to a child "we can sit here all night till you eat your vegetables." Except you're the child too.

I'm not saying no mistake gets through. For example, I added condition (c) in "A Way to Detect Bias" after readers pointed out that I'd omitted it. But in practice you can catch nearly all of them.

There's a trick for getting importance too. It's like the trick I suggest to young founders for getting startup ideas: to make something

you yourself want. You can use yourself as a proxy for the reader. The reader is not completely unlike you, so if you write about topics that seem important to you, they'll probably seem important to a significant number of readers as well.

Importance has two factors. It's the number of people something matters to, times how much it matters to them. Which means of course that it's not a rectangle, but a sort of ragged comb, like a Riemann sum.

The way to get novelty is to write about topics you've thought about a lot. Then you can use yourself as a proxy for the reader in this department too. Anything you notice that surprises you, who've thought about the topic a lot, will probably also surprise a significant number of readers. And here, as with correctness and importance, you can use the Morris technique to ensure that you will. If you don't learn anything from writing an essay, don't publish it.

You need humility to measure novelty, because acknowledging the novelty of an idea means acknowledging your previous ignorance of it. Confidence and humility are often seen as opposites, but in this case, as in many others, confidence helps you to be humble. If you know you're an expert on some topic, you can freely admit when you learn something you didn't know, because you can be confident that most other people wouldn't know it either.

The fourth component of useful writing, strength, comes from two things: thinking well, and the skillful use of qualification. These two counterbalance each other, like the accelerator and clutch in a car with a manual transmission. As you try to refine the expression of an idea, you adjust the qualification accordingly. Something you're sure of, you can state baldly with no qualification at all, as I did the four components of useful writing. Whereas points that seem dubious have to be held at arm's length with perhapses.

As you refine an idea, you're pushing in the direction of less qualification. But you can rarely get it down to zero. Sometimes you don't even want to, if it's a side point and a fully refined version would be too long.

Some say that qualifications weaken writing. For example, that you should never begin a sentence in an essay with "I think," because if you're saying it, then of course you think it. And it's true that "I think x" is a weaker statement than simply "x." Which is exactly why you need "I think." You need it to express your degree of certainty.

But qualifications are not scalars. They're not just experimental error. There must be 50 things they can express: how broadly something applies, how you know it, how happy you are it's so, even how it could be falsified. I'm not going to try to explore the structure of qualification here. It's probably more complex than the whole topic of writing usefully. Instead I'll just give you a practical tip: Don't underestimate qualification. It's an important skill in its own right, not just a sort of tax you have to pay in order to avoid saying things that are false. So learn and use its full range. It may not be fully half of having good ideas, but it's part of having them.

There's one other quality I aim for in essays: to say things as simply as possible. But I don't think this is a component of usefulness. It's more a matter of consideration for the reader. And it's a practical aid in getting things right; a mistake is more obvious when expressed in simple language. But I'll admit that the main reason I write simply is not for the reader's sake or because it helps get things right, but because it bothers me to use more or fancier words than I need to. It seems inelegant, like a program that's too long.

I realize florid writing works for some people. But unless you're sure you're one of them, the best advice is to write as simply as you can.

I believe the formula I've given you, importance + novelty + correctness + strength, is the recipe for a good essay. But I should warn you that it's also a recipe for making people mad.

The root of the problem is novelty. When you tell people something they didn't know, they don't always thank you for it. Sometimes the reason people don't know something is because they don't want to know it. Usually because it contradicts some cherished belief. And indeed, if you're looking for novel ideas, popular but mistaken beliefs are a good place to find them. Every popular mistaken belief creates a dead zone of ideas around it that are relatively unexplored because they contradict it.

The strength component just makes things worse. If there's anything that annoys people more than having their cherished assumptions contradicted, it's having them flatly contradicted.

Plus if you've used the Morris technique, your writing will seem quite confident. Perhaps offensively confident, to people who disagree with you. The reason you'll seem confident is that you are confident: you've cheated, by only publishing the things you're sure of. It will seem to people who try to disagree with you that you never admit you're wrong. In fact you constantly admit you're wrong. You just do it before publishing instead of after.

And if your writing is as simple as possible, that just makes things worse. Brevity is the diction of command. If you watch someone delivering unwelcome news from a position of inferiority, you'll notice they tend to use lots of words, to soften the blow. Whereas to be short with someone is more or less to be rude to them.

It can sometimes work to deliberately phrase statements more weakly than you mean. To put "perhaps" in front of something you're actually quite sure of. But you'll notice that when writers do this, they usually do it with a wink.

I don't like to do this too much. It's cheesy to adopt an ironic tone for a whole essay. I think we just have to face the fact that elegance and curtness are two names for the same thing.

You might think that if you work sufficiently hard to ensure that an essay is correct, it will be invulnerable to attack. That's sort of true. It will be invulnerable to valid attacks. But in practice that's little consolation.

In fact, the strength component of useful writing will make you particularly vulnerable to misrepresentation. If you've stated an idea as strongly as you could without making it false, all anyone has to do is to exaggerate slightly what you said, and now it is false.

Much of the time they're not even doing it deliberately. One of the most surprising things you'll discover, if you start writing essays, is that people who disagree with you rarely disagree with what you've actually written. Instead they make up something you said and disagree with that.

For what it's worth, the countermove is to ask someone who does this to quote a specific sentence or passage you wrote that they believe is false, and explain why. I say "for what it's worth" because they never do. So although it might seem that this could get a broken discussion back on track, the truth is that it was never on track in the first place.

Should you explicitly forestall likely misinterpretations? Yes, if they're misinterpretations a reasonably smart and well-intentioned person might make. In fact it's sometimes better to say something slightly misleading and then add the correction than to try to get an idea right in one shot. That can be more efficient, and can also model the way such an idea would be discovered.

But I don't think you should explicitly forestall intentional misinterpretations in the body of an essay. An essay is a place to meet honest readers. You don't want to spoil your house by putting bars on the windows to protect against dishonest ones. The place to protect against intentional misinterpretations is in end-notes. But don't think you can predict them all. People are as ingenious at misrepresenting you when you say something they don't want to hear as they are at coming up with rationalizations for things they want to do but know they shouldn't. I suspect it's the same skill.

As with most other things, the way to get better at writing essays

is to practice. But how do you start? Now that we've examined the structure of useful writing, we can rephrase that question more precisely. Which constraint do you relax initially? The answer is, the first component of importance: the number of people who care about what you write.

If you narrow the topic sufficiently, you can probably find something you're an expert on. Write about that to start with. If you only have ten readers who care, that's fine. You're helping them, and you're writing. Later you can expand the breadth of topics you write about.

The other constraint you can relax is a little surprising: publication. Writing essays doesn't have to mean publishing them. That may seem strange now that the trend is to publish every random thought, but it worked for me. I wrote what amounted to essays in notebooks for about 15 years. I never published any of them and never expected to. I wrote them as a way of figuring things out. But when the web came along I'd had a lot of practice.

Incidentally,

Steve

Wozniak did the same thing. In high school he designed computers on paper for fun. He couldn't build them because he couldn't afford the components. But when Intel launched 4K DRAMs in 1975, he was ready.

How many essays are there left to write though? The answer to that question is probably the most exciting thing I've learned about essay writing. Nearly all of them are left to write.

Although the essay is an old form, it hasn't been assiduously cultivated. In the print era, publication was expensive, and there wasn't enough demand for essays to publish that many. You could publish essays if you were already well known for writing something else, like novels. Or you could write book reviews that you took over to express your own ideas. But there was not really a direct

path to becoming an essayist. Which meant few essays got written, and those that did tended to be about a narrow range of subjects.

Now, thanks to the internet, there's a path. Anyone can publish essays online. You start in obscurity, perhaps, but at least you can start. You don't need anyone's permission.

It sometimes happens that an area of knowledge sits quietly for years, till some change makes it explode. Cryptography did this to number theory. The internet is doing it to the essay.

The exciting thing is not that there's a lot left to write, but that there's a lot left to discover. There's a certain kind of idea that's best discovered by writing essays. If most essays are still unwritten, most such ideas are still undiscovered.

Notes

[1] Put railings on the balconies, but don't put bars on the windows.

[2] Even now I sometimes write essays that are not meant for publication. I wrote several to figure out what Y Combinator should do, and they were really helpful.

Thanks to Trevor Blackwell, Daniel Gackle, Jessica Livingston, and Robert Morris for reading drafts of this.